

vmware®

Carbon Black App Control

[View Details](#)



User Guide

Product Version: 8.5.0

Document Date: August 2020

Copyrights and Notices

Copyright © 2004-2020 VMware, Inc. All rights reserved. Carbon Black is a registered trademark and/or trademark of VMware, Inc. in the United States and other countries. All other trademarks and product names may be the trademarks of their respective owners.

THERE IS NO WARRANTY FOR THE PROGRAM, TO THE EXTENT PERMITTED BY APPLICABLE LAW EXCEPT WHEN OTHERWISE STATED IN WRITING. THE COPYRIGHT HOLDERS AND/OR OTHER PARTIES PROVIDE THE PROGRAM "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. THE ENTIRE RISK AS TO THE QUALITY AND PERFORMANCE OF THE PROGRAM IS WITH YOU. SHOULD THE PROGRAM PROVE DEFECTIVE, YOU ASSUME THE COST OF ALL NECESSARY SERVICING, REPAIR OR CORRECTION.

We acknowledge the use of the following third-party software in the VMware Carbon Black App Control product:

Portions of this software created by gSOAP are Copyright © 2001-2004 Robert A. van Engelen, Genivia inc. All Rights Reserved. See **Note 1** below for additional details.

This product includes PHP, freely available from <http://www.php.net>. Copyright © 1999 - 2015 The PHP Group, All rights reserved. See **Note 1** below for additional details.

This product includes third-party software licensed under the MIT License (MIT) Copyright (c) <2020> See **Note 1** below for additional details.

Portions of this software use RadControls for WinForms, Copyright © 2010-2014, Telerik Corporation. All Rights Reserved. Warning: This computer program is protected by copyright law and international treaties. Unauthorized reproduction or distribution of this program, or any portion of it, may result in severe civil and criminal penalties, and will be prosecuted to the maximum extent possible under the law.

This product contains Smarty and SimpleSAMLphp, which are copyrighted software licensed under the Lesser General Public License v3. Copies of the GPL and LGPL licenses can be found at <http://www.gnu.org/licenses/gpl-3.0.html> and <http://www.gnu.org/copyleft/lesser.html>. You may obtain the Minimal Corresponding Source code from us for a period of three years after our last shipment of this product, which will be no earlier than 2016-01-30 by writing to GPL Compliance Division, VMware Carbon Black, 1100 Winter Street, Waltham, MA 02451.

NOTE 1

SOFTWARE FROM THE FOLLOWING ORGANIZATIONS OR INDIVIDUALS IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT OWNER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE. THIS STATEMENT APPLIES TO:

- GENIVIA INC
- PHP DEVELOPMENT TEAM
- Copyright (c) 2010 Terence Parr
- Copyright jQuery Foundation and other contributors
- Copyright (c) 2012-2020, CodePlex Foundation
- Copyright (c) 2004-2020 Jaroslaw Kowalski <jaak@jkowalski.net>, Kim Christensen, Julian Verdurmen
- Copyright 2015 Ben Plum
- Copyright (c) 2007-2015 Ariel Flesler <aflesler@gmail.com>
- Copyright (c) 2010 Kelvin Luck
- Copyright © 2013 - present by Luigi Berrettini and others:
<https://github.com/luigiberrettini/NLog.Targets.Syslog/graphs/contributors>
- Copyright (c) 2009 Eduardo Lundgren (edu@rdo.io) and Richard D. Worth (rdworth@gmail.com)
- Copyright (c) 2014 Christian Bach
- Copyright (c) 2007-2016. The YARA Authors. All Rights Reserved.
- Font data copyright Google 2012
- Copyright © 2005-2008 Thomas Fuchs (<http://script.aculo.us>, <http://mir.aculo.us>)
- Prototype is Copyright © 2005-2007 Sam Stephenson. It is freely distributable under the terms of an MIT-style license.
- Copyright (c) 2011-2020 The Bootstrap Authors
- HoverIntent.js
- jQuery.selectbox

- MochiKit
- NewtonSoft JSON Copyright (c) 2007 James Newton-King

VMware Carbon Black

1100 Winter Street, Waltham, MA 02451 USA

Tel: 617.393.7400

Fax: 617.393.7499

Web Site: <http://www.carbonblack.com>

Support E-mail: support@carbonblack.com

User Exchange (Carbon Black Community): <https://community.carbonblack.com>

Before You Begin

This preface provides a brief introduction to this user guide.

Sections

Topic	Page
Intended Audience	5
Terminology	5
What this Documentation Covers	7
Community Resources	11
Contacting Support	11

Intended Audience

This documentation provides information for administrators, incident responders, and others who will operate the VMware Carbon Black App Control Console. Staff who manage Carbon Black App Control activities should be familiar with the Microsoft Windows operating system, web applications, desktop infrastructure (especially in-house procedures for software roll-outs, patch management, and anti-virus software maintenance), and the effects of unwanted software. In addition, if you intend to use features that integrate App Control and Active Directory, you should be familiar with Active Directory concepts and use. Although not necessary for day-to-day users, knowledge of SQL Server management is required for the administrator of the App Control database server at your site.

App Control administrators should also be familiar with the operating systems of clients managed by the App Control server, as well as the software installed on them.

Terminology

The following table defines some of the key terms you will need to understand App Control and its features:

Term	Definition
App Control Server	Computer running the App Control Server software on a supported Windows platform.
App Control Agent	Agent software installed on computers on your network; the agent runs independently but reports to the App Control Server.
App Control Console	The console, which can be displayed remotely with a web browser, is the user interface and management center for all App Control Server management activities.
Enforcement Level	The protection level applied to computers running the App Control Agent. A range of levels from High (Block Unapproved) to None (Disabled) enable you to specify the level of file blocking required.
Computer	Computer that runs the App Control Agent. Each App Control-managed computer is protected by the agent, which both provides information and receives protection updates when it is connected to the App Control Server. The App Control Agent can be installed on both physical and virtual machines.
Template	Computer that has the App Control Agent pre-installed and will be used to clone one or more computers.
Policy	Each computer protected by App Control is assigned a policy that defines its security characteristics. Computers with the same security requirements can share the same policy.
Computer Initialization	File inventory initialization process for new computers that come online to the App Control Server. During initialization, each file on the fixed drives of the new machine is evaluated and classified by the server.

Term	Definition
Login Account	<p>To use the App Control Console, users must have a login account. Role-based accounts tailored to users' responsibilities determine what they can do on the system.</p> <p>Note that users of computers running the App Control Agent do not need console accounts. The server requires no direct interaction with users of computers it monitors and protects.</p>
Executables and Scripts	<p>An executable is any file that contains executable code. App Control examines the <i>content</i> of each unknown file that appears on a computer in its network, determines whether it contains executable code, and, if so, categorizes it according to executable type.</p> <p>App Control also has rules that identify and manage <i>scripts</i>, and you can define additional rules for script identification.</p> <p>The App Control Server keeps an inventory of executables and scripts, and provides rules that control whether they are allowed to run. The files that are inventoried are sometimes referred to as "interesting files." Files not identified as executables or scripts are not inventoried, although you might be able to control access to them with custom rules, such as <i>file integrity rules</i>.</p>
File State	<p>The App Control classification that determines how executables are tracked and permitted or not permitted to be run. Top-level file states includes approved, banned, and unapproved (neither approved nor banned) states. Files have global and local file states, and these may vary in some cases.</p>
Software Approval	<p>App Control features for approving legitimate software. Approved software is allowed to run without user or administrator intervention, even on computers "locked down" under high protection.</p>
Reputation	<p>Information that provides guidance about whether a file should be approved or banned. Carbon Black File Reputation, when integrated with the App Control Server, provides reputation data for a large database of files and file publishers.</p>
Notifier	<p>A dialog box or transient panel that can appear when a App Control rule blocks an action. Notifiers may contain information about why the action was blocked, and in some cases give the user the option of allowing the action or requesting approval from an administrator. Notifiers are configured and saved by name, and can be attached to different rules.</p>
Approval Request	<p>A request by a user whose action was blocked for access to a file or device. Requests can be handled informally through email or websites outside of App Control, or using the approval request management feature in notifiers and the App Control Console.</p>
Drift Report	<p>A report that can help determine how far one or more computers have "drifted" from a baseline of files (by having files added, removed or changed). This can help determine level of compliance with company policies on acceptable files, and also identify files that should be approved and added to an updated baseline.</p>

Term	Definition
Live Inventory	App Control's near-real-time database of all files of interest on all local drives on all computers running the App Control Agent (removable and remote drives are not tracked).
Baseline and Snapshot	A reference point that can be used to determine drift of computers' file inventory from the reference, which might indicate potential risk for those computers. A baseline can be a named table of files, called a Snapshot, or it can be the current set of files on a reference computer.
Indicator Set	Groups of rules called "indicators" that aid in detecting particularly threatening or suspicious activity on systems reporting to your server.
Health Indicator	A rule that checks whether certain parameters on the App Control Server and SQL Server meet the operating requirements and reports its results to the System Health page.
Event	Records of actions related to App Control activities, including files blocked, unapproved files executed, system management processes and actions by console users. Events may be examined in the console and exported to other analytical tools such as Syslog servers or data analysis systems.
Event Rule	A rule that takes a particular action when a specified event is recorded on the App Control Server. Actions include changing file states, uploading files from endpoints, and sending files to third-party detonation engines.
Unified Management	In an organization with multiple App Control Servers, Unified Management allows one server to control many common management functions for itself and any of the other connected App Control Servers.

What this Documentation Covers

This is your guide to day-to-day administration and security monitoring tasks: monitoring executable files on your network using App Control; configuring the App Control Server; managing computers running the App Control Agent; and managing App Control Console users. It covers the following:

Chapter	Description
1 Overview	Describes the App Control architecture, key management concepts, and operation strategies.
2 Using the Console	Describes how to log in to the system and navigate to features using the App Control Console. It includes descriptions of common menus and buttons.
3 Managing Console Login Accounts	Describes how to create, manage, and delete login accounts. Also describes the role-based access privileges of different types of user accounts, and how to use Active Directory accounts as App Control Console accounts.
4 Creating and Configuring Policies	Describes policies, which define the protections for groups of computers; includes policy settings, Enforcement Levels, and how to change them.
5 Managing Computers	Describes how to configure, deploy, and install the App Control Agent. Also describes how to get information about managed computers.
6 Managing Virtual Machines	Describes special considerations for managing virtual machines created from template computers.
7 File, Publisher, and Application Information	Describes where and how you get information about files seen by agents reporting to your App Control Server. Includes descriptions of the detailed global and local file state information provided by the server. Also describes information provided about publishers and applications discovered and inventoried by App Control.
8 Approving and Banning Software	Describes different methods of approving and banning files, and when to use them.
9 Deleting Files	Describes how to use the App Control console to delete files from one or more endpoints.
10 Reputation Approval Rules	Describes how to use Carbon Black File Reputation trust settings to automatically approve files and publishers.
11 Managing File-Signing Certificates	Describes how approve and ban files by approving or banning specific certificates associated with a publisher.
12 Managing Devices	Describes how to set up rules to control access to files on devices connected to computers.
13 Custom Software Rules	Describes how to create “custom rules” that affect what happens when there is an attempt to execute or write files at specified paths. Also describes how to export rules from one server and import them to another.
14 Script Rules	Describes how to add files to the list of those controlled by script rules.

Chapter	Description
15 Registry Rules	Describes how to create rules that affect what happens when there is an attempt to modify the Windows Registry at specified paths.
16 Memory Rules	Describes how to create rules that affect what happens when there is an attempt by one process to access or alter another process.
17 Expert Rules	Describes the expert interface to Custom, Registry, and Memory Rules. This interface is for use in consultation with Carbon Black Support or Services only.
18 Rapid Configs	Describes how to enable and configure sets of rules that can be used to accomplish tasks such as application optimization, operating system and application hardening, and approval of files delivered by software distribution systems.
19 Event Rules	Describes how to create rules that take an action when specified events are reported to the App Control Server.
20 Endpoint Notifiers and Approval Requests	Describes how blocked file notifiers work on agent computers and describes how to customize notifiers. Also describes configuration and management of approval requests from users.
21 Events, Alerts and Meters	Describes how to carry out day-to-day monitoring operations. Instructions include how to use App Control reports and events to identify changes in network file activity and respond appropriately. Also describes how to set up email alerts for App Control-monitored activity, and how to meter execution of specific files.
22 Monitoring Change: Baseline Drift Reports	Describes how to use the Baseline Drift Report feature to monitor change in file inventory over time.
23 Advanced Threat Detection	Describes App Control's advanced threat indicators, which can be used to detect threatening or suspicious activity on systems reporting to your server.
24 Using and Customizing Dashboards	Describes special graphic display pages, called Dashboards, that summarize key information about managed computers and the files on them.
25 Locating Files	Describes the Find Files feature, which can locate specific executable files on computers running the agent on your network.
26 System Configuration	Describes configuration settings, including integration with other servers (including VMWare Carbon Black EDR), backup procedures, product update procedures, optional CB Reputation hash-identification services, agent-server communication security, and other configuration options.

Chapter	Description
27 Unified Management of Multiple Servers	Describes the Unified Management features that allow one App Control Server to control many common management functions on multiple servers.
28 Monitoring System Health	Describes the System Health page, which provides information about factors that affect the health of your App Control environment, including compliance with the hardware and software requirements, SQL Server configuration, and other health and performance data.
A Live Inventory SDK: Database Views	Describes the set of available read-only views into the "live inventory" database of files on your managed computers.
B App Control API	Describes the App Control API, a RESTful API that may be used to write code to interact with App Control, either using custom scripts or from other applications, including network security platforms.
C App Control Connector	Describes the optional, separately licensed connector for integrating third-party network security devices (Check Point, Palo Alto Networks) with App Control. It also describes integration of CB Inspection file analysis services.
D Diagnostic Files	Describes how to upload and access agent diagnostic files. Also describes server diagnostic files available through the console.
E Uploading Files from Agents	Describes the optional, separately licensed features for uploading files from agents to the server.
F Exporting Data for External Analysis	Describes the optional, separately licensed features for sending endpoint data collected by the App Control Server to external analysis tools such as Splunk.

Other App Control Documentation

You will need some or all of the following documentation to accomplish tasks not covered in this user guide. They are available on the Carbon Black customer portal.

Some of these documents are updated with every new released build while others are updated only for minor or major version changes:

- **VMware Carbon Black App Control Operating Environment Requirements** – This describes the hardware and software platform requirements for App Control Server, the SQL Server database that stores App Control data, and the App Control Agent.
- **VMware Carbon Black App Control Installation Guide** – This includes instructions for initial installation of the App Control Server and for upgrades of the server from previous releases. Note that installation of *agents* is described in this document.
- **VMware Carbon Black App Control Release Notes** – This document is specific to the version and build of App Control Server you received. It contains information about new features, corrective content, and known issues with the release.

- **VMware Carbon Black App Control Events Guide** – This document provides a detailed inventory of events recorded by the App Control Server and includes instructions for integrating event data with third-party SIEM systems via Syslog.
- **VMware Carbon Black Supported Agent Operating Systems** – The supported operating systems for the current version of the App Control Agent are listed on the Carbon Black User Exchange at <https://community.carbonblack.com/t5/Documentation-Downloads/Supported-Carbon-Black-sensors-and-agents/ta-p/33041>.

Community Resources

The Carbon Black User Exchange website at <https://community.carbonblack.com> provides access to information shared by Carbon Black customers, employees and partners. It includes information and community participation for users of all VMware Carbon Black products including App Control (formerly CB Protecton) and VMWare Carbon Black EDR (formerly CB Response).

When you log in to this resource, you can:

- ask questions and provide answers to other users' questions
- “vote” to bump up the status of product ideas
- download user documentation
- participate in the Carbon Black developer community by posting ideas and solutions or discussing those posted by others
- view the training resources available for Carbon Black products

You must have a login account to access the User Exchange. Contact your Technical Support representative if you need to get an account.

Contacting Support

For your convenience, Carbon Black Technical Support offers several channels for resolving support questions:

Technical Support Contact Options

Carbon Black User Exchange: <https://community.carbonblack.com>

Email: support@carbonblack.com

Phone: 877.248.9098

Fax: 617.393.7499

Reporting Problems

When you call or e-mail technical support, please provide the following information to the support representative:

Required Information	Description
Contact	Your name, company name, telephone number, and email address
Product version	Product name and version number
Hardware configuration	Hardware configuration of the server or computer the product is running on (processor, memory, and RAM)
Document version	For documentation issues, specify the version of the manual you are using. The date and version of the document appear on the cover page, or for longer manuals, after the Copyrights and Notices section of the manual.
Problem	Action causing the problem, error message returned, and any other appropriate output
Problem severity	Critical, serious, minor, or enhancement

Contents

Copyrights and Notices	2
Before You Begin	4
Intended Audience	5
Terminology	5
What this Documentation Covers	7
Other App Control Documentation	10
Community Resources	11
Contacting Support	11
Reporting Problems	11
1 Overview	38
What is App Control?	39
How App Control Works	43
Files Tracked by App Control	44
System & Platform Architecture	44
App Control Server	45
Integrating App Control with Active Directory	46
Unified Management of Multiple Servers	46
App Control Agent	46
Trust Rating from Carbon Black File Reputation	46
File State, Approving and Banning	47
Global State	47
Local State	47
File Approval Methods	48
File Ban Methods	48
Custom Rules	49
Rapid Configs	49
Security Policies and Levels	49
Policy Settings	50
Modes and Enforcement Levels	50
App Control Licensing and Modes	50
Operating Strategies	51
2 Using the Console	52
Console Access	53
Browser Recommendations	53
Logging In	53
Custom Login Banners	55
Logging Out	55
Login, Server, Version and Alert Information	56
The Home Page	57
Using the Main Menu	60
Left Navigation Menu and Breadcrumbs	65
Console Tables	66
Table Data Control Links	67

Table Column Resizing	67
Row Action Buttons	68
Toggle Switches for Enabling and Disabling Rules	68
Checked Row Action Menus	69
Row Selection Scope	70
Row Rank Arrows	71
“Add” Buttons	72
Pages, Tabs and Saved Views	72
Filter Options	73
Show Columns Options	76
Tabs	77
Table Length, Scrolling and Selection	77
Default and Saved Views	78
Exporting App Control Server Data to Files	80
Details Pages and Object Previews	80
Menus on Details Pages	82
Object Previews in Table Data	82
Shortcut Links	83
Preference Settings for Console Users	83
Using Online Help	86
3 Managing Console Login Accounts	88
Login Account Management	89
User Roles and Permissions	89
Upgrades from Previous Versions	91
Enabling Console Access via AD Accounts	92
AD Login Account Format	93
Adding, Deleting, and Changing AD Login Accounts	93
Changing AD User Details Displayed in the Console	94
Creating Login Accounts in the Console	95
Logging In Using SAML	97
Changing Passwords and Other Account Details	98
Deleting Login Accounts	100
Disabling Login Accounts	101
Managing Console User Roles	101
Creating a New User Role	102
User Role Permissions	106
Editing a User Role	110
Mapping AD Groups to Roles	111
AD Role Mapping Summary	112
Creating AD Mapping Rules	112
Disabling a Role	115
Deleting a Role	116
4 Managing Computers	117
Computer Configuration Overview	118
Pre-Installation Activities	118
Installation and Initialization	119
Post-Installation Activities	119

Permissions for Computer Management Features	120
Assigning Computers to a Policy	121
Assigning Policy by Active Directory Mapping	122
AD Policy Mapping Summary	122
Creating AD Mapping Rules	123
Mapping Rule Ranking	128
AD Object Browser Options	128
Computer Registration and AD Mapping	129
Clearing the Server AD Cache	130
Viewing AD Computer Details in the Console	130
Uploading Agent Installers and Rules to the Server	131
Viewing Agent / Rule Versions and Package Generation Status	133
Downloading Agent Installers	134
Installing App Control Agents	136
Preparing for New Agent Installation	136
Considerations When Installing a App Control Agent using GPO	137
Installing the Agent on a Windows Computer	138
Conditions Requiring Reboot after Installation	139
Command Line Installations of Windows Agents	139
Installing the Agent on a Mac Computer	142
Allowing the Agent Kernel Extension (High Sierra or later)	143
Installing the Agent on a Linux Computer	144
Verifying the Installation	146
Verifying Installation on the Agent Computer	146
Upgrading App Control Agents	147
Feature Limitations for Non-Upgraded Agents	147
Enabling Automatic Agent Upgrades	147
Upgrading Immediately from the Console	148
Manually Upgrading Agents	149
Manually Upgrading Windows Agents	149
Manually Upgrading Mac Agents	151
Manually Upgrading Linux Agents	152
Agent Upgrade Status	153
Uninstalling App Control Agents	154
Uninstalling the Agent from a Windows Computer	154
Uninstalling the Agent from a Mac Computer	155
Uninstalling the Agent from a Linux Computer	155
Viewing the Table of Computers	156
Agent Policy Status	157
Actions on Selected Computers	158
Viewing Complete Details for One Computer	159
Moving Computers to Another Policy	170
Restoring Computers from the Default Policy	171
Moving a Computer to Local Approval Mode	173
Adding Computers	173
Deleting Computers	173
Duplicate Computers	174
Operating System Updates on Agents	175
Operating System Updates on Windows Agents	175

Enabling Trusted Directory Approval of WIM Files	176
Operating System Updates on Mac Agents	177
Upgrading to Mojave with an Agent Installed	178
Operating System Updates on Linux Agents	178
5 Creating and Configuring Policies	179
Policy and Enforcement Level Overview	180
Creating Policies	181
Policy Settings	186
Advanced Settings	187
Rules Affecting Policies	191
Template Policy and Default Policy	192
Default Policy	192
Template Policy	192
Resetting a Policy to Template Policy Settings	193
Tamper-Protection Setting	194
Editing a Policy	194
Related Views in Policy Details	196
Enforcement Levels	197
How Enforcement Levels Affect Policy Setting Enforcement	199
Special Enforcement Level for Local Approval	201
Changing Policy Enforcement Levels	201
Locking Down all Computers	203
Deleting Policies	205
6 Managing Virtual Machines	207
Overview	208
Creating a Template Computer	209
Viewing Templates in the Computers Table	210
Viewing and Editing Template Details	211
Deploying Clones	214
Viewing Clones in the Computers Table	214
Finding the Clones for a Template	214
Finding the Template for a Clone	215
Server Backlog for Clones	215
Making Changes to a Template	216
Deleting a Template	217
Configuring Clone Inventory	217
Choosing an Inventory Option	218
Deleting Clones	219
Manual Cleanup of Clones	220
Automatic Cleanup for All Clones	220
Automatic Clone Cleanup for One Template	221
Converting a Template to a Regular Computer	221
7 File, Publisher, and Application Information	223
Overview	224
Viewing File Tables	225
File Catalog	225

Files on Computers	227
Showing Individual Files	227
Initialized Files	228
Menus on the File Tables Pages	229
Finding Computers With or Without Specified Files	229
Excluding Tracking of Microsoft Support Files	230
Files Instances Affected	232
Changes that Affect OS Inventory Tracking	232
Information about Excluded File Instances	233
File Groups	234
Viewing Details Pages	235
File Details Page	236
File Instance Details Page	242
Menus on the File Details and File Instance Details Pages	245
Summary of File Views	247
Global File State	249
Flags	249
Local File State	250
Local State Details	251
Publisher Information	253
Application Information	256
Viewing Application Data	256
8 Approving and Banning Software	260
What is App Control Software Approval?	261
Platform Considerations for Rule Specifications	263
What are App Control Software Bans?	263
File Ban Options	265
Approving by Updater	266
Updater History	270
Automatic Cloud Management of Updaters	270
Alerts for Updater Changes from the Cloud	271
Approving by Trusted Directory	271
Windows Trusted Directories	272
Installers and Archives in Trusted Directories	272
Mac and Linux Trusted Directories	273
Creating a Trusted Directory	273
Verifying Trusted Directories	276
Tracking Analysis Progress in Trusted Directories	276
Verifying Approval of Windows Packages	278
Custom Rules for Installer Access	278
Removing or Disabling Directory Trust	278
Approving by Trusted User or Group	279
How Groups are Specified	279
Creating a Trusted User or Group	280
Removing Trust from a User or Group	280
Approving or Banning by Publisher	281
Publisher Approvals	281
Publisher Bans	282

Managing Bans and Approvals from the Publishers Tab	282
Managing Bans and Approvals from the Publishers Details Page	284
Adding Publishers	285
Removing Publisher Approvals	285
Removing Publisher Bans	286
Finding All Files from a Publisher	286
Determining Which Certificates Can Approve Files	286
Approval with Expired Certificates	288
Excluding Certificate Algorithms	288
Minimum Key Size	288
Countersignature Options	288
Revocation Checks	289
Locally Approving Files	289
Automatic Local Approval on Enforcement Level Change	290
Which Files Are Locally Approved On Transition	291
Locally Approving Individual Files	292
Removing Local Approval	292
Locally Approving Files Not Yet in File Catalog Inventory	293
Locally Approving Transient or Deleted Files	293
Locally Approving All Unapproved Files on a Computer	294
Moving Computers to Local Approval Mode	294
Moving Online Computers into Local Approval Mode	295
Restoring Online Computers from Local Approval Mode	297
Using Timed Policy Overrides	297
Marking a File as an Installer/Not an Installer	300
File-Specific Rules: Approvals and Bans	301
Report Only Bans	303
Approvals and Bans of MSI Files by Hash	303
Creating an Approval or Ban from the Software Rules Page	304
Editing and Deleting File Rules	306
Creating File Approvals and Bans from Table Pages	307
Creating Global Approvals and Bans	308
Custom Approvals and Bans	308
Warnings when Creating or Editing Bans	309
Approving and Banning Files from the File Details Page	310
Approving or Banning Lists of Files	310
Enabling Bans to Stop Running Processes	312
9 Deleting Files	314
Overview	315
Scope and Limits	315
Permission to Delete Files	315
Timing of File Deletion	316
Files Protected from Deletion	316
Requesting File Deletion from Table Pages	317
Requesting File Deletion from Details Pages	320
Automating File Deletion Requests	321
Monitoring File Deletions	324

10 Reputation Approval Rules	326
Overview	327
Trust Ratings for Files and Publishers	327
File Trust Ratings	327
Publisher Trust Ratings	328
Reputation Approval Strategy	328
Setting the Trust Level for Approvals	329
How File Reputation Approvals Work	329
Removal of Reputation Approval for a File	330
How Publisher Reputation Approvals Work	330
Removal of Reputation Approval for a Publisher	330
Reputation Approvals and Other App Control Rules	331
Creating Exceptions for Files and Publishers	331
Disabling Reputation Approvals for a File	331
Disabling Reputation Approvals for a Publisher	332
Enabling Reputation Approvals	333
Modifying and Disabling Reputation Approvals	335
Views Related to Reputation Approvals	335
11 Managing File-Signing Certificates	337
Overview	338
Summary of Certificate Management Features	339
Viewing Certificate Information	339
Certificates Table	339
Searching, Sorting and Grouping on the Certificates Table	343
Certificate Details	344
Related Views Menu on Certificate Details	345
Viewing Certificates for a Publisher	346
Certificate Fields in File/File Instance Details	347
Certificate Alerts	347
Certificate Events	348
Certificates in External Views	348
Using Certificates for Enforcement	348
Certificate Approval Configuration Choices	349
Certificate Types	350
Path Position and Agent Differences	350
Approving or Banning Certificates for a Publisher	351
Certificate Global State	352
Mixed and By-Policy States	358
Certificate Ban Setting in Policies	358
Interactions with Other Rules	358
How Certificate Global State Affects Global File State	359
Agent Version and Global File State	359
12 Managing Devices	360
Overview	361
Devices Managed by App Control	361
Enabling Per-Policy Device Control	362
Managing Specific Devices	365

Viewing Device Information	365
Managing Devices by Model	366
Viewing Device Models in the Device Catalog	366
Viewing Details for One Device Model	367
Approving and Banning Device Models	369
Managing Device Instances	371
Viewing Instances in the Device Catalog	371
Viewing Details for One Device Instance	372
Approving or Banning Device Instances	374
Managing Computer-Device Attachments	375
Viewing Devices on Computers	376
Viewing Details for One Computer-Device Attachment	377
13 Script Rules	378
Overview	379
What is a Script?	379
What App Control Script Rules Do	380
Pre-configured Script Rules	380
Windows Script Rules Changes on Upgrade	382
Script Rules and Other App Control Rules	382
Shell Scripts Identified by Content	383
Policy Settings for Script Rules	383
Creating a Custom Script Rule	384
Editing a Script Rule	387
Disabling or Deleting a Script Rule	388
Viewing Rule Status on Computers	389
Script Rule Examples	389
Example: Windows Batch Scripts	389
Example: Linux Shell Scripts	390
14 Custom Software Rules	392
Overview	393
Rule Types	393
Rule Scope	394
File and Process Matching	394
Pre-configured Rules	395
Internal Rules in the Custom Rule Table	395
Specifying the Notifier for a Custom Rule	395
Custom Rules in Visibility Mode	396
Creating a Custom Rule	396
Editing a Custom Rule	399
Copying a Custom Rule	400
Custom Rule Fields	401
Specifying Execute and Write Actions	404
Specifying Paths and Processes	407
Specifying a File or Directory	408
Platform-Specific Syntax	408
Using Wildcards in Rules	409
Automatic Path Conversions	409

Specifying Devices in Paths in Windows Rules	409
Using Macros in Rules	410
Common Path Macros	410
Using Windows KNOWNFOLDER GUIDs in Macros	414
OnlyIf Macros	415
Additional Macros	420
Windows Registry Macros	422
Entering Multiple Paths or Processes	423
Specifying Processes	423
Specifying Users or Groups	424
Rule Ranking	424
Rule Ranking and Internal Rules	427
Enabling and Disabling Custom Rules	428
Deleting Custom Rules	430
Viewing Rule Status on Computers	430
Exporting and Importing Rules	431
Exporting Rules	432
Importing Rules	433
Selecting Rules to Import	433
Differences in Settings for Imported Rules	436
Custom Rule Types and Examples	439
File Integrity Control	439
Trusted Paths	441
Execution Control	443
File Creation Control	445
Performance Optimization	446
Pairing Ignore and Block Rules	447
15 Registry Rules	449
Overview	450
Rule Scope	450
Sample Rules	451
Exporting and Importing Registry Rules	451
Specifying the Notifier for Registry Rules	451
Creating Registry Rules	452
Editing a Registry Rule	454
Copying a Registry Rule	454
Registry Rule Fields	455
Specifying a Write Action	458
Specifying Registry Paths	459
Using Wildcards	459
Specifying Keys or Values	460
Specifying Processes in Registry Rules	460
Specifying Processes or Directories	462
Using Wildcards	462
Automatic Process Path Conversions	462
Specifying Devices in Process Path	462
Using Macros	463
Entering Multiple Paths or Processes	463

Specifying Users or Groups	464
Rule Ranking	465
Disabling or Deleting Registry Rules	466
Viewing Rule Status on Computers	466
Sample Registry Rules.	467
Example: Report Changes to Internet Explorer Trusted Zone.	467
Autostart Rules	469
16 Memory Rules	470
Overview	471
Rule Scope	471
Exporting and Importing Memory Rules	472
Specifying the Notifier for Memory Rules	472
Creating Memory Rules	473
Editing a Memory Rule	474
Copying a Memory Rule	475
Memory Rule Fields	476
Specifying the Rule Action	479
Specifying the Rule Permissions	480
Specifying Target and Source Processes	481
Specifying a File or Directory.	481
Using Wildcards.	482
Automatic Path Conversions	482
Specifying Devices in Paths	482
Using Macros.	483
Entering Multiple Target or Source Processes	483
The Source Process Menu	484
Specifying Users or Groups	484
Rule Ranking	485
Disabling or Deleting Memory Rules	486
Viewing Rule Status on Computers	487
17 Expert Rules	488
Overview	489
Enabling the Expert Rules Interface.	489
Expert Rule Definitions.	492
Expert Rule Operations	492
Expert Rule Actions	495
Mutually Exclusive Actions	496
Tags and Tagging Actions in Expert Rules	499
Tag Syntax Requirements.	501
Built-in Tags	502
Tag Persistence	502
Expert Rule Examples	502
Example: Allow Execution in a Folder when Visual Studio is Running	503
Example: Tag a Process and Report its Children	503
Example: Promote an Installer and Demote its Children	504
Switching to or from Expert Rule Mode	504

18 Rapid Configs	505
Overview	506
Rule Scope	506
Viewing Rapid Configs	507
Rapid Config Details	510
Configuring and Enabling Rapid Configs	511
User-Configured Rapid Configs	513
Specifying Notifiers for Rapid Configs	516
Specifying Paths and Processes	517
Automatic Rapid Config Updates	518
Alerts for Rapid Config Changes from the Cloud	519
19 Event Rules	520
Overview	521
Events That Can Trigger Rule Actions	521
Actions A Rule Can Take	521
Simulating the Effect of a Rule	522
Re-Applying a Rule to Past Events	522
Enabling, Disabling, and Deleting Event Rules	522
Disabling Processing of All Event Rules	523
Testing a Rule before Enabling	524
Creating and Editing Event Rules	525
Editing an Event Rule	531
Edit Event Rule Page Menus	531
Event Rule Ranking	532
File and Process Properties in Event Rule Definitions	532
Carbon Black File Reputation Trust and Threat Data	532
File Prevalence	532
File Metadata	532
File Extension	533
Analysis Results Options	533
Global Bans for Non-Cataloged Files	533
How Event Rule Approvals Affect Endpoints	534
Event Rule History and Processed Events List	534
Sample Event Rules	536
Sample Rule: Analyze files from approval requests	536
Sample Rule: Resolve approval requests for clean files	536
Sample Rule: Analyze downloaded files	537
Sample Rule: Report malicious files	537
20 Endpoint Notifiers and Approval Requests	538
Notifiers: What Users See	539
Prompt Notifiers	539
Block-only Notifiers	541
Block Notifiers on Windows Computers	542
Block Notifiers on Mac and Linux Computers	542
Notifier Components	543
Notifier Tray Icon and History Window	543
App Control Notifier History Window	544

The Notifiers Page	545
Assigning Notifiers to Settings and Rules	545
Assigning Notifiers to Policy Settings	545
Policy Settings with Notifiers	546
Assigning Notifiers to Custom, Registry and Memory Rules	547
Assigning Notifiers to Rapid Configs	548
Customizing and Creating Notifiers	548
Related Views on the Edit Notifier Page	551
Creating a New Notifier	552
Editing Notifier Text	552
Using Tags in Notifier Text	552
Conditional Messages for Block vs. Prompt	554
Informational Tags as Conditional Operators	556
Editing the Notifier Link	557
Tags in Notifier Links	557
Editing the Notifier Source Line	559
Specifying a Custom Notifier Logo	560
Image File Requirements	561
Logo-Related Events	562
Changing the Logo Image	562
Suppressing the Notifier Logo in a Policy	562
Resetting a Notifier to Initial Settings	562
Resetting a Policy to Initial Notifiers	562
Disabling App Control Notifiers	562
Notifiers in Windows Session Virtualization	564
Approval Requests and Justifications	566
Enabling Requests and Justifications	566
Submitting Requests and Justifications	568
Managing Requests and Justifications	569
Approval Requests Summary	569
Approval Request Details	570
Reviewing and Resolving Requests and Justifications	572
Request Management Work Flow Shortcuts	576
Opening Rule Details from the Rule Information Panel	578
Managing Duplicate and Related Requests	579
Notifying Users of Approval Request Resolution	580
Approval Request and Justification Details	583
Customizing the Request/Justification Interface in Notifiers	587
21 Events, Alerts and Meters	589
Monitoring Prerequisites	590
Event Reports	590
Using the Home Page Event Reports Portlet	591
Viewing Reports on the Events Page	591
Object Previews in Events Tables	595
Taking Action on Files in Event Reports	596
Customizing Event Reports	596
Using the Event Search Box	596
Editing Event Reports	600

Adding Command Line Information to Event Reports	601
Caching Events for Later Viewing	602
Viewing and Taking Action on the Cached Events Page	604
Removing an Event Cache	605
Viewing Install Event Details	606
Viewing Event Archives	607
Using App Control Alerts	608
Creating Alerts	612
Informational Tags for Event Alert Messages	617
Editing Alerts	618
Alert Priority	618
Deleting Alerts	619
How Alerts are Triggered	619
Mail Notification for Triggered Alerts	620
Reminder Mail for Triggered Alerts	621
Manual and Automatic Alert Resets	621
Viewing Alert Instances and History	623
Managing Alert Email Subscriptions	624
Detecting Agent Issues with Computer Security Alerts	625
Criteria Triggering a Security Alert	626
Alerts for File Prevalence	627
Prevalence Alerts	627
Monitoring Specific File Executions	629
Creating a Meter from the File Details Page	633
22 Monitoring Change: Baseline Drift Reports	634
Baseline Drift Overview	635
How Drift and Risk are Measured	636
Viewing and Managing Baseline Drift Reports	637
Viewing Baseline Drift Report Results	638
Report Results: Computer View	639
Report Results: File Views	639
Drift by Files: Individual or Top-Level	641
Drift by Files: Associated Files Report	642
Drift by Files on a Single Computer	643
Responding to Drift Report Results	643
Adding Drift Results to a Snapshot	644
Creating and Editing Reports	646
Advanced Baseline Drift Report Options	649
Advanced Options: File Filter Options	649
Advanced Options: File Comparison Method	650
Advanced Options: Report Detail Level	651
Using Filters in Target and Baseline Definitions	651
Drift in Multi-Platform Environments	652
Managing Snapshots	653
Creating and Modifying Snapshots	653
Viewing and Editing Snapshots	655
Managing Files in Snapshots	655
Deleting Snapshots	656

Displaying Baseline Drift Reports in Graphs	656
Creating Baseline Drift Alerts	657
23 Advanced Threat Detection	660
Overview	661
Indicator Sets for Threat Detection	662
Indicator Set Details	664
Indicator Set Exceptions	666
Indicator Set Exception Details	668
Updates to Indicator Sets	670
Alerts for Tracking Indicator Set Updates	671
Monitoring Threat Reports	671
Threat Views on the Events Page	672
Fields in Threat-Related Events Views	672
Reviewing Threat Event Reports	673
Showing and Modifying View Parameters	674
Threat Events in Syslog Output	674
Exporting Threat Event Data to CSV Files	675
Threat Views on the Files Pages	675
Threat-Related Alerts	676
Responding to Threats	677
Responding to Threats with Event Rules	678
24 Using and Customizing Dashboards	680
Dashboards Overview	681
Dashboard Elements	683
Using Portlets	683
Getting More Detailed Data	684
Portlet Toolbar Buttons	684
Collapsing, Expanding, and Exploding Portlets	685
Entering Information into Portlets	685
Other Portlet Controls	686
Viewing Other Dashboards	686
Changing Dashboard Appearance	688
Changing Dashboard Layout	689
Portlet Distribution in Layouts	690
Changing Dashboard Width	690
Changing Dashboard Background Color	690
Moving Portlets	691
Creating, Editing and Managing Dashboards	691
Shared Dashboards	693
Creating a New Dashboard	694
Copying a Dashboard	695
Editing a Dashboard	695
Managing the Default Home Page	697
Deleting a Dashboard	697
Managing Dashboards from the Dashboards Page	698
Creating and Customizing Portlets	699
Portlet Types and Subtypes	699

System Portlets	699
Editing Portlet Details	700
Deleting Portlets	700
Creating Custom Portlets	701
Using Tables in Portlets	705
Table-only Portlets	705
Supplemental Tables in Portlets	707
Using Filters in Portlets	708
Nesting Groups of Expressions	711
25 Locating Files	712
Find Files Overview	713
Initiating Find Files from Other Pages	713
Defining a Search on the Find Files Page	714
Finding Files by Name	714
Adding a Pathname to a File Search	716
Finding Files by Hash	716
Using Find Files Results	717
Special Cases in Results	718
Files on Offline Computers	718
Deleted Files	718
Files on Deleted Computers	719
Files on Computers Still Initializing or Synchronizing	720
Saved Views for File Searches	720
26 System Configuration	721
Overview	722
The General Configuration Tab	723
Viewing Server Status and Options	724
Configuring Active Directory Integration	726
Configuring Agent Management Privileges	727
Connection Status and Agent Management Choices	729
Specifying User Password Requirements	730
Event Management Options	732
Managing the App Control Event Database	733
Setting Limits for Event Deletion	733
Enabling Daily Event Archiving	734
Moving the Database to an External Server	734
Setting up External Event Logging	734
Logging Events to a Syslog Server	734
Logging Events to a Supplemental SQL Server	736
Securing Agent-Server Communications	739
Security Status	740
Current Certificate Details	740
Verifying that the Server Name and Certificate Match	742
Importing a Certificate	742
Enabling Certificate Verification	743
Advanced Configuration Options	744
Adding a Login Banner to the Console	749

Backing Up the App Control Server	750
Restoring the App Control Server	753
Configuring Alert and Approval Request Mail	754
Configuring Standard Email for Notifications	756
Configuring Secure Email for Notifications	757
Specifying a Global Alert Subscriber	758
Managing App Control Licenses	759
Viewing Your App Control License Limits and Use	759
License Warnings	761
Adding Licenses	761
Confirming License Addition	762
Activating Carbon Black File Reputation	763
Carbon Black File Reputation Availability Status	765
Deactivating the Carbon Black File Reputation Connection	766
Using a Proxy Server for Carbon Black File Reputation	766
Carbon Black File Reputation Synchronization	767
Activating Carbon Black EDR Integration	768
Creating a Carbon Black EDR User for the Integration	769
Configuring Unified Management	769
Configuring SAML Logins	769
Integrating App Control with an IdP	770
Allowing Non-SAML Logins for Specified User Roles	772
Logging In Using SAML	773
Deleting or Editing an Identity Provider	774
27 Unified Management of Multiple Servers	777
Overview	778
Unified Management Features	778
Configuring Unified Management	780
Enabling Unified Management and Adding Servers	780
Adding Client Servers to Unified Management	781
Server Information on the Unified Management Page	783
Disconnected and Unreachable Servers	784
Authentication Errors	784
Configuration Information on Managed Servers	785
Creating Unified Management Console Accounts	785
Authenticating a User on Client Servers	786
Editing Client Server Configuration	788
Disabling Unified Management	791
Unified Management of Rules	791
Managing Unified Rules from the Software Rules Page	792
Copying Existing Rules to Other Servers	798
Editing a Unified Rule	801
Permissions Needed for Unified Rule Editing	802
Changing a Unified Rule to a Local Rule	803
Disabling and Enabling Unified Rules	803
Managing Unified File Rules from File Table and Details Pages	803

28 Monitoring System Health	806
Overview	807
Enabling System Health Indicators	808
Disabling System Health Indicators	808
Viewing the System Health Page	809
Navigating on the System Health Page	810
Health Indicator State	811
System Health Alerts	812
System Health Events	813
A Live Inventory SDK: Database Views	815
Performance Considerations	815
Upgrading from a Previous Version	816
Schema Overview: bit9_public	817
Specifying a Schema User	817
Schema Views and Diagram	817
Schema Diagram for bit9_public	819
Details of Database Views	821
ExComputers	821
ExInfo	823
ExMeters	824
ExEvents	825
ExFileCatalog	826
ExFileInstances	829
ExDeletedFileInstances	831
ExFileInstanceGroups	832
ExApprovalRequests	833
Sample Queries	836
B App Control API	839
Overview	840
API Authentication and Access Control	840
Available Objects	841
Using the App Control API to Add a Connector	842
C App Control Connector	843
Overview	844
Preparing to use the Connector	845
Enabling Carbon Black Inspection	845
CB Inspection Connector Configuration	846
Enabling Palo Alto Networks Integration	846
Integrating Palo Alto Networks Appliances for Notifications	847
Palo Alto Networks Notification Appliance Status	849
Modifying or Deleting an Appliance Integration	850
Integrating with the WildFire Cloud for Analysis	850
Integrating with the WildFire Public Cloud	851
WildFire Public Cloud Query Limits	851
Integrating with a WildFire Private Cloud Device	852
Enabling VMware Carbon Black Cloud Integration	853
Using the Links to the Carbon Black Cloud	855

Enabling Check Point Integration	857
Integrating Check Point Log Servers with App Control	857
Custom Import Filters for Check Point	861
Check Point Log Server Status	863
Modifying or Deleting a Log Server Integration	864
Integrating with Check Point for File Analysis	864
Connecting to a Threat Emulation Appliance	865
Connecting to the ThreatCloud Emulation Service	866
ThreatCloud Emulation Lookup Limits	866
Enabling Automatic Threat Emulation Lookups	867
Enabling Console Account Permissions	867
External Notifications	867
Action Menu on External Notifications Table Page	871
Saved Views on the Notifications Table Page	872
Notification Table Access from File Details Pages	872
Choosing Correlation Level for External Notifications	872
Notifications from Multiple Analysis Environments	874
External Notification Details	874
Total Files Tab	875
Known Files Tab	876
Files On Computers Tab	877
Registry Keys	877
More Details Tab	878
History Tab	878
Showing Related Notifications	878
Showing XML Details	879
External Console Access	879
Managing Notification Status	879
Banning Externally Reported Malware	880
Manually Banning Files	880
Special Rules for Reporting or Banning Malware	881
Registry Rules	881
Custom Rules for Directory Control	882
Analysis of Suspicious Files on Endpoints	883
Monitoring Files Submitted for Analysis	884
Analysis Status	885
Actions on the Analyzed Files tab	886
Logging of Connector-related Events	886
Additional Log Information	889
D Diagnostic Files	890
Overview	891
Uploading Agent Diagnostic Files	891
Canceling or Retrying an Upload	892
Viewing Diagnostic Files	892
Deleting Uploaded Diagnostic Files	894
E Uploading Files from Agents	895
Overview	896
Enabling Access to File Upload Features	896

Scheduling Uploads	897
Starting Uploads of Inventoried Files from Tables	897
Starting Uploads from the File Instance Details Page	898
Starting Uploads by Path from the Computer Details Page	899
Viewing the Uploads Table	900
Diagnostic Files	902
Downloading Uploaded Files	903
File and Path Information for Uploaded Files	903
Upload Configuration Options	904
Deleting Uploaded Files	904
Changing the Uploaded File Location	904
F Exporting Data for External Analysis	906
Overview	907
Preparing to Use External Analytics	907
Data Format and Management	908
Data Volume for Exported Analytics	909
Limiting Export Directory Size	909
Local vs. Network Log Files	909
Enabling External Analytics in the App Control Console	910
Editing or Disabling the External Analytics Integration	914
Adding a Custom Rule to Ignore Analytics Log Files	914
Enabling an External Tool for Data Analytics	914
Enabling Splunk to Collect App Control Data	915
Configuring the Splunk Server for App Control Access	915
Installing the Splunk Forwarder and App on the App Control Server	916
Viewing App Control Data in External Analytics Tools	917
Linking to an External Tool from the App Control Console	917
Using the Splunk App for App Control	918
Dashboards in the Splunk App for App Control	918
Field Mappings to CIM in the Splunk App for App Control	924
Index	926

List of Tasks

How do I . . .

- To access the full XML details for an External Notification: 879
- To activate VMware Carbon Black Cloud integration: 855
- To add (create) a custom rule: 397
- To add (create) a custom script rule: 384
- To add (create) a memory rule: 473
- To add (create) a new notifier: 552
- To add (create) a registry rule: 452
- To add (create) a unified rule from the Software Rules page: 792
- To add (create) an event rule: 526
- To add a client server to the Unified Management server: 781
- To add a publisher: 285
- To add a subscriber to the email notification list for one alert: 625
- To add an identity provider to App Control: 771
- To add App Control as a service provider for an identity provider: 770
- To add files to a snapshot from a baseline drift report: 645
- To add new App Control licenses by entering the key: 761
- To add new App Control licenses by filename: 762
- To allow the Carbon Black, Inc. kernel extension after agent installation or upgrade: 143
- To allow the kernel extension during agent installation or upgrade: 143
- To allow trusted directory approval of WIM files: 176
- To allow users with selected User Roles to log in locally (without SAML): 772
- To approve one device model from the Device Model Details page: 370
- To approve one or more device instances from the Device Catalog: 374
- To approve one or more device models from the Device Catalog: 369
- To approve or ban a certificate: 351
- To approve or ban a single file using the File Details page: 310
- To approve or ban one device instance (Device or Attachment Details page): 375
- To approve or ban one publisher by policy (Publisher Details page): 284
- To approve or ban software from one or more publishers for all policies: 283
- To assign a notifier to a policy setting: 545
- To authenticate a Unified Management user: 787
- To automatically approve files installed by application updaters: 269
- To cancel diagnostic file uploads: 892
- To change a console password and other login account details: 99
- To change all unapproved files on a computer to Locally Approved: 294
- To change Enforcement Level for a policy in Control mode: 202
- To change permissions or other properties of a console user role: 111
- To change the rank of a memory rule: 485
- To change the rank of a registry rule: 465
- To change the rank of a rule in an unfiltered table sorted by rank: 425
- To change the rank of a rule in any table: 426
- To change the target location for uploaded files: 905
- To check the status of a trusted directory: 276
- To clear the server cache and update AD information: 130
- To configure and enable a policy-specific Rapid Config: 514
- To configure and enable a special login banner: 749
- To configure automatic clone cleanup for a specific template: 221

To configure automatic deletion of uploaded files: 904
To configure email using standard (unsecure) mail: 756
To configure integration of a Check Point log server with App Control: 857
To configure the App Control Server to use SMTP/TLS for notifications: 757
To configure the clone inventory setting for a template: 219
To convert a template computer back to an agent-managed computer: 222
To copy a custom rule: 400
To copy a memory rule: 475
To copy a registry rule: 454
To copy rules to another server under Unified Management: 798
To create a baseline drift alert: 658
To create a Baseline Drift Report: 646
To create a Carbon Black EDR integration user and API Token (summary steps): 769
To create a console login account: 95
To create a custom approval or ban for one or more files on a Files page: 309
To create a custom portlet: 701
To create a Custom Rule from a Notification Details page: 882
To create a global approval or ban for one or more files on a Files page: 308
To create a global cleanup rule for offline clones: 220
To create a new console user role: 103
To create a new dashboard: 694
To create a policy: 182
To create a prevalence alert for a file from its File Details page: 628
To create a Registry Rule from a Notification Details page: 881
To create a Saved View on the Find Files page: 720
To create a snapshot (or add to one) from a file table: 654
To create a snapshot (or add to one) from all files on a computer: 653
To create a software execution meter from a File Details page: 633
To create a template computer: 209
To create a Unified Management login account: 786
To create an AD policy mapping rule: 125
To create an AD role mapping rule: 114
To create an advanced Indicator Set Exception: 667
To create an alert: 613
To create an API user and get its API token: 840
To create an event cache for later viewing: 603
To create an Event Rule that deletes files (example): 321
To create an Expert Rule applied only to operations with a specific tag: 500
To create and configure an approval or ban for a single file: 304
To create and save a view of a console table: 79
To create approvals or bans for a list of hashes: 311
To create Indicator Set Exceptions (default method): 667
To customize an existing notifier: 548
To customize and save an event report as a Saved View: 599
To delete a computer from a App Control Server: 174
To delete a console user role: 116
To delete a custom rule: 430
To delete a dashboard: 697
To delete a login account: 100
To delete a memory rule: 486
To delete a policy: 205
To delete a registry rule: 466

To delete a script rule: 388
To delete a snapshot: 656
To delete a template computer from the console: 217
To delete an alert: 619
To delete event rules in the Event Rules table: 523
To delete one or more approval or ban rules: 306
To delete or edit a Check Point log server integration: 864
To delete or edit a Palo Alto Networks appliance integration: 850
To delete the identity provider for a App Control Server: 774
To disable a custom rule from the Add/Edit Custom Rule page: 429
To disable a custom rule from the Custom Rule table page: 429
To disable a login account: 101
To disable a login banner: 750
To disable a memory rule: 486
To disable a registry rule: 466
To disable a script rule: 388
To disable all event rule processing: 524
To disable automatic local approval of unapproved files on Enforcement Level change: 291
To disable external event logging: 739
To disable notification for a setting in a policy: 563
To disable notification for a specific custom, registry, or memory rule: 563
To disable or re-enable tracking of Microsoft-signed support file instances: 231
To disable reputation approval for a file: 332
To disable reputation approval for a publisher: 332
To disable Unified Management: 791
To display a pre-configured Saved View: 79
To display AD integration configuration options: 726
To display agent management configuration options: 728
To display online documentation from the console: 86
To display server status information: 724
To display the drift by files for a single computer: 643
To display the files view of a baseline drift report: 641
To display the Home page daily event summary: 591
To display the System Configuration page: 722
To do a Windows agent upgrade manually or using third-party tools: 150
To download an agent installer: 134
To download an uploaded file: 903
To download table data to a file: 80
To edit a custom rule: 399
To edit a dashboard: 696
To edit a memory rule: 475
To edit a policy: 195
To edit a portlet on the currently displayed dashboard: 700
To edit a registry rule: 454
To edit a script rule: 388
To edit a unified rule: 802
To edit an approval or ban rule: 306
To edit an existing event report: 600
To edit any portlet from the Edit Dashboard table: 700
To edit the configuration for a client server: 788
To edit the configuration of an identity provider for a App Control Server: 775
To edit the details of a self-signed communications security certificate: 741

To edit, enable or disable an alert: 618

To enable a configured Rapid Config from Rapid Config table page: 512

To enable a custom rule from the Add/Edit Custom Rule page: 429

To enable a custom rule from the Custom Rule table page: 429

To enable agents to verify the server communication certificate: 743

To enable alerts for Rapid Config updates delivered from the cloud: 519

To enable alerts for Updater changes delivered from the cloud: 271

To enable and configure Carbon Black File Reputation: 763

To enable and select policies for a configured Rapid Config: 512

To enable Approval Requests and/or Justifications for a notifier: 567

To enable automatic approval request email responses: 580

To enable device control for a policy: 364

To enable event logging to a Syslog server: 735

To enable External Analytics features in the App Control Console: 910

To enable external event logging to an additional SQL server: 736

To enable file uploads to a Check Point threat emulation appliance: 865

To enable file uploads to a WildFire private cloud for analysis: 852

To enable file uploads to the Check Point cloud for analysis: 866

To enable file uploads to the WildFire public cloud for analysis: 851

To enable integration of Palo Alto Networks alerts with App Control: 847

To enable one subscriber to receive all alert emails: 759

To enable or disable cloud updates of Rapid Configs: 519

To enable or disable cloud updates of updaters: 270

To enable or disable immediate termination of banned processes: 313

To enable or disable rules in the Event Rules table: 522

To enable special notifier routing for session virtualization: 565

To enable System Health Indicators on a App Control Server: 808

To enable the AD Mapping interface: 123

To enable trusted directory approvals of WIM file contents: 273

To enable Unified Management on a server: 780

To enable use of AD logins on the App Control console: 92

To enable, disable or simulate the effect of an Event Rule: 523

To enter a Reg macro: 422

To exclude tracking of exported analytics files: 914

To export rules to a file: 432

To filter results in a table: 73

To generate a code to place a computer in temporary local approval mode: 298

To immediately upgrade one or more agents from the console: 148

To import a new certificate for agent-server communications security: 742

To import rules from a file: 437

To initiate a diagnostic file upload from one agent: 891

To initiate a file upload from a file table: 898

To initiate a file upload from the Computer Details page: 899

To initiate a file upload from the File Instance Details page: 898

To initiate diagnostic file uploads for one or more agents: 891

To install a new agent on a Linux computer: 145

To install a new agent on a Mac computer: 142

To install a new agent on a Windows computer: 138

To install the Splunk App for App Control on the App Control Server: 916

To install the Splunk App for App Control on the Splunk Server: 915

To install the Splunk Forwarder on the App Control Server: 916

To list all computers that have an application: 257

To locally approve individual file instances from a table of files: 292

To locate instances of a file by name: 715

To lock down all computers: 204

To log in to the App Control Console: 53

To log out of the App Control Console: 55

To make file hash bans override custom rules that allow execution: 428

To manage the status of a notification: 879

To manually ban files reported as malware in an external notification: 881

To manually upgrade a Linux agent: 152

To manually upgrade a Mac agent (agent version 7.2.0 or later): 151

To mark a file as an installer: 300

To mark a file as not an installer: 300

To meter execution of specified file(s): 630

To modify or disable the reputation approvals feature: 335

To move a computer in the Default policy to another policy: 172

To move a computer to another policy: 170

To move one online computer to Local Approval (Computer Details page): 297

To open a dashboard: 688

To open the details page for one Approval Request or Justification: 571

To open the Expert Rules interface for Custom Rules: 489

To open the Expert Rules interface for Memory or Registry Rules: 489

To open the External Notification Details page for one notification: 874

To open the External Notifications table in the console: 868

To open the Uploaded Files page: 900

To open, review, and resolve an approval request: 573

To permanently delete a portlet: 701

To place one or more online computers in Local Approval mode: 296

To reconnect a disconnected client server: 790

To re-enable reputation approval for a file: 332

To re-enable reputation approval for a publisher: 333

To remove a certificate approval or ban: 352

To remove a server from unified management: 790

To remove an event cache: 605

To request deletion of a file via the File Instance Details page: 320

To request deletion of all instances of a file from all computers via the File Details page: 321

To request deletion of all instances of selected files via the File Catalog page: 318

To request deletion of selected file instances via the Files on Computers or Find Files page: 317

To restore Local Approval mode computers to their previous policy: 297

To restore the App Control Server to its most recent configuration: 753

To return to the top-level Computer view from computer drift details view: 643

To return to the top-level files view from an associated files report: 643

To save an existing dashboard under another name: 695

To set up the Splunk server to receive Splunk Universal Forwarder messages: 915

To show, hide or rearrange information that appears in table columns: 76

To specify a custom logo for a notifier (Windows only): 560

To specify password requirements for console users: 730

To submit files to an external service for analysis: 883

To temporarily disconnect a client server: 789

To test the effects of a rule with Simulate only mode: 524

To trust users to install software on High Enforcement computers: 280

To uninstall the agent: 154

To update a template computer: 216

To upload installers for rule files and agent packages to a server: 131

To use a Timed Policy Override code on a Mac and Linux computer: 299

To use a Timed Policy Override code on a Windows computer: 299

To use a trusted directory for automatic software approval: 273

To use an Expert Rule to apply tags to an object: 499

To use the App Control Server database backup mechanism: 751

To verify connected computer is running the agent and visible to the server: 146

To view a baseline drift report: 638

To view a snapshot: 655

To view all attachments between a specific device and a specific computer: 376

To view all device models detected by App Control: 366

To view all unique device instances detected by App Control: 371

To view an existing event report: 595

To view an updater's history: 270

To view and change configurable certificate approval options: 287

To view and edit Advanced configuration options: 744

To view certificates in the Publisher Details page: 346

To view complete details for one publisher: 253

To view current agent and rule installer versions: 133

To view diagnostic files: 892

To view files associated with a top-level file in a drift report: 642

To view health indicator events in the console: 813

To view the App Control Licensing configuration page: 759

To view the Applications page: 256

To view the Approval Requests and Justifications table: 569

To view the Certificates Details page for one certificate: 344

To view the Certificates table: 340

To view the Computer Details page for a computer: 159

To view the list of discovered or added publishers: 253

To view the Login Account: User Roles page: 102

To view the rules that affect a policy: 191

To view the System Health page: 809

To view the table of Baseline Drift Reports: 637

To view the table of computers managed by your App Control Server: 156

To view the table of Rapid Configs: 507

To view the template computers in the table of computers: 211

To view threat reports on the Events page: 672

To view, enable or disable Indicator Sets: 664

Chapter 1

Overview


This chapter introduces App Control, explains key concepts and system architecture, and provides a summary of features you can use to prevent unauthorized or malicious file activity on your endpoints.

Sections

Topic	Page
What is App Control?	39
How App Control Works	43
System & Platform Architecture	44
File State, Approving and Banning	47
Security Policies and Levels	49
Operating Strategies	51

What is App Control?

App Control is an industry-leading application control product, used to lock down servers and critical systems, prevent unwanted changes and ensure continuous compliance with regulatory mandates. Combining a trust-based and policy-driven approach to application control with real-time threat intelligence, App Control continuously monitors and records all endpoint and server activity to prevent, detect and respond to cyber-threats that evade traditional security defenses. With open APIs and a broad partner ecosystem, App Control provides exceptional flexibility to seamlessly integrate with both in-house and third-party tools.



Best Possible Protection – With App Control, administrators can stop attacks before they occur. Leveraging App Control's proactive "Default-Deny" prevention capability, App Control can lock down systems to stop malware, ransomware, zero-day, and non-malware attacks.

Instant Visibility – Once installed, the App Control Agent provides administrators with real-time visibility into all executable-type files running across their environment. Working with the Carbon Black File Reputation, the App Control Agent provides administrators with trust ratings and actionable intelligence to easily identify and automatically take action against those files most likely to be malicious.

Continuous Compliance – The cost of compliance is outpacing any other spending in IT. App Control makes regulatory and policy compliance easier and less costly with built-in file integrity monitoring, device control, and powerful change control.

Simplified Software Approval – App Control has many trust mechanisms to help simplify the approval of software. Software can be automatically approved by IT-driven trust mechanisms via software distribution systems, patch management solutions, and application auto-updates. Software can also be automatically approved by cloud-driven trust, such as software reputation and publisher trust using our Carbon Black File Reputation or via file analysis services.

Rapid Configs – App Control includes pre-built rule sets that provide advanced threat detection and prevention delivered from the cloud. These rule sets can detect and prevent malicious activity across endpoints in an organization's environment with minimal effort. Some of the Rapid Configs that we've released include OS Hardening, Browser Protection, and MS Office Protection.

Open API Architecture – App Control's open architecture helps organizations integrate with the entire security stack to automate and simplify the security process. Through its RESTful API and broad partner integration ecosystem, App Control provides organizations with unmatched openness and extensibility to integrate their security solutions for improved automation, reporting and faster security response times, via third-party security products (SIEM, Network, Endpoint, Operations) or custom in-house tools.

Using App Control, you can:

- Stop malicious software by blocking known viruses, trojans, application exploits, and custom and targeted attacks
- Stop zero-day threats by allowing only approved software to run
- Create rules to monitor and control access to the Windows registry
- Create rules to stop “living-off-the-land” attacks that use PowerShell and other scripting tools
- Create memory rules to monitor and control access to specific processes on Windows computers
- Create file integrity monitoring and control rules to prevent or report access to critical, non-executable system configuration files
- Reduce the burden of compliance through streamlined audits, activity monitoring, violation notification, and policy enforcement
- Use Carbon Black File Reputation service to identify and classify the risk associated with the software discovered in your environment using reputation services, and to automatically approve files or publishers considered trusted by the service
- Prevent data theft and leakage by auditing and controlling the transfer of sensitive data to attached storage devices on Windows and Mac computers
- Create rules to approve or ban file execution on storage devices by model or serial number on Windows and Mac computers
- Monitor drift away from a baseline of files to minimize risk, identify needed remediations, maintain compliance, and reduce support costs
- Monitor threats using advanced threat indicators, App Control events, file details, and alerts.
- Automate file- and computer-related actions based on incoming events.
- Use the OpenAPI to integrate third-party network, endpoint, SIEM, and analytic security products and services with the App Control Server for notifications and analysis.
- Export App Control data for use by external analytics products such as Splunk.

[Table 1](#) shows complementary App Control features that provide visibility into what files are on your computers, give you control of unauthorized software and hardware, and allow flexible management of computers at your site:

Table 1: App Control Features

Feature	Description
Live File Inventory and Baseline Drift Tracking	App Control can track all files of interest on all computers all the time. This near-real-time inventory means that App Control can provide a wide variety of information about these files, and about the rate and nature of change across your organization. One benefit of this information is Baseline Drift Reports, which report changes in the file inventory on one or more computers. Another is the ability to locate all instances of a specified executable file that exist on the (fixed) local drives of managed computers.
Carbon Black File Reputation File Identification & Reputation Services	Carbon Black File Reputation identifies and classifies files. It assigns a Trust Factor to files based on a variety of sources, including the source of the file, its prevalence on computers running the App Control Agent, results of anti-virus scanning, and whether it has a legitimate digital certificate. You can automatically approve files or publishers that meet a certain trust threshold.
Event Tracking	App Control keeps an up-to-date database of file-related events, as well as other activities involving the App Control Server or managed computers. From this data, you can view predefined or custom reports that can give visibility into changes to your environment and significant App Control Server operations. You also can trigger alerts based on certain events. Events can be exported to Syslog for integration with SIEM systems, to data analytics systems, and to CSV files.
Modes	Active App Control Agents can be operated in one of two modes: Visibility mode provides the file and event tracking features of App Control, but does not enforce file or device bans or other security restrictions. Control mode blocks banned files and allows you to choose one of three Enforcement Levels to determine how unapproved files (i.e., files neither approved nor banned) are treated. Control policies can be configured to enforce other file and device security rules.
Enforcement Levels and Policies	<p>Enforcement Levels and policies work in combination to control file and device activity on specific computers. Depending upon the Enforcement Level you choose, execution of banned files as well as unapproved (neither approved nor banned) files can be blocked. Enforcement Levels range from very restrictive to no enforcement.</p> <p>Policies are rule sets that include an Enforcement Level and other settings, such as the ability to block or control the behavior of some removable devices on Windows and Mac computers. All computers managed by App Control have an assigned policy.</p>

Feature	Description
Flexible and Emergency Lockdown	<p>You can run different groups of computers at different security levels. For example, you may choose to run some computers at High Enforcement Level, which prevents computers from executing unapproved files that were not present when the App Control Agent was installed, while allowing other computers greater privileges.</p> <p>If necessary, you can implement an emergency lockdown to move all computers to High Enforcement during attacks or high threat periods. You can return the systems to their previous security level when you believe the threat is contained.</p>
File Integrity Monitoring and Control	App Control allows you to create custom software rules that apply to specified files or paths. These include File Integrity rules, with which you can monitor, and if you choose, restrict modifications to a specific folder or folders matching your specification.
Software Rules: Bans	Bans enable you to specify files (by name or hash) to be blocked for some or all computers at your site. You can ban files individually, and also can ban all files identified on a list of hashes you provide. You also can ban all files from a specified publisher.
Software Rules: Approvals	Several complementary software approval methods enable you to approve legitimate software to run on all computers, on groups of computers (i.e., by policy) or to <i>locally</i> approve software to run on a single computer. You can integrate approval rules with the Carbon Black File Reputation service to automatically approve files meeting a specific Trust level according to analysis by the service.
Registry Rules	You can specify rules to protect specific registry key/value patterns from alteration on Windows computers.
Memory Rules	You can specify rules to protect a process from access or alteration by any (or specified) other process(es) or user(s) on Windows computers.
Rapid Configs	Rapid Configs are sets of rules that can be used to accomplish tasks such as application optimization, operating system and application hardening, and approval of files delivered by software distribution systems.
Device Rules: Approvals and Bans	You can approve or ban file execution and writing on detected storage devices on Windows and Mac computers. You can approve and ban device models or specific, individual devices, and you can apply the rules to some or all computers.
Notifiers and User-Initiated Approval Requests	When a App Control rule blocks file access, you can display a notifier that explains the block to the user. The notifier can provide an optional file approval request method that lets you track and respond to requests directly in the App Control Console.
Detection: Advanced Threat Indicators	You can enable advanced threat indicators that will trigger events when suspicious conditions occur, and you can fine-tune these indicators by creating exceptions for events that you consider benign.

Feature	Description
Event-Triggered Actions	You can create Event Rules that specify an action to be performed when a file- or computer-related event occurs that matches filters you define. You also can create an alert that reports when a specified event rule is triggered.
File Deletion	You can delete files on Windows endpoints through the App Control Console and create Event Rules that will automatically delete files reported as malicious.
Integration with Network Security Devices	You can integrate the App Control Server with one or more network security devices or services from third-parties, including Check Point and Palo Alto Networks. You also can integrate with Carbon Black's CB Inspection service to send files for detonation and analysis.
Access via the App Control API	You can use the RESTful API to write code to interact with App Control via custom scripts or from other applications. API code can be consumed over the HTTPS protocol using any language that can create get URI requests, post/put JSON requests, and interpret JSON responses.
Integration with External Data Analyzers	You can export events, file operations data, and file catalog data for use by external analytics products such as Splunk.
System Health Monitoring	You can opt in to System Health indicators that monitor and report on factors affecting the operation of this App Control Server, such as compliance with the operating environment requirements.
Unified Management	If you have more than one App Control Server, you can use Unified Management to designate one server to control many common management functions on any of your connected App Control Servers.

How App Control Works

App Control tracks executable files and monitors their prevalence and execution. *Initialization*, the inventory of existing files by App Control, begins immediately after installation of the App Control Agent on a computer. Each file found on a fixed, local drive of a computer during the initial inventory is *locally approved* on that computer unless it has been already banned on the App Control Server. Local approval does not change the global state of a file.

After initialization, new unidentified files that appear on the fixed, local drives of computers running an agent are classified as having a state of *Unapproved*, both globally and locally, on the computer on which they were found. A file keeps its *Unapproved* state until it becomes *Approved* or *Banned*. Once a file has been approved, it is allowed to execute but continues to be tracked.

App Control features several automatic file approval methods (trusted directories, approved publishers, trusted users, pre-configured updaters for Windows computers, reputation approvals, and bulk approval of files from a list of hashes) that make it easy to approve new software without having to do it file-by-file. You also can manually mark individual files as approved or banned.

Other App Control features monitor activity on your computers, which might help you decide on what files to approve or ban. The App Control Server can tell you:

- Whether a file exists on your computers
- Which computers have the file
- Where and when the file first arrived in your environment
- What is known about the source, category, trust level, and threat of the file
- Whether and when a file has executed, and on which computers
- Whether a file has propagated and, if so, whether it has been renamed
- On Windows and Mac computers, whether attached storage devices (including USB, SCSI, and others) exist on your network, when they first were discovered, and on what computer
- How the inventory of files on computers has changed over time

Files Tracked by App Control

In the App Control Console and throughout this manual, you will often see the term “files.” What constitutes a “file” depends upon the App Control feature:

- For App Control’s live inventory, a “file” is an *executable* or *script* file. When you install the App Control Agent on a computer, it analyzes all files on the fixed, local drives of the system, determines which of them are executables or scripts, and keeps an inventory of these files. Non-executable files are ignored once they are identified. App Control determines that a file is an executable by the content of the file, not its file extension. App Control determines that a file is a script by a combination of factors, and users can add to or modify these script definitions. Only executable and script files can be approved or banned. Certain configuration settings can exclude special cases of these files from tracking and inventory.
- For File Integrity Monitoring, access to *non-executable data and configuration files* can be tracked if you register the files with App Control through a File Integrity Control rule. Once a file or path is covered by such a rule, any attempt to access it generates an auditable event in App Control, and if you choose, the attempt is blocked.

System & Platform Architecture

The App Control architecture consists of the following components:

- App Control Server software provides central file security management, event monitoring, and a live inventory of files of interest on all agent systems.
- App Control Agent software runs on servers, desktops, laptops, virtual machines and fixed-function devices. It monitors files and either blocks or permits their execution based on security policy settings. It also reports new executable and script files to the App Control Server and enforces other rules you configure.
- App Control includes an API that allows programmers to write code to interact with App Control, either using custom scripts or from other applications.
- App Control also may be integrated with third-party products. This includes external analytics products such as Splunk and network security products such as those from Check Point and Palo Alto Networks.

App Control is part of a platform which also includes:

- Carbon Black File Reputation, which compares new files introduced on computers running the App Control Agent and Carbon Black EDR sensor to a database of known files, providing information on threat level, trust factor, and software categorization. If you choose, you can use trust information to automatically approve files in App Control.
- Carbon Black EDR, which provides incident response and threat hunting capabilities. IT continuously records and centralizes all endpoint activity, giving Incident Responders, SOC analysts, and MSSPs complete, real-time information for identifying the root cause of an attack, hunting anomalous behavior, and isolating threats. If you choose, you can configure App Control to receive file information and watchlist events from Carbon Black EDR.
- Carbon Black Defense, which is a powerful, lightweight, next-generation anti-virus solution that protects endpoint, provides visibility into attack patterns and behavior, and includes simple tools for response and remediation.
- Carbon Black Inspection, which can be integrated with both App Control and Carbon Black EDR to provide detonation and analysis services for files uploaded from managed endpoints.

App Control Server

App Control Server software runs on standard, server-class Windows computers. It can be run on a dedicated system or as a virtual machine. The App Control Server manages policies and rules, including software and device approvals and bans, and provides visibility into events and file activity on computers running App Control Agents. The App Control Console, a convenient, web-based user interface, provides access to the App Control Server from any connected computer.

The App Control Server database uses SQL Server, either on the same machine as App Control Server or on separate hardware. Key App Control data is accessible externally through a series of published views in the database that are part of the Live Inventory SDK. App Control events also can be output to a Syslog server or data analytics system for further analysis.

Integrating App Control with Active Directory

You may have already defined and named users, computers, and groups by using Microsoft Active Directory. The App Control Server can take advantage of your Active Directory environment to set access privileges for users of the App Control Console, assign security policies to computers, provide user and computer metadata, and designate certain groups or users to be able to install software (and have it automatically approved) on agent-managed computers.

Unified Management of Multiple Servers

If you have multiple App Control Servers, you can centralize the management of those servers. Unified Management allows you to specify that one server can control many common management functions not only for itself but for any of your other connected App Control Servers. You might choose this option to allow regional IT managers or security personnel to manage their own endpoints but have certain functions centrally managed on all servers in your organization. You might also choose to have different servers to manage different types of endpoints (for example, servers, desktops, POS systems).

See [Chapter 27, “Unified Management of Multiple Servers,”](#) for more details.

App Control Agent

App Control Agent software runs on client computers. It monitors file and process activity and communicates with the App Control Server when necessary. On Windows and Mac computers, it also monitors connected storage devices and registry activity. Even when disconnected from the server, the agent continues to enforce the last specified bans and security policies it received. When a disconnected computer running the App Control Agent reconnects, the agent receives policy and rule updates from the server and communicates relevant file activity that occurred during the time it was off the network.

The App Control Agent runs silently in the background until it blocks a file, at which point it can display a message to the computer user, explaining why the file was not permitted to execute. Depending on the file state, the agent's security level, and other configuration choices, App Control may also let the user on the client computer choose to run a blocked file. You also can enable mechanisms for users to request approval of blocked files, either informally via email or using a formal request process built into and tracked by App Control.

Trust Rating from Carbon Black File Reputation

Carbon Black File Reputation is a web service from by Carbon Black that helps identify and classify software discovered on your computers by comparing it to an extensive database of known files. Based on weighted analysis, Carbon Black File Reputation further assigns a threat level (malicious, potentially malicious, unknown, or clean) and a trust rating (0-10 or unknown) to each file. The App Control Server can include this information in its live file inventory so that you immediately know the threat status and other key information about files on your systems. If you have Carbon Black File Reputation enabled, you can “analyze” any file in the App Control Server inventory to get whatever information is available.

A file's trust rating goes beyond the information available from one anti-virus scan. It is based on a series of factors, including how long and on how many computers the file has been seen, whether it has a trusted digital certificate, and the results of scanning by multiple anti-virus programs.

For example, a file that scans as clean on anti-virus programs, has a trusted digital certificate from a known publisher, and appears on many computers for a long period of time might have a trust rating of 10, highly trusted. Another file that also produces clean anti-virus scans but has only recently been seen, is on very few computers, and does not have a digital certificate might only get a trust rating of 2, low trust.

You can use the trust rating to automatically approve files, either based on their own trust rating or the rating of their publisher. By using Reputation Approvals, administrators can enforce their chosen security posture as it relates to file or publisher trust level and approve high trust software with no administrative overhead.

File State, Approving and Banning

Several key feature groups work together in App Control to secure computers on your network. At the heart of this security capability is the ability to classify files according to their *state*. Groups of security rules, called policies, control how different groups of computers treat files in different states. This section describes the primary file states – approved, banned, and unapproved – and how they can be changed.

Global State

The App Control Server maintains a central database of unique files (determined by hash) for all executable files tracked on computers running the App Control Agent. You can view the *global state* of these files in the File Catalog. Global state determines what the file is allowed to do on agent-managed computers with different Enforcement Levels.

Global state is a combination of:

- *File State*, which indicates the approval/ban state of the file itself, and
- *Publisher State*, which is the approval state of the file's publisher (if known).

A file can have a global state of:

- *Approved* – for all computers
- *Approved by Policy* – approved for some computers, unapproved for others
- *Banned* – for all computers
- *Banned by Policy* – banned for some computers, unapproved for others
- *Unapproved* – for all computers
- *Mixed* – banned for some computers but approved for others

Global State cannot be modified directly, but can be modified by changing the *file state* or *publisher state*. App Control provides a variety of ways to modify the file state. See [Chapter 8, “Approving and Banning Software,”](#) for details. [Chapter 7, “File, Publisher, and Application Information,”](#) shows additional details for files tracked by App Control.

Local State

While the App Control Server keeps a global state for a file, each *instance* of a file on a computer running the App Control Agent has its own *Local State*, which indicates what the file is allowed to do on the computer it was found on, depending upon its Enforcement Level.

Files with a Global State of Unapproved may have different local states. In particular, you can locally approve a file by various methods, as long as that file was not globally banned. App Control includes local file state information in its Files on Computers inventory of all tracked file instances.

A file can have a local state of:

- *Approved*
- *Banned*
- *Unapproved*
- *Deleted* (the file has been deleted recently and will be removed from the database on next update)

In addition to its primary state, each file instance has Local File Details (see [Chapter 7, “File, Publisher, and Application Information”](#)) that may identify the source of its approval or other decisions made about it in App Control. These details are primarily used by Carbon Black Support.

File Approval Methods

Software approval ensures that users of computers running the App Control Agent can freely install and run known-good applications regardless of App Control settings and Enforcement Level in effect. Approving files also can reduce time devoted to tracking files you are not concerned about. App Control supports several complementary methods for approving software on computers:

- When you need to pre-approve applications to run on all computers, you can designate trusted directories, publishers, or updaters to automatically generate approvals.
- When you want to protect against advanced threats and would like to reduce the number of files you need to approve individually, you can enable automatic reputation approvals of files based on file or publisher trust in Carbon Black File Reputation.
- You can approve an individual file by hash, either for all computers or by policy. In addition, you can create multiple individual file approvals by importing a list of file hashes you want to approve.
- When you need to approve software for installation on selected individual computers, either designate trusted users (or groups) to perform installations, or choose one of App Control’s local approval methods.

See [“What is App Control Software Approval?”](#) on page 261 for more details.

File Ban Methods

In Control mode, App Control lets you ban specific files from executing on all computers, or on computers associated with specified policies. You can ban files using the following methods:

- *File-name bans* are platform-specific (Windows, Mac, Linux). For the named platform, they ban execution of named files on either *all* systems on running the App Control Agent or on all systems in policies you specify.
- *Hash bans* prevent files matching a unique hash from executing regardless of the file name used. They are enforced for all platforms, either on all systems running the App

Control Agent or on systems in policies you specify. You can ban more than one file in a single operation by importing a list of hashes.

- *Publisher bans* prevent files identified as being from a specified publisher from executing. They are enforced either on all Windows systems running the App Control Agent or on systems in policies you specify.

See [“What are App Control Software Bans?”](#) in [Chapter 8, “Approving and Banning Software,”](#) for more details.

Custom Rules

In addition to the variety of ban and approval rules described above, App Control provides other ways to protect your computers, allow needed software to run, and optimize performance.

Custom Rules allow you to designate one or more paths, either at the directory or the file level, at which certain activities are allowed or blocked. In some cases this involves changing the state of files, but in others it simply allows, blocks, or disables certain behavior on a case-by-case basis without any global rule changes. You can use Custom Rules for File Integrity Control, to create a Trusted Path for your installation directories, to reduce tracking of files in directories known to be safe or not of interest, and for many other purposes you can configure.

See [Chapter 14, “Custom Software Rules,”](#) for more details.

Rapid Configs

A Rapid Config (short for “rapid configuration”) is a named set of rules that can be used to achieve goals such as optimizing the interaction of App Control and a specific application, hardening of operating systems and applications, and approval of files created or delivered by certain tools or pathways. For example, a Rapid Config for a specific application might include rules that ignore writes to specified folders while approving certain files so that the application can be more easily used.

See [Chapter 18, “Rapid Configs,”](#) for more details.

Security Policies and Levels

App Control policies are groups of protection rules shared by targeted groups of computers running the App Control Agent – every computer running a App Control Agent must belong to a policy. You create policies based on your security and organizational requirements. For example, you might base policy membership on functional role (e.g., marketing, customer service, IT); location; or type of computer (e.g., laptop, desktop, server).

Each policy has its own App Control Agent installer, which is automatically generated on the server when you create the policy. Each installer automatically assigns a policy to each agent it installs. However, if you choose, you can have the App Control Server assign a policy based on Active Directory data for the user and/or computer running the agent each time the computer with the App Control Agent connects to the server.

See [Chapter 5, “Creating and Configuring Policies”](#) for details on policies.

Policy Settings

Policy settings define the way you want App Control to manage a particular group of computers. There are three main categories of settings:

- *Basic Policy Definitions* – These include the policy name and other descriptive information, whether computers in this policy allow agent upgrades, whether live file inventory is activated for these computers, and the basic security level (the Mode and Enforcement Level) for the policy. Modes and Enforcement Levels are described in more detail below.
- *Device Settings* – Device settings control the way a App Control policy treats removable devices. You can make different rules to control read, write, and execute operations on devices, and you can specify that approved and banned devices are treated differently than devices that have not been classified.
- *Advanced Settings* – Advanced policy settings primarily control whether computers in a policy have certain file types blocked. The possible values are Active, Off, and Report Only.

In addition to these, most of the other rules in App Control can be limited to certain policies if you choose, or can be made to apply to computers in all policies.

See [Chapter 5, “Creating and Configuring Policies”](#) for full details on policy settings.

Modes and Enforcement Levels

The Enforcement Level in a security policy controls whether unapproved files (applications that may be unidentified and that have not been approved or banned) are allowed to execute. The availability of different Enforcement Levels enables you to choose a setting for each policy that suits the security and user requirements for the group of computers associated with that policy.

App Control offers three different modes of operation: Agent Disabled, Visibility, and Control. Disabled agents neither enforce rules on nor report information from their computers. Agents in Visibility mode collect and report information but do not enforce rules.

Control mode offers the full range of features, including tracking of files and device activities, and enforcement of bans and other rules that protect your computers. If a file has been banned, it is blocked at *all* Enforcement Levels in Control mode. Control mode Enforcement Levels differ primarily in how they treat unapproved files:

- *High (Block Unapproved)* – Only approved files are allowed to execute.
- *Medium (Prompt Unapproved)* – Approved files are allowed to execute. Attempts to execute Unapproved files cause a notifier dialog to display, in which the user can decide whether to Allow or Block them.
- *Low (Monitor Unapproved)* – Approved and Unapproved files are allowed to execute without prompting. The activity of these files is still monitored by App Control.

In some cases, a computer can have different Enforcement Levels when it is connected vs. when it is disconnected.

App Control Licensing and Modes

App Control Server can be licensed at two feature levels that parallel the available Modes described in the previous section:

- *Visibility* – This provides all of App Control’s file and event tracking and reporting capabilities, but not control features such as file and device blocking.
- *Suite* – This provides the Visibility and Control features of App Control.

License keys determine the number of agents allowed to run in each mode. You can mix licenses on the same server, having, for example, 20 Visibility licenses and 20 Suite licenses. In addition, you can purchase the upgrade at any time to bring Visibility licenses up to Suite level. Keep in mind that if you have *no* Suite licenses, Control features are not available and certain elements of the console documented in this manual will not appear.

See [“Managing App Control Licenses”](#) on page 759 for information for more information on how licenses work in App Control.

Operating Strategies

Your overall App Control operating strategy depends on whether you are only interested in getting *visibility* into file activity on your network or whether you need to exercise a degree of *control* over the use of software and devices. It also could vary according to whether you want all of your computers operating at the same security level or you need to control some more than others. In addition, your strategy might change over time, perhaps due to greater experience with App Control, different threat levels, or the frequency with which your privileged users need to run new software that is not managed by IT.

Different operating strategies will require different amounts of preparation and maintenance. You might want to create a reference system – one computer that has all of the applications you want to approve for all of your users and has no applications you don’t want executed on your users’ computers. You can use this system to create a baseline for analyzing any drift of files on other computers, or over time.

Your Carbon Black Support or Services representative can help you develop an operating strategy appropriate for your environment. In addition, you can log in to the Carbon Black [User Exchange](#) to research strategies suggested by Carbon Black or used by your fellow Carbon Black users. Here is small sample of relevant postings on the User Exchange:

- [App Control Deployment Strategies for Code and Application Developers](#)
- [App Control High Enforcement with a Self-Service Install Folder](#)
- [Carbon Black Enterprise Protection: Value on “Day 1”](#) (video)

You can search the User Exchange for other topics of interest, for example by using “strategy” or “enforcement” as keywords.

Chapter 2

Using the Console

This chapter covers the basics of using the App Control Console: how to log in and out, how to navigate in the user interface from the Home page and menu system, and how to view the information App Control makes available to you through tables, details pages, and dashboards. Mastering the information and tasks in this chapter will give you a head start on all other App Control activities described in this guide.

The Unified Management option adds user interface elements not described here. See [Chapter 27, “Unified Management of Multiple Servers.”](#) for instructions on setting up and using this capability.

Important

The console user interface is documented based on users having full permissions. Features available to a specific user depend upon that user’s account privileges. Any permissions that are turned off will remove related user interface elements. Consider making users with restricted permissions aware of this so that they are not confused by the absence of features described in App Control help. See [Chapter 3, “Managing Console Login Accounts”](#) for details.

Sections

Topic	Page
Console Access	53
Logging Out	55
The Home Page	57
Using the Main Menu	60
Left Navigation Menu and Breadcrumbs	65
Console Tables	66
Details Pages and Object Previews	80
Menus on Details Pages	82
Preference Settings for Console Users	83
Using Online Help	86

Console Access

Access to App Control features is primarily through a browser-based user interface called the App Control Console. You can log in to the console from a web browser on any computer with access to your server, including the App Control Server itself.

Browser Recommendations

Although other browsers with HTML frame support should work, the following browsers are recommended:

- Microsoft Internet Explorer Version 10.0 or higher
- Latest Mozilla Firefox release
- Latest Chrome release
- Latest Safari version (on OS X only)

See the separate *Operating Environment Requirements* document for this release for a list of officially supported browsers.

In Internet Explorer, you may need to adjust your overall security settings or set the console address to be part of your Local Intranet or Trusted Sites zone in order to access the console. You might also need to add “about:blank” to your trusted sites. The security settings are accessed by choosing **Tools > Internet Options** in Internet Explorer and clicking on the **Security** tab.

Logging In

In addition to the connection, browser, and screen requirements described in previous sections, you will need a App Control username and password to log in to the console.

To log in to the App Control Console:

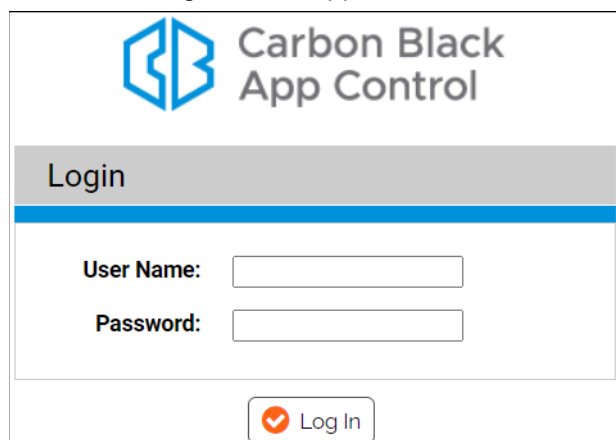
1. From a supported web browser, enter the URL for the App Control Server name you chose during installation, usually the server’s fully qualified domain name or a configured alias:
`https://server_name.domain.extension`
2. If you see a certificate dialog, accept the digital certificate presented for the server. A certificate is required by the web server to support SSL and HTTPS connections.
 - a. If you provided one at installation time, your company’s certificate appears. Otherwise, you see a self-signed certificate created during server installation. You can accept the Carbon Black certificate without compromising security.
 - b. If your browser displays a warning about the certificate, you can safely ignore the warning and click through the remaining confirmation screens.

Note

To avoid future certificate warnings:

- In Firefox, accept the certificate permanently.
- In Internet Explorer, click through the warning, click the Certificate Error button in the IE toolbar, and install the self-signed certificate.
- In Safari, click **Show Certificate** on the warning and check the *Always trust...* box for the App Control Console certificate, and click **Continue**.

The console login screen appears:



The screenshot shows the Carbon Black App Control login interface. At the top left is the Carbon Black logo, followed by the text 'Carbon Black App Control'. Below this is a grey bar with the word 'Login' in white. Underneath is a white form area with two input fields: 'User Name:' and 'Password:'. At the bottom of the form is a 'Log In' button with a checkmark icon.


3. Enter your user name and password. For first-time login, enter the default user name (**admin**) and password (**admin**). For security, change the default password according to the instructions in [“Changing Passwords and Other Account Details”](#) on page 98.
4. Click the **Log In** button.
5. The console Home page appears. The first time any user logs in to the console after installation, there may be a noticeable delay in display of the Home Page. Subsequent logins will be faster for all users.

Unified Management Note: If you have multiple App Control Servers and enable Unified Management, you can use single sign-on from the management server to access its client servers. See [Chapter 27, “Unified Management of Multiple Servers.”](#)

Custom Login Banners

If regulatory requirements or your organization's policies require that users see special text when they access your information resources, this can be configured on the Advanced Options tab of the App Control System Configuration page.

Table 114, “Advanced (Configuration) Options”, on page 746, shows the fields for enabling and configuring a banner and “Adding a Login Banner to the Console” on page 749 provides instructions.




Logging Out

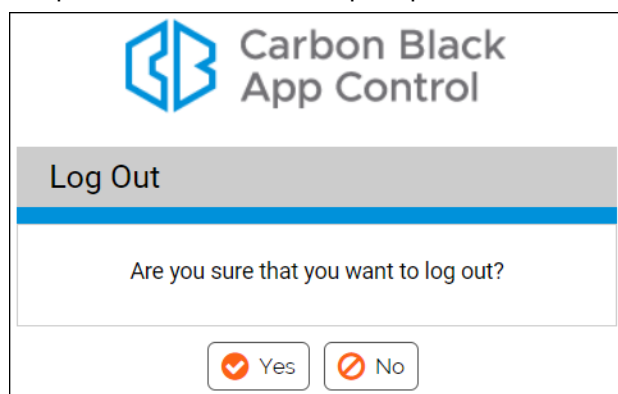
On every page of the console, a log out command is available on the user name menu in the upper right banner area of the web page. Logging out ends your console session.

To log out of the App Control Console:


1. In the user name menu at the top right corner of console menu, choose **Log Out** (with the default login, this menu will show as “Admin”):



2. Respond to the confirmation prompt:



Login, Server, Version and Alert Information



The top banner on console pages includes the main menu, described in [“Using the Main Menu”](#) on page 60. It also shows the following information:

- on the far right, the name of the currently logged in console user, which is also a menu for changing the User Settings (preferences) for that user and also for logging out
- on the right, below the menu bar, the version number of the App Control software you are running
- on the left, the name (or in some cases, the IP address) of the App Control Server
- the number of App Control alerts currently triggered (if any) in each of three categories with separate color symbols: High (red), Medium (orange) and Low (yellow); hovering over the symbol or number shows the alert name if there is a single alert in that category or the alert level if there are multiple alerts

The Home Page

The Home Page provides quick access to common tasks and information. When you log in for the first time, the App Control Home page opens, with the console main menu at the top of the window:

The screenshot shows the Carbon Black App Control console interface. At the top is a navigation bar with the logo, user 'APC-12709', and menu items: Home, Reports, Assets, Rules, and Tools. Below the navigation bar is a 'Home Page' dropdown and a help icon. The main content area is divided into several panels:

- Alerts:** A table showing a 'New Certificate Alert' with details like Type (Certificate Alert), Priority (Low), Date Triggered (Jun 24 2020 02:12:14 PM), and Instances (8). A 'Reset All Alerts' link is present.
- Top X:** Search filters for 'Find top: 10' items, 'Blocks by Computer', and 'Max age: Last Day'. Includes 'Search' and 'Clear' buttons.
- Find Computer:** Search options for 'Computer name or IP' or 'User name'. Includes a search input field and 'Search'/'Clear' buttons.
- Find Files or Events:** Search filters for 'Computer: Any Computer', 'User: Any User', 'Filename: All Files', 'Exact match' checkbox, and 'Max age: Last Day'. Includes 'Search' and 'Clear' buttons.
- Change Policy:** Options to change policy for a specific computer, including 'From existing policy' and 'To new policy' dropdowns. Includes 'Change' and 'Clear' buttons.
- Event Reports:** A summary for the period 'Jun 24 2020 11:25:46 AM to Jun 25 2020 11:25:46 AM' with a table of reports.
- Licensing:** A table showing license usage for 'Visibility' and 'Control'.
- Emergency Lockdown:** A section with a 'LOCK DOWN' button and text: 'Click on the button below to move all your connected computers not under High Enforcement Level to High Enforcement Level.'

The Home Page is a *dashboard*, a configurable page on which you can add and delete *portlets* containing information or controls. See [Chapter 24, “Using and Customizing Dashboards,”](#) for more details on how to use and modify the Home Page and other dashboards. [Table 2](#) below describes the default contents of the Home Page. Note that the Home Page can be modified, so your Home Page may have different portlets:

Table 2: Home Page Quick Access Portlets

Portlet	Links/Buttons	Description
Alerts	Reset/Reset All Alerts	Shows any triggered App Control Alerts that have not been reset, and provides a Reset button for each so you can clear them if you choose. It also provides links from each alert to its Alerts History page for more details about that alert.
Top X	Search/Clear	Shows a table of the top items in various categories – for example, the 10 computers with the most blocked files in the past day. You can specify the number of items to show (default is 10) and the time period over which to look for them (default is 1 day). In the results, clicking on a name (e.g., a computer name) opens a details page for that item. Clicking on a number usually displays the Events page filtered to show events matching your Top X query.
Find Files or Events	Search/Clear	Finds files and events (file blocks, unapproved files, or all events) associated with computers, users or file names you specify. For file name searches, when the "Exact Match" box is checked, only that single file is listed in the results (if found). When the box is not checked, all files containing the string you enter in the box are listed in the results. The Max Age dropdown allows you to determine the time period over which to conduct the search; it defaults to “Last Day”.
Event Reports	New installations	Displays a table of all new file installations that have taken place during the past day (24 hours up to the time you display the page) on Windows computers managed by this App Control Server. Platform Note: Installations on Mac systems are not included in this <i>New installations</i> table. However, the files that are installed appear in tables that show <i>new files</i> .
	New unapproved files	Displays a table of all new unapproved files that have appeared on computers managed by this App Control Server during the past day (24 hours up to the time you display the page).
	Blocked files (by bans)	Displays a table of all banned files that have been blocked on computers managed by this App Control Server during the past day (24 hours up to the time you display the page).

Portlet	Links/Buttons	Description
Event Reports (cont.)	Blocked files (by unapproved status)	Displays a table of all new, unapproved files that have been blocked as a result of the Unapproved Executables setting. The report covers one day (24 hours up to the time you display the page).
Licensing	Manage your licenses	<p>Displays the total number of App Control Agent licenses available on your server and the number in use. If some licenses are for Visibility and some for Control, shows the number for each type.</p> <p>Clicking the Manage your licenses link opens the Licensing panel of the System Configuration page, where you can add App Control licenses, and can configure and activate Carbon Black File Reputation.</p>
Find Computer	Search/Clear	<p>Entering a string that matches all or part of the name or IP address of a computer running a App Control Agent displays a list of matching computers. If you click on a computer in the results, its Computer Details page appears. Computer details include currently Enforcement Levels and connection status. Tabbed views also show details such as last logged in user, agent version, and System Details (if available).</p> <p>Computer name searches are not case sensitive.</p>
Change Policy	Change/Clear	<p>Changes the current security policy of a specified computer. Enter the name or IP address of the computer whose policy you want to change in the upper box. Its current policy is shown. Enter the policy you want to change <i>to</i> in the lower box. Once you click Change, the computer moves to the new policy and stays there unless you explicitly move it again.</p>
Emergency Lockdown	Lockdown/Restore	<p>Lockdown switches all connected computers managed by this App Control Server to High (Block Unapproved) Enforcement Level. Placing computers in High Enforcement Level during high-threat periods helps ensure that no new executable files are permitted to run.</p> <p>When computers are under emergency lockdown, Restore returns them to their pre-lockdown state. If they were in High Enforcement Level prior to the emergency lockdown, they remain in that state.</p> <p>Note: Lockdown <i>does not</i> affect systems that are in Local Approval mode.</p> <p>If you do not have any Control licenses, Lockdown is disabled, but Restore is still available in case machines were locked down at a time when you <i>did</i> have full licenses.</p>

Using the Main Menu



The console main menu at the top of each page allows you to navigate to other console pages. The menu is organized by logical task-groupings, and in most cases shows a submenu of choices when you move the mouse over one of the top-level labels. Clicking on an item in the main menu without making a submenu choice opens the page for the first item on the submenu. A blue bar appears under the currently active menu item.



Table 3: App Control Console Main Menu Choices

Section	Description
<servername>	The server name is always shown in the menu bar. If Unified Management is enabled, the name dropdown provides a shortcut to Unified Management configuration and access to other servers by the management server. See Chapter 27, “Unified Management of Multiple Servers,” for more details.
Home	<p>By default, the console displays the <i>Home</i> page when you log in. Clicking Home in the menu bar returns to this page from other pages.</p> <p>The Home Page provides quick access to information about files, events, computers, and licenses. It also lets you change the policy of a computer or initiate a network-wide lockdown if needed.</p> <p>The Home Page is a <i>dashboard</i> that you can customize to deliver different information in different forms. A dropdown menu on the <i>Home Page</i> lists other dashboards to which you have access. See Chapter 24, “Using and Customizing Dashboards,” for more details.</p> <p>You can change the page that appears first when you log in to the console. See “Preference Settings for Console Users” on page 83.</p>
Reports	<p>Events are messages resulting from activities monitored by or related to App Control. On the Events page, Saved Views provide custom reports for certain types of events, and you can filter any view to create your own report. Events include files blocked, unapproved files executed, and system changes made by console users. For file-related events, you can link directly from an event to the file details. See “Event Reports” on page 590.</p> <p>Cached Events displays a subset of events that a user has chosen to cache for faster display. See “Caching Events for Later Viewing” on page 602.</p> <p>Dashboards displays the Dashboard List page. A dashboard displays information about your App Control installation and the assets it manages through a series of “portlets.” You can drill down for more details about files, computers, events and alerts. The Home Page is a special dashboard. Users can create and optionally share their own dashboards and portlets. See Chapter 24, “Using and Customizing Dashboards.”</p>

Section	Description
Reports (continued)	<p>Baseline Drift displays a page with two tabs:</p> <ul style="list-style-type: none"> • The Baseline Drift tab shows any available reports that analyze the “drift” from a specified baseline file inventory, allows you to run the reports, and allows you to create and configure new reports. • The Snapshot tab on the Baseline Drift page shows any named file lists, called “Snapshots,” that you have created for use in baseline drift analysis. There are several places in the console from which you can create a Snapshot. <p>See Chapter 22, “Monitoring Change: Baseline Drift Reports.”</p> <p>External Notifications displays notifications from network security devices, such as those from Check Point and Palo Alto Networks. If a notification references files or computers shown in your endpoint data, you can correlate data from the two sources. See Appendix C, “App Control Connector.”</p>
Assets	<p>Computers shows a table of computers managed by your server. You can filter the table of computers by various categories. For the computers in the table, you can change the security policy to apply and also temporarily put the computer into Local Approval mode. See Chapter 4, “Managing Computers.”</p> <p>Files displays two tabbed lists of files on your App Control-managed computers:</p> <ul style="list-style-type: none"> • File Catalog is a list of all <i>unique</i> files that have been discovered by agents reporting to your App Control Server. • Files on Computers is a list of all <i>instances</i> of tracked files discovered by agents reporting to your App Control Server. <p>In addition, you can use the Saved Views menu to further specify the files you want to see. Views include Banned Files, New Unapproved Files, Malicious Files, Categorized Files, and Installed Programs.</p> <p>Platform Note: Installed Programs shows Windows programs only.</p> <p>You can use custom filters on the Files page to locate specific files and ban or approve them (locally or globally) as appropriate.</p> <p>See “Viewing File Tables” on page 225.</p> <p>Applications shows two tabbed lists of applications detected on App Control computers reporting to your server:</p> <ul style="list-style-type: none"> • Application Catalog is a table of all unique applications that have been discovered by computers reporting to your server. • Applications on Computers is a list of all instances of applications on computers reporting to your server. <p>See “Application Information” on page 256.</p>

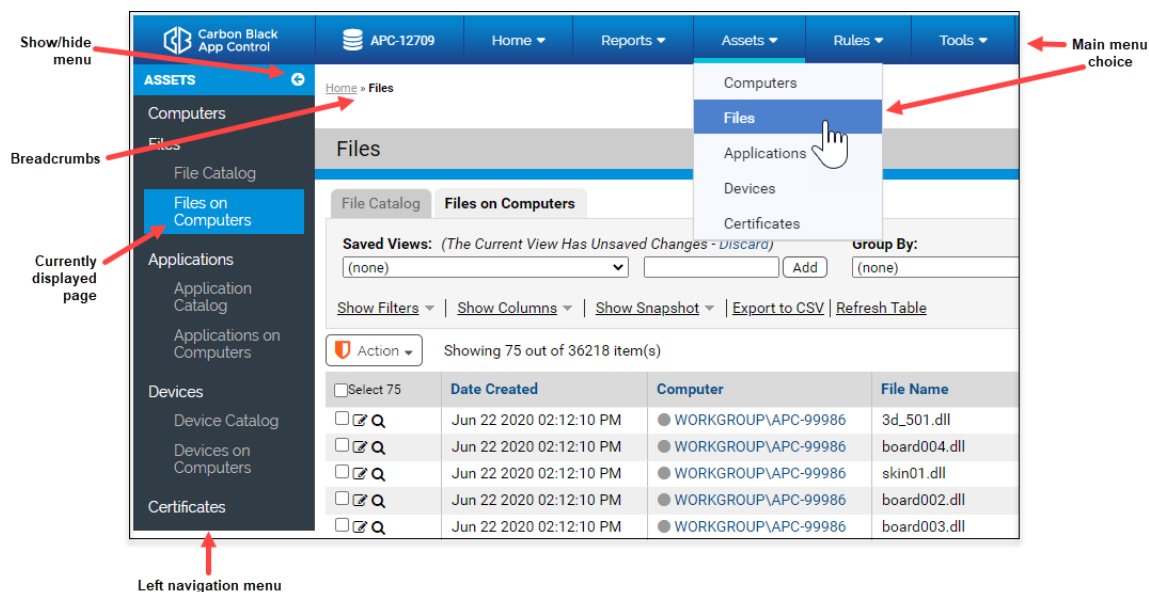
Section	Description
Assets (continued)	<p>Devices displays two tabbed lists of removable devices detected on Windows and Mac computers reporting to your server:</p> <ul style="list-style-type: none"> • Device Catalog has two views. One is a list of all unique device <i>models</i> that have been discovered by agents on computers reporting to your App Control Server; the other lists all <i>instances</i> (i.e., unique serial numbers) found. • Devices on Computers is a list of all unique <i>attachments</i>, which are defined as pairings of one computer and one device. <p>You can globally approve a device so that client computers can access files on it when other devices are restricted. You can ban a specific device so that files on it are never allowed to execute. See Chapter 12, “Managing Devices.”</p> <p>Platform Note: Device discovery and control are currently available on Windows and Mac agents.</p>
Rules	<p>Policies shows the table of existing policies (named sets of security rules) and allows you to edit these policies or create new ones. It also provides a link to the App Control Agent download page.</p> <p>Each policy automatically generates its own agent installation file when created. The installation file used for an agent determines the initial policy for that computer, but computers can be moved to another policy or deleted from the policy when retired from service. See Chapter 5, “Creating and Configuring Policies.”</p> <p>A Mappings tab is added to the Policies page if Active Directory (AD) integration with the App Control Server is enabled on the System Configuration page, and the App Control Server and an AD server inhabit the same AD Forest.</p> <p>Clicking this tab opens the Active Directory Policy Mappings page, where you can set rules by which computers running the App Control Agent are assigned to policies according to one of the AD groups the computer (or its user) belongs to. See “Assigning Policy by Active Directory Mapping” on page 122.</p> <p>Notifiers displays the table of existing blocked file or action notifiers that can be associated with policies and their settings. You can add, delete, and modify notifiers on this page. Notifiers can be configured to appear on an endpoint running the App Control Agent when an action is blocked on that endpoint. See Chapter 20, “Endpoint Notifiers and Approval Requests.”</p> <p>Software Rules displays several categories of App Control rules for approving or banning files and controlling access to computer functions. Each tab shows existing rules, and may allow editing, deleting, creating, and/or enabling or disabling of rules:</p> <ul style="list-style-type: none"> • The Updaters tab lists updaters available to your server. Enabling an updater for a program or set of programs permits end-users to install application updates when they become available for download via the application update program. Platform Note: Updaters are platform-specific. • The Rapid Configs tab lists sets of rules that can be used to accomplish tasks such as application optimization, OS and application hardening, and approval of files delivered by software distribution systems. Platform Note: Rapid Configs are platform-specific.

Section	Description
Rules (continued)	<p>Software Rules (continued)</p> <ul style="list-style-type: none"> • The Publishers tab lists software vendors for which App Control can confirm one or more valid digital certificates. Publishers can be approved or banned through this page. • The Users tab lists users or groups trusted with permission to install files on any computer they log into with their credentials. • The Directories tab lists authorized approval directories in which all software is approved. • The Files tab lists individual file approvals and bans. • The Custom tab lists custom rules, such as specifying how and where files are allowed to execute or write and, whether a file is tracked by App Control. • The Memory tab lists App Control rules controlling retrieval of information about, modification of, and execution (or termination) of specified processes. Platform Note: This feature applies to Windows agents only. • The Registry tab lists App Control rules controlling creation, modification, and editing in the Windows Registry. Platform Note: This feature applies to Windows agents only. • The Scripts tab lists rules that define which files are tracked and controlled as scripts in App Control. • The Reputation tab appears if Carbon Black File Reputation is enabled on the System Configuration/Licensing page. Reputation-based file and publisher approvals can be enabled and disabled on this tab. <p>Event Rules displays the Event Rule table. Event rules specify an action to be performed when an event matches filters you define.</p> <p>Indicator Sets displays the Indicator Set table. An Indicator Set is a group of advanced threat detection rules that can be enabled to increase the visibility of suspicious activities.</p>
Tools	<p>Meters enable you to monitor the number of executions of files you specify, and the users and computers executing them.</p> <p>Alerts provide notifications in the console and via email when certain conditions occur. Alerts can be made policy-specific.</p> <p>Find Files enables you to locate all instances of an executable file on computers running the App Control Agent on your network. You can make similar searches from the Files page using filters, but Find Files is pre-configured for this purpose.</p> <p>Approval Requests displays a list of file approval requests received from users on computers running the App Control Agent. Requests are created when a user is blocked from a file action and requests that the file be approved. The Approval Requests page shows request status along with information about the file and the requester.</p>

Section	Description
Tools (continued)	<p>Requested Files displays a page with three tabs, each of which is a table of files. The tabs are:</p> <ul style="list-style-type: none"> • Uploaded Files – Shows the list and the status of files that a user requested to be uploaded to the server from an agent computer. • Analyzed Files – Shows the list and the status of files that a user or rule requested to be sent to an external device for analysis. • Diagnostic Files – Shows the list and the status of diagnostic files that a console user requested to be uploaded to the server from an agent computer.
 (configuration)	<p>Login Accounts displays the Login Accounts page for creating and managing users of the <i>console</i>. Note that login accounts are not needed for the users of computers running the App Control Agent.</p> <p>System Configuration provides access to pages for tasks including the server configuration; managing log files; securing communications with agents; configuring backups; downloading software updates; and configuring optional App Control services, including integration with Active Directory. System configuration features are available only to administrator-level login accounts.</p> <p>System Health displays the System Health page, which provides a summary of the state of factors affecting the operation of this App Control Server plus more detailed information about specific factors, such as compliance with the operating environment requirements for a server.</p> <p>Update Agent/Rule Versions displays the drag-and-drop interface for updating agent installation packages and their associated rules on the server. See “Uploading Agent Installers and Rules to the Server” on page 131 for more information.</p>
 (help home page)	<p>Clicking the blue question mark button in the main console menu displays the home page for App Control help in a separate browser window. To go directly to information about the page you are on, click the black question mark button next to the page title.</p>
<username>	<p>The name of the currently logged in user is shown on the far right of the menu bar and provides a menu with two choices:</p> <p>User Settings enables each user (including ReadOnly users) to change their password, choose the first page seen upon login, determine the default number of rows on table pages, enable resizable columns, and specify whether the console maintains customizations to a page between visits.</p> <p>Logout logs the user out of the App Control Console.</p>

Left Navigation Menu and Breadcrumbs

For any console page other than a dashboard, a navigation menu is displayed on the left side of the page. This navigation menu shows the page choices available under the section of the console main menu you currently are in. For example, if you click **Assets** in the top menu and choose **Files** from the menu, the Files page opens showing the tab displayed when you were last on the page. To the left of the table, a menu appears showing the tab choices under Files, and you can click on either of these choices to display its associated view. You can collapse or expand the left navigation by clicking on the boxed arrow button in the upper right of the menu.



When you navigate to a console page, a trail of “breadcrumbs” is shown in the upper left of the page, indicating the path to your current page. In the illustration above, **Home > Files** is the path to the page shown. You can navigate back to a previous location on the path by clicking on it.

Console Tables

Much of the file and computer information you see while using the console appears in tables. Console tables list each primary item on the page (for example, each file on a Files page) in its own row with data related to the item. You can control many aspects of the “view” you have of the information in these tables, and if you like a particular view, you can name it and save it. While the emphasis in this section is on *viewing*, console tables also include many of the controls you use to take *action* on files and computers. These actions are described in detail in later chapters.

Note

This section describes the tables currently used on most console pages. Dashboard pages have different layout and buttons. See [Chapter 24, “Using and Customizing Dashboards”](#) for a description of dashboard elements.

The Files page illustrates many of the typical elements in console tables.

	First Seen Date	First Seen Name	Publisher or Company	Product Name	Trust	Global State
<input type="checkbox"/>	Jul 16 2012 01:26:16 PM	shlwapi.dll	Microsoft Corporation	Microsoft® Windows® Operating Syst.	10	Approved
<input type="checkbox"/>	Jul 16 2012 01:26:16 PM	magicdisc.exe	MagicISO, Inc.	MagicDisc	4	Unapproved
<input type="checkbox"/>	Jul 16 2012 01:26:16 PM	smss.exe	Microsoft Corporation	Microsoft® Windows® Operating Syst.	10	Approved
<input type="checkbox"/>	Jul 16 2012 01:26:17 PM	msxml3.dll	Microsoft Corporation	Microsoft(R) MSXML 3.0 SP11	10	Approved
<input type="checkbox"/>	Jul 16 2012 01:26:17 PM	taskhost.exe	Microsoft Corporation	Microsoft® Windows® Operating Syst.	10	Approved
<input type="checkbox"/>	Jul 16 2012 01:26:17 PM	comctl32.dll	Microsoft Corporation	Microsoft® Windows® Operating Syst.	10	Approved
<input type="checkbox"/>	Jul 16 2012 01:26:17 PM	dimsjob.dll	Microsoft Corporation	Microsoft® Windows® Operating Syst.	10	Approved
<input type="checkbox"/>	Jul 16 2012 01:26:17 PM	taskschd.dll	Microsoft Corporation	Microsoft® Windows® Operating Syst.	10	Approved
<input type="checkbox"/>	Jul 16 2012 01:26:17 PM	ntmarta.dll	Microsoft Corporation	Microsoft® Windows® Operating Syst.	10	Approved
<input type="checkbox"/>	Jul 16 2012 01:26:17 PM	version.dll	Microsoft Corporation	Microsoft® Windows® Operating Syst.	10	Approved
<input type="checkbox"/>	Jul 16 2012 01:26:17 PM	dnsapi.dll	Microsoft Corporation	Microsoft® Windows® Operating Syst.	9	Approved
<input type="checkbox"/>	Jul 16 2012 01:26:17 PM	netprofm.dll	Microsoft Corporation	Microsoft® Windows® Operating Syst.	10	Approved

Tables feature various buttons and menus that enable you to configure results and execute actions. In addition to the Help button that appears on every page, console pages that show tables may include:

- [Table Data Control Links](#)
- [Table Column Resizing](#)
- [Row Action Buttons](#)
- [Checked Row Action Menus](#)
- [“Add” Buttons](#)

Table Data Control Links

On many console table pages, a row of text links above the table head allows you to take actions on table data. [Table 4](#) show the possible Table Data Control links (not all appear on all pages).

Table 4: Table Data Control Links

Link Text	Action
Show/Hide Filters	Shows or hides the Filters panel, which lets you narrow the number of results returned in the table.
Show/Hide Columns	Shows or hides the Column Settings panel, which lets you specify which columns are displayed and in what order.
Show/Hide Snapshot	Shows or hides the Snapshot panel, which allows you to add selected files to an existing “snapshot” of files or create a new snapshot. Snapshots can be used to measure Baseline Drift. See “Managing Snapshots” on page 653 for more information.
Export to CSV	Saves the information displayed in the current table to a file, using the standard download method for the current browser. Exported data is formatted as a CSV (comma-separated-value) file suitable for opening as a spreadsheet. Time values output to CSV files are shown in the local time defined by the Server Timezone setting on the System Administration page.
Refresh Page	Refreshes the page view to show the most current data available from the App Control Server. This can be useful if you have been on a page for a long period of time or the page contains information known to change frequently.

Table Column Resizing

One way to control the width of a table is to add or remove columns using the Show Columns link. You also have the option of resizing table columns. This feature is available when you see vertical borders between columns. You can enable and disable re-sizable columns on the User Settings page, which you access by choosing **loginname > User Settings** in the console menu.

You change column width by hovering the mouse cursor over a column border, holding down the left mouse button, and moving the mouse. If you make a column narrower than its contents, text is abbreviated with an ellipsis (...) at the end.

Row Action Buttons

Rows in dynamic tables include information about objects such as client computers, devices, events, reports, or files. Many tables include buttons at the far left of each row that operate on that row.

<input type="checkbox"/>	Date Created	Computer	File Name	Publisher or Company
<input type="checkbox"/>	Jun 17 2016 10:05:11 PM	● MYCORP\SERVER1	setup.exe	Google Inc
<input type="checkbox"/>	Jun 17 2016 10:05:11 PM	● MYCORP\SERVER1	chrome.exe	Google Inc
<input type="checkbox"/>	Jun 17 2016 10:05:11 PM	● MYCORP\LAPTOP2	adapter.dll	Google Inc
<input type="checkbox"/>	Jun 17 2016 10:05:07 PM	● MYCORP\SERVER6	d3dcompiler_47.dll	Microsoft Corporation

Table 5: Common Row Action Buttons

Button	Label	Action
	View Details or Edit Properties	Displays details of an item in a row. If the item has editable properties, clicking this button opens its editor.
	Delete	Removes the item in its row from the table and deletes it from your App Control database.
	View Report/History/Etc.	Displays a report, history, or other information corresponding to the item in a row.
	Find File	Displays the Find Files page and automatically uses the file name or hash of the file in the current row as the search parameter.

Note

Different tables include different combinations of row action buttons (not necessarily all of them), as appropriate for the types of information displayed. Some tables have page-specific buttons not shown above.

Toggle Switches for Enabling and Disabling Rules

In some tables, rules that you are allowed to disable and enable are shown on the table pages with a toggle switch in the Status column. If the toggle button is on the right and the background is green, the rule is enabled. If the toggle button is on the left and the background is white, the rule is disabled. Clicking on the button changes the status. Rules whose status you cannot change in the table show their status in text rather than with the switch.

	Rank ▲	Status	Platform	Rule Type	Name
<input type="checkbox"/>	47	<input checked="" type="checkbox"/>	Windows	Advanced	Track system profile
<input type="checkbox"/>	48	<input type="checkbox"/>	Linux	Performan...	[Sample] Ignore Puppet master and agent
<input type="checkbox"/>	49	<input checked="" type="checkbox"/>	Windows	Performan...	Cb Enterprise Response Sensor Optimization
	50	Disabled	All Platf...	Internal	Block executables run from a network drive
	51	Enabled	All Platf...	Internal	Block executions from banned removable devices
	52	Disabled	All Platf...	Internal	Block executions from unapproved removable devices

Checked Row Action Menus

On many pages, there is an Action menu with commands that take action on any checked rows in the table on that page. For example, if you are on the File Catalog tab of the Files page and you check the box next to “abc.exe”, the Action menu allows you to globally approve or ban the file, remove an approval or ban if one exists, acknowledge the file, or analyze it in Carbon Black File Reputation.

Files

File Catalog Files on Computers

Saved Views: (The Current View Has Unsaved Changes - Discard) (none) Add Group By: (none) Ascending Max Age: (none) Show Individual Files

Show Filters | Show Columns | Show Snapshot | Export to CSV | Refresh Table

Action Showing 625 out of 36218 item(s)

Vendor or Company ▲	Product Name	Trust	Global State
Microsoft Corporation	Microsoft® Windows® Operating Syst...	10	Approved
Microsoft Corporation	Microsoft® Windows® Operating Syst...	10	Approved
Microsoft Corporation	Microsoft® Windows® Operating Syst...	10	Approved
Microsoft Corporation	Microsoft® Windows® Operating Syst...	10	Approved
Microsoft Corporation	Microsoft® Windows® Operating Syst...	10	Approved
Microsoft Corporation	Microsoft® Windows® Operating Syst...	10	Approved
Microsoft Corporation	Microsoft® Windows® Operating Syst...	10	Approved
Microsoft Corporation	Microsoft® Windows® Operating Syst...	10	Approved
Microsoft Corporation	Microsoft® Windows® Operating Syst...	10	Approved
Microsoft Corporation	Microsoft® Windows® Operating Syst...	10	Approved
Microsoft Corporation	Microsoft® Windows® Operating Syst...	9	Approved
Microsoft Corporation	Microsoft® Windows® Operating Syst...	10	Approved

View Carbon Black Reputation Data

Version	File Name	Vendor	Product Name	Trust	Global State
Jul 16 2012 01:26:17 PM	version.dll	Microsoft Corporation	Microsoft® Windows® Operating Syst...	10	Approved
Jul 16 2012 01:26:17 PM	dnsapi.dll	Microsoft Corporation	Microsoft® Windows® Operating Syst...	9	Approved
Jul 16 2012 01:26:17 PM	netprofm.dll	Microsoft Corporation	Microsoft® Windows® Operating Syst...	10	Approved

The choices on the Action menu vary according to the page you are on and in some cases the options you have configured.

Unified Management Note: The name of your server appears on many of the menu commands to indicate that the action is taking place on that server. If you are using Unified Management to manage multiple servers, some of the Action menu commands allow you to choose which servers an action applies to.

Row Selection Scope

You can check the box next to one or more rows and apply the Action menu command to all rows that you have selected. In addition, there is a box above the top row that allows you to check all of the rows that are currently loaded for viewing.

On all table pages with checkboxes, you can select the row by clicking the box itself. On the Events, Files, and Software Rules pages, clicking anywhere in the row selects the row and checks the box.

The actual range of rows you can check and act on at one time depends on whether the page you are on uses dynamic scrolling or fixed-length pages, and whether items on the page are grouped:

- **Fixed-length pages** – On fixed-length pages, any action you take on checked items affects only the checked items on the current page. For example, if a console table has three pages and you check items on page 2 and then go back to page 1, the checkmarks are cleared from page 2. If you check some items on page 1 and then choose Approve Globally on the Action menu, for example, only checked items you see on page 1 are approved, even if you previously checked items on other pages.

This also means that when you check the checkbox in the table head, it checks all the items (or all the items that can be acted upon) in the rows on the current page, regardless of page length. It does not check the rows on any other page.

- **Pages with grouping** – When items on a page are grouped, only items in an expanded group (i.e., the group members are visible) can be checked and acted upon. If the group is collapsed (i.e., only the *group name* is showing), none of the items in the group are checked.

The screenshot shows the 'Events' table in the console. At the top, there are 'Saved Views' and 'Action' menus. Below the table, there are navigation controls and a search bar. The table has columns for 'Timestamp', 'Severity', 'Type', and 'Subtype'. A 'Select all on page box checked' checkbox is checked. A 'Collapsed group: Type: Computer Management' is shown with a minus sign and a red box around it. An 'Expanded group: Type: Discovery' is shown with a plus sign and a red box around it. Below the groups, two rows are visible, both with checked checkboxes. The top row is 'Sep 12 2016 09:43:52 PM Notice Discovery New unapproved file to corr' and the bottom row is 'Sep 12 2016 09:43:51 PM Notice Discovery New unapproved file to corr'.

- **Dynamically scrolled pages** – If a page uses dynamic scrolling, using the Action menu affects all checked rows that have been loaded (scrolled to) during the current visit to the page, regardless of how long the “page” is. Likewise, checking the top box in the table to check all items will check every row that has been loaded. This behavior may be confusing given the dynamic nature of these pages, so extra care should be taken when using the Action menu when it could affect row that you cannot currently see.

Dynamically scrolled pages provide a count of how many rows have been loaded during a visit to that page. In addition, the top checkbox tells you how many rows will be checked if you click it. For example, if the top box says “Select 100,” that means that 100 of the currently loaded are selectable, and if you click that box, the boxes next to all 100 will be checked. If you then scroll down so that 150 rows are loaded,

the top box is un-checked and its label changes to “Select 150,” but the initial 100 rows that you checked remain checked. Checking the top box and then clicking it again to un-check the box deselects all rows.

Keep in mind that the count for the top box only shows *selectable* rows. The checkbox for a row is not selectable (and is grayed out) if the item in that row cannot be acted upon by the commands on the action menu. For example, if the Action menu on a rules page only includes commands to enable and disable rules, you cannot check the box next to certain rules unless the rule has been configured.

The screenshot illustrates the behavior of the selection controls in the console. It shows three panels of a table with columns for 'Date Created' and 'File Name'. Each panel has an 'Action' dropdown and a 'Showing 100 out of 6984 item(s)' indicator.

Panel 1: Dynamically scrolled page with no checked rows
 The 'Select 100' checkbox is unchecked. The table lists 100 items, including parity.exe, chrome.exe, mmc.exe, parityreporter.exe, vmttoolsd.exe, and emet_agent.exe.

Panel 2: Dynamically scrolled page with all loaded rows checked
 The '100 Selected' checkbox is checked. The table lists the same 100 items as in Panel 1.

Panel 3: Additional rows loaded after checking first 100
 The 'Showing 150 out of 6984 item(s)' indicator is present. The 'Select 150' checkbox is unchecked. The table lists 150 items, including iprestr.dll, iphlpapi.dll, and filter.dll.

Row Rank Arrows

On some tables, the ranking of rows affects how App Control processes rules. For example, on the Custom Rules page, rule number 1 is processed before rule number 2, etc. These tables show rank numbers for each row, and also can be sorted in rank order.

On table where rank matters, there are arrows in each row (except for special cases) that allow you to move rules so that their rank is higher or lower. In addition, you can drag and drop a row to change its rank in most of these tables, and you can also specify a move location by holding down the left mouse button on the rank number of a rule.

“Add” Buttons

On pages where you can create a new instance of something, such as a policy or alert, there will be a button for adding that item. For example, if you wanted to create a new alert, you would go to the Alerts page and click the **Add Alert** button to open a form allowing you to configure the new alert. These Add buttons generally appear in the upper left area of the page.

	Name	Type	Enabled	Priority
+ Priority: High				
- Priority: Low				
<input type="checkbox"/>	Updater Modified Alert	System Alert	Yes	Low
<input type="checkbox"/>	Justification Alert	Approval Request Alert	Yes	Low
<input type="checkbox"/>	Reset ! New Certificate Alert	Certificate Alert	Yes	Low
<input type="checkbox"/>	Indicator Set Alert	Event Alert	Yes	Low

Pages, Tabs and Saved Views

Each console page that contains tables provides a specific type of information, such as a table of files, a table of computers, or a table of events. On many pages, you can choose among different “views,” which limit the data on that page to certain parameters, and you can create new views that suit your need. A table page may have one or more of the following features:

- **Tabs** switch you from one major subset of information on the page to another. For example, on the Files page, one tab shows the Files Catalog and another shows the Files on Computers table.
- **Filters** allow you to limit data in a table to items matching criteria you specify. For example, you can filter a files table to show only those with a particular approval or ban state, or only those with a particular Threat level. Filters can be used with or without saving the views they create.
- **Column controls** allow you to show different information about each item in a table. For example, you can eliminate a column showing the date a file was created but add one that indicates whether anyone has executed the file. As with filters, special column configurations can be incorporated into Saved Views or just used in passing.
- **Saved Views** can filter out unwanted items from the table and also can change the types (columns) of data shown for each item. The App Control Console provides pre-configured Saved Views, and you also can create your own. Not all pages have Saved Views.
- **Group By** gives you a menu of choices for different ways to group information in a table. For example, on the Computers page, you can group by Policy, which creates a list of policies, each of which you can click on to show all computers in that policy.

- *Max age* allows you to limit the results shown in a table to those covering a period of time you select on the menu.

You can choose to have the console return each page to its default view when you navigate away from it and come back, or you can have the console “remember” your most recent page view choices and apply them when you next visit the page. See [“Preference Settings for Console Users”](#) on page 83 for more details.

Filter Options

Filters let you narrow information displayed in a table so that you can more easily find the data you need. You can select one or more attributes, which correspond to information in table columns, and then enter attribute values on which to search. The operators you can use with the filters (“equal to”, “greater than”, etc.) vary according to the attribute you select. Depending on the filter you choose, its values can be text, numbers, or dates. For attributes that accept date values, the console displays a date box.

To filter results in a table:

1. Click **Show Filters** to open the Filters dialog. The Show Filter link changes to Hide Filter.

The screenshot shows the 'Files' console interface. At the top, there's a 'Files' header and a 'File Catalog' tab. Below that, there's a 'Files on Computers' section. It includes 'Saved Views' (with a dropdown set to '(none)' and an 'Add' button), 'Group By' (with a dropdown set to '(none)'), and 'Max Age' (with a dropdown set to '(none)'). There are also links for 'Hide Filters', 'Show Columns', 'Show Snapshot', 'Export to CSV', and 'Refresh Table'. The 'Filters' section has an 'Add filter' dropdown and 'Apply', 'Cancel', and 'Reset' buttons. Below the filters, there's an 'Action' dropdown and a status bar indicating 'Showing 75 out of 36218 item(s)'. A table is displayed with columns: 'Date Created', 'Computer', 'File Name', and 'Publisher'. The first row shows 'Jun 22 2020 02:12:10 PM', 'WORKGROUP\APC-99986', and '3d_501.dll'.

2. In the Add Filter menu, select one or more filter attributes you want to use to limit information displayed in the table.

The screenshot shows the 'Files' console interface with the 'Filters' dialog open. The 'Add filter' dropdown is now populated with two filters: 'Category is Media Players' and 'File Type is Application'. The 'Apply' button is highlighted. The table below shows the same 75 items out of 36218, with the same columns as the previous screenshot.

3. For each filter attribute, select the appropriate operators and enter values (if required).
4. To filter results by the selected attributes, click the **Apply** button.
5. To return to a display of unfiltered results, click the **Reset** button.
6. To close the Filters panel, click **Hide Filters**.

The default operator varies depending upon the attribute you choose, sometimes for performance reasons. For example, “is” is the default operator for File Name in order to limit the amount of data matching the filter.

You usually can add multiple filters of the same type clicking the plus button to the right of an existing filter of that type (or by picking the same filter from the menu again). Two filters of the same type are treated as an either/or operation. For example, if you add a File Name filter for filenames containing “alpha” and another for filenames containing “beta”, the table will show files containing either “alpha” or “beta” in the name.

For the “value” field (the data that you want to match), many filters do auto-completion as you enter in characters. For example, if you type in “Abc” in a Product Name filter with a “is” operator, the console displays a menu of all product names that begin with “Abc”, and you can pick one from the menu rather than typing in the entire name. Auto-completion is generally limited to filters using the “is” or “is not” operator.

Some filter criteria are more efficient than others. In general, a filter that allows searching for an exact match rather than requiring a string analysis will be much faster and have less likelihood of database timeouts. For example, if you want to find all files with a particular extension (such as .exe) using the File Name filter and choosing “ending with .exe” is very inefficient. In this case, use the Extension filter. Searching for a file using the “containing” operator (such as, “File Name contains setup”) is particularly inefficient.

Filters apply only to the level of information *currently displayed* in a table. For example, if you display a list of file groups (the default) rather than individual files, a filter that looks for First Seen Name containing the text “abc” will only match the names of *installer files* containing that string. It will not match individual files installed by another file. On the other hand, if you click the *Show individual files* box with the same filter in effect, any file containing the filter string installed by the installer will appear in the table.

Notes

- You can click the Show Filters button and the Show Columns button to show both panels at the same time. This combination might provide more insight into how you would like to modify a particular table.
- To save a view that you would like to use regularly, create a new Saved View. See [“Default and Saved Views”](#) on page 78.

Filter Shortcuts

Some pages in the console include filter shortcuts that automatically fill in some or all of the information in the Filters dialog. These shortcuts are available on the Events, Files, and Software Rules pages. These shortcuts are available on the Events, Files, and Software Rules pages for columns in which a funnel icon appears when you move the cursor over the right end of a heading or cell.

- **Column head filters** – If you move the cursor over the right end of a column head, a funnel icon appears. Clicking on the funnel icon opens the Filters dialog with the column head as the filter. You can then fill in the operator (“is”, “contains”, etc.) and the value you want to use, and then **Apply** the filter.

- Column cell filters** – If you move the cursor over the right end of a specific cell in a column, a funnel icon appears. Clicking on the funnel icon opens the Filters dialog with the column head as the filter field and the contents of the cell as the value. An operator appropriate to the data type is also automatically added, *and the filter is applied immediately*.

Notes

- Filters created through shortcuts are *added to* any existing filters on the page. These new filters may conflict with definitions already in place, in which case an error dialog appears.
- Time-based filters might change the Max Age setting in a way that alters what you see if you Reset the filters.

Show Columns Options

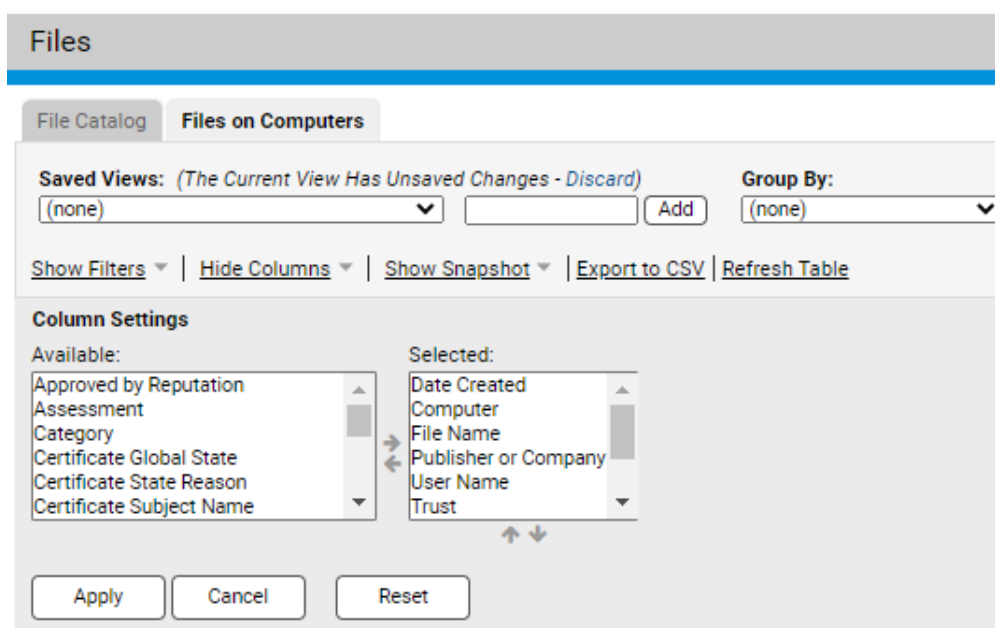
The Show Columns link opens a Column Settings panel where you specify which columns are displayed and in what order for a particular table:

- Items in the *Selected* column are displayed in the table.
- Items in the *Available* column are not displayed in the table.

Because there is a very large number of possible columns for most pages, not all columns are shown by default, and there are different column defaults for different console pages. You can reset any table to its initial, default columns.

To show, hide or rearrange information that appears in table columns:

1. Click **Show Columns**. The Column Settings panel appears and the Show Columns button becomes a Hide Columns button:



2. To hide a currently displayed column:
 - a. Select a column heading in the Selected list.
 - b. Click the left-arrow icon to move the column heading into the Available list.
 - c. To accept changes and update the table display, click the **Apply** button.
3. To display a currently hidden column:
 - a. Select a column heading in the Available list.
 - b. Click the right-arrow icon to move the column heading into the Selected list.
 - c. To accept changes and update the table display, click the **Apply** button.
4. To change column order:
 - a. Select a column heading in the Selected list.
 - b. Click the up arrow or down arrow (below the Selected list) to change the position of the column in the table. The top-to-bottom item order in the list corresponds to a left-to-right orientation of columns in the table. You can only move items that are visible in the table (i.e., column headings that appear in the Selected list).

- c. To accept changes and update the table display, click the **Apply** button.
5. To restore the table to the default settings for the current view, click the **Reset** button

Notes

- You can open the Filters and Columns Settings panels at the same time. The combination of the two might provide more insight into how to best modify a particular table.
- If you use column controls to configure a view that you think you would like to use regularly, you can name it so you can access it again as a Saved View. See “Default and Saved Views” on page 78.
- On the Events, Files, and Software Rules pages, you can use drag-and-drop to rearrange columns.

Tabs

Tabs switch you from one major grouping of information to another within a page. For example, on the Files page, you can click the File Catalog tab, which (if not modified) shows all of the *unique* files (i.e., not each instance of the same file) discovered on App Control Agent-controlled computers on your network. The other tab on that page, Files on Computers, shows all *instances* of all tracked files found on your computers. In some cases, different actions are available on a page when you change tabs.

Table Length, Scrolling and Selection

Certain tables, most notably, the File Catalog and Files on Computers pages, use “dynamic scrolling” (sometimes called “infinite scrolling”), in which more items load as you scroll to the last displayed row. In this mode, there are no defined-length “pages” *per se*, and the length of the table in the current view is whatever number of rows you have scrolled to. At the bottom of the current view, the number of rows displayed is shown, and as you scroll down, more load.

On dynamic scrolling pages with grouping turned on, including pages that display Unified Management data by server, the page length of an expanded group is initially limited to 25 rows so that you can easily scroll to and expand the next group heading. You can either click on the next group heading to expand it or click one of the *Click to load more* links to see more rows for the currently displayed group.

		Sep 6 2016 09:49:23 PM	oleaut32.dll	Microsoft Corporation		10		<Initialization files>
		Sep 6 2016 09:49:23 PM	crypt32.dll	Microsoft Corporation		10		<Initialization files>
		Sep 6 2016 09:49:23 PM	cryptbase.dll	Microsoft Corporation		8		<Initialization files>
		Sep 6 2016 09:49:23 PM	cryptsp.dll	Microsoft Corporation		10		<Initialization files>
		Sep 6 2016 09:49:23 PM	bcrypt.dll	Microsoft Corporation		10		<Initialization files>
		Sep 6 2016 09:49:23 PM	bcryptprimitives.dll	Microsoft Windows		10		<Initialization files>
<div style="display: flex; justify-content: space-between;"> Click to load more Click to load more Click to load more </div> <div style="display: flex; justify-content: space-between;"> File State: Approved 119 item(s) </div>								

On other pages, for example, the Computers page, the bottom of a table page shows the total number of items in the table and the number of pages in the table. It also provides page navigation buttons for moving between pages in the table and a menu for changing the number of rows displayed per page.

<input type="checkbox"/>	<input type="checkbox"/>	MYCORP\LAPTOP-2	●	Up to date	Up to date	High (Block Unapproved)
<input type="checkbox"/>	<input type="checkbox"/>	MYCORP\LAPTOP-3	●	Up to date	Up to date	High (Block Unapproved)
<input type="checkbox"/>	<input type="checkbox"/>	MYCORP\LAPTOP-4	●	Up to date	Up to date	High (Block Unapproved)
<input type="checkbox"/>	<input type="checkbox"/>	MYCORP\SERVER-1	●	Up to date	Up to date	High (Block Unapproved)

Navigation controls: 1 2 3 4 5 (highlighted with a red box)

963 items (highlighted with a red box) Page 3/39 (highlighted with a red box) 25 rows per page (highlighted with a red box)

With both scrolled and fixed-length views, if you request an extremely large table, the total number of items in the table (i.e., on all pages, not just the currently displayed page) will show as an approximation, such as *More than 10000 items*, and display the first page of the table. This allows the console to optimize page loading time and also indicates that you might want to request a table with a more manageable set of data. Consider modifying the view, for example, by changing the *Group By* choice, or sorting by a different column.

In rare cases, especially with a very large number of App Control Agents and/or an underpowered database server, requesting a table with an extremely large amount of data may cause the App Control Server to time out. Use the techniques mentioned above to reduce the data set.

Default and Saved Views

Each page and tab has a default view, which is unfiltered and shows data columns assumed to be most commonly of interest. To get exactly the view you want, you might modify several different table parameters. So that you do not have to recreate these modifications every time you view a page, the console allows you to *name* and *save* views on most pages. Once you have named a view, you can get to it again simply by choosing it on the Saved Views menu. When you choose **(none)** on the Saved View menu, you reset the page to the system default view.

ReadOnly accounts cannot create new Saved Views. They can access pre-configured Saved Views and those created by other users.

Most console pages come with pre-configured Saved Views in addition to **(none)**. Although you cannot change pre-configured views, you can use them as templates to create your own new Saved Views.

Files

File Catalog Files on Computers

Saved Views: (The Current View Has Unsaved Changes - Discard) Add Group By: (none) Max Age: (none)

(none) Applications by Publisher/Company Snapshot Export to CSV Refresh Table

Banned Files

Categorized Files

Installed Programs

Malicious Files

Threat Report - Suspicious Files by Name

Unapproved Files

Action Showing 75 out of 36218 item(s)


Select	Date Created	Computer	File Name	Publisher
<input type="checkbox"/>	Jun 22 2020 02:12:10 PM	● WORKGROUP\APC-99986	3d_501.dll	

To display a pre-configured Saved View:

1. Go to the page and tab (if any) you want to view.
2. In the Saved View panel, make a choice from the Saved Views menu. The view is displayed as soon as you release the mouse button.

Depending upon the view you choose, you might see different columns in the table, or only information matching a filter (for example, only files with status “Banned”).

In any view, including (*none*), you can make your own modifications via the filters and column controls, and also through a variety of other shortcuts on the page that let you set a time period, maximum number of items per page, and grouping. Once you modify a view from its original form, the Saved Views panels shows that you have unsaved changes until you either save the changes or reset the view to another Saved View. Changes to system-provided views must be saved to a different name.



To create and save a view of a console table:

1. Go to the page and tab (if any) you want to view.
2. If you want to start with an existing view as your template, choose that view from the Saved Views menu.
3. Use **Show Columns** to show the columns you want.
4. Use **Show Filter** to include or exclude items from the table.
5. To view only items newer than a particular date or time, use the Maximum Age menu. (You also can create more complex date/time filters on the Filters menu).
6. To show items listed by a group name rather than the item name, choose a Group from the *Group By* menu and choose the order in which you want them displayed (*Ascending* or *Descending*). For example, to group files by Publisher, choose Publisher. The table initially shows the groups, but if you click on a group name, it expands to show the individual items in that group.
7. On pages that show tables of files, if you want to see individual files installed by an installer rather than the installer file name only, click the *Show individual files* checkbox in the bottom right of the page.
8. If you want more or fewer rows displayed per page, choose a different number from the *rows per page* menu in the bottom right of the page. If you choose *page* in the right menu of this line, the change affects only the page you are on (e.g., only the Computers page). If you choose *all pages* in the right menu of this line, the change affects every page in the console for which you have not specified a length.
9. Once you have exactly the view you want, type a name representing this view into the right box in the Saved View panel and click the **Add** button. Your new view is now saved and available by name from the Saved Views menu.

Note that even if you do not create a Saved View, the console can remember the most recent view (filters and columns choices) for each page, so if you navigate away from the page and come back, you will see your most recent view until you make an alternate view

selection. Once you choose a different view, however, any changes to the current view are lost.

If you choose, you can set a user preference that does not remember your most recent view of a page, instead resetting to the console default view when you navigate away from a page. See [“Preference Settings for Console Users”](#) on page 83 for more details. Also, even if your preference is to remember changes, if you do not want any modifications remembered in a particular visit to the page, you can click on the **Discard** link next to the message about unsaved changes, and this returns the view to its saved format.

Exporting App Control Server Data to Files

The console file export tool downloads data to a file in comma-separated-value format. Downloaded data is presented according to the current column and filter configuration for online display.

If you download the file to a Windows system, it has a .CSV extension. On Mac systems using the Safari browser, the downloaded file has the standard CSV format but has a .CSV.XLS extension.

To download table data to a file:

1. Click **Export to CSV**. The standard download dialog box for your browser appears.
2. Follow the instructions presented in the dialog box to download the file:
 - a. Choose to open the file or save the file to disk.
 - b. If you save the file to disk, select a location and optionally rename the file.

Details Pages and Object Previews

In many console tables, you can get more details about the item in a row by clicking a View Details button or (if it is highlighted in blue) clicking the name of an object in the table. Details pages include:

- File Details pages
- Computer Details pages
- Publisher Details pages
- Certificate Details pages
- Device Details pages
- External Notification Details pages
- Indicator Set Details pages
- Approval Request Details pages

For example, clicking the details button next to a file name in the Files Catalog brings you to a File Details page, which shows more information about the file. See [Chapter 7, “File, Publisher, and Application Information”](#) for more on the file details available in the console.

Files

File Catalog | **Files on Computers**

Saved Views: (The Current View Has Unsaved Changes - Discard) | Group By: | Max Age: | Add | (none) | (none) | (none)

Show Filters | Show Columns | Show Snapshot | Export to CSV | Refresh Table

Action | Showing 100 out of 36218 item(s)

Select	Date Created
<input type="checkbox"/>	Jun 23 2020 05:33:47 PM
<input checked="" type="checkbox"/>	Jun 23 2020 05:33:47 PM
<input type="checkbox"/>	Jun 23 2020 05:33:48 PM
<input type="checkbox"/>	Jun 23 2020 05:33:50 PM
<input type="checkbox"/>	Jun 23 2020 06:02:35 PM
<input type="checkbox"/>	Jun 23 2020 05:33:52 PM
<input type="checkbox"/>	Jun 23 2020 06:05:41 PM

File Instance Details

Details for file on computer: WORKGROUP\APC-99986

File Name: sshd.exe
Date Created: Jun 23 2020 05:33:47 PM
File Path: c:\cygwin64\usr\sbin\
Computer: WORKGROUP\APC-99986
Platform: Windows
User Name: (none)
Local State: Approved
Local State Details: Locally Approved
Detached Publisher: (none)
Executed: Yes
Present At Initialization: Yes
Top-Level File: No
Deleted: No
Root File Name: (none)

General

First Seen Name: sshd.exe
First Seen Date: Jun 23 2020 05:34:14 PM
Last Updated: Jun 23 2020 05:34:14 PM
First Seen Path: c:\cygwin64\usr\sbin\
First Seen Computer: WORKGROUP\APC-99986
First Seen Platform: Windows
Extension: exe
Global State: Unapproved
Global State Details: File is unapproved, Publisher is not present, Certificate is not present
Flags: (none)
Installer / Updater: No
File Prevalence: File exists on 1 computer

File Properties

Publisher: (none)
Company: (none)
Product Name: (none)
Product Version: (none)
File Size: 731,155 bytes
Description: (none)
File Type: Application
SHA-256: E666018D4A22B5424385D375280A2718A3525E68CF1B448D4F7037BFA40C77EB
MD5: BB9DF523CCB9A89A56AAE1D06A62DE9
SHA-1: 8D6CE6E6F2D452DCEB392008FB80F369A8D3009

External Analysis Results

Check Point:
FireEye:
Microsoft SCEP:
Palo Alto Networks WildFire:

History

Jun 23 2020 05:34:14 PM The file appeared on WORKGROUP\APC-99986 during initialization

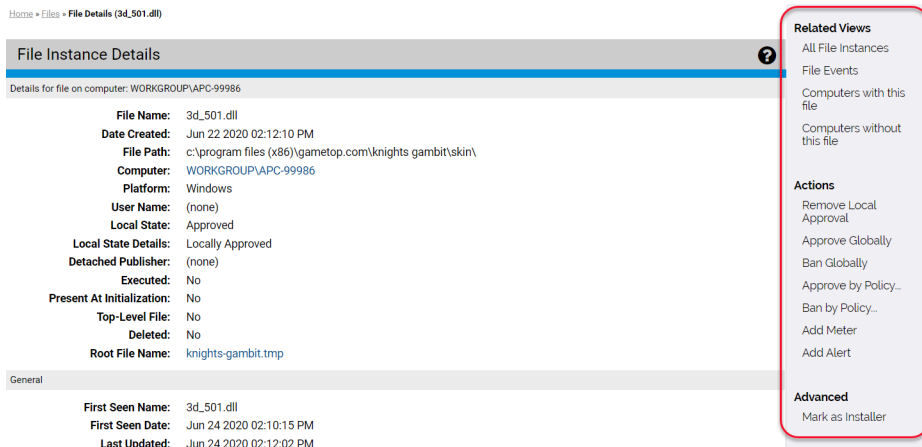
Note

Double-clicking a table row for objects with details pages, including the Files, Software Rules and Approval Requests tables, has the same effect as clicking the View Details button – that is, it opens the details page for that row.

Menus on Details Pages

Some console pages have menus to the right of the main content. These menus may include one or more of the following sections:

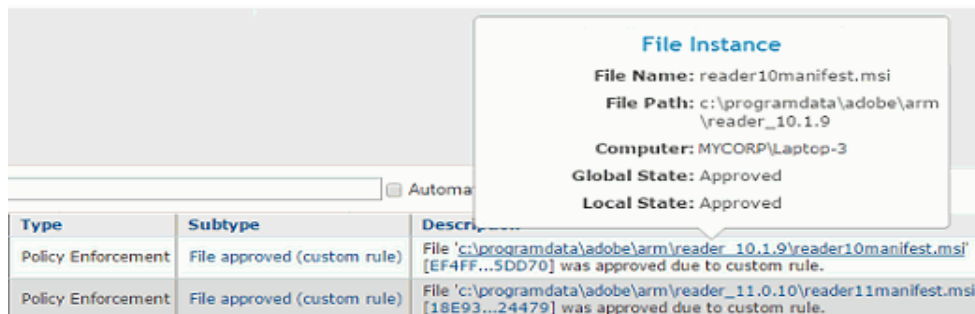
- **Related Views** links send you to pages related to the current page. For example, the File Instance Details page includes a link to a table of all computers with the file.
- **Actions** commands take actions related to the content of the page. For example, the File Instance Details page includes commands to ban or approve the current file.
- **Advanced** commands are less common or require consultation with Carbon Black Support for proper use.
- **External Pages** links are available if other products are integrated with the server and they are configured for direct links to their information from the console. For example, a Computer Analytics link could take you to the Splunk console.



Object Previews in Table Data

As the sections above have described, details pages provide a significant amount of information about objects cataloged in the App Control database, and one of the ways to get to details pages is to click on highlighted information in a table. In some cases, you might want more than the name of a highlighted object but not all of the information provided by its details page. Object previews provide summary information for many highlighted objects without requiring that you navigate away from the current page.

To see an object preview, move the mouse cursor over a highlighted item without clicking. For example, this is what a File Instance preview looks like when you move the cursor over a file name in the Events page.



The following items in a table have object previews (if they are highlighted):

- Files in catalog
- File Instances
- Certificates
- Computers
- Devices
- Publishers
- Policies
- Approval Requests

Shortcut Links

On many console pages, there are blue highlighted shortcut links that bring you to pages showing information related to the page you are on. For example, on the Computers page, clicking on a computer name takes you to the Computer Details page for that system while clicking on the policy name takes you to the Edit Policy page.

Select	Date Created	Computer	File Name	Publisher or Company
<input type="checkbox"/>	Jun 23 2020 05:33:47 PM	WORKGROUP\APC-99986	parity.exe	Carbon Black, Inc.
<input type="checkbox"/>	Jun 23 2020 05:41:08 PM	WORKGROUP\APC-99986	chrome.exe	Google LLC
<input type="checkbox"/>	Jun 23 2020 05:33:47 PM	WORKGROUP\APC-99986	vmtoolsd.exe	VMware, Inc.

On some pages, the link is a quick way to search for information that might otherwise require creation of a complex query on another page. For example, on the Edit Policy page, there is a link that shows you all computers in the policy.

Preference Settings for Console Users

The User Settings page allows each console user to change their password, the page they see first when they log in, and whether changes they make to page views are preserved when they navigate away and return to a page. To view the User Settings page, choose **loginname > User Settings** in the main menu.

Systems with Unified Management activated also see the following options:

Changes to the User Settings page apply to the currently logged in console user, and can be specified by any user, including those with ReadOnly access. Table 6 shows the effect of changes specified on this page.

Table 6: User Settings Page Choices

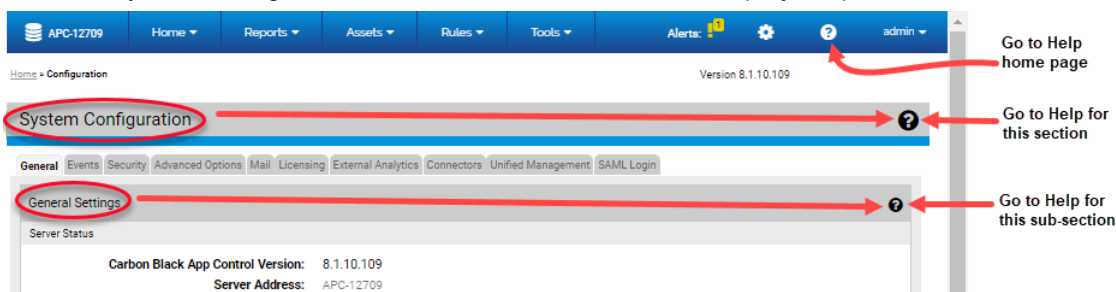
Panel:Field	Description
Change Password: Existing Password	Enter your current console password. Not available for accounts created through Active Directory.
Change Password: New Password	Enter the new password. Password requirements for length and special characters are specified by your administrator. Special characters include: > < & % () @ . [] { } ; ; ^ = ! ' " ` ~ , Not available for accounts created through Active Directory.
Change Password: Confirm Password	Enter the same password that you entered in the New Password field. It must be identical. Click Save when done. Not available for accounts created through Active Directory.

Panel:Field	Description
Display Preferences: Remember Page Settings	<p>Allows current user to choose whether page settings are saved (both within and between sessions). This setting applies to all console pages for the current user</p> <p><i>If checked</i>, all page configuration, including filters, columns, and group by settings, is remembered when you navigate away from a page (or logout) and come back to it.</p> <p><i>If not checked</i>, pages return to console defaults when you navigate away from them, and you lose any special layout you applied to them.</p> <p>In the Action menu, Reset Current Settings returns pages to their defaults without requiring you to un-check this box.</p>
Display Preferences: Resizable Table Columns	<p>Allows current user to enable or disable resizable table rows for console tables. Enabled by default. See “Table Column Resizing” on page 67 for more information.</p>
Display Preferences: Set Rows per Page	<p>Allows current user to set the standard number of rows per page to be shown on pages that display tables of information. When changed, this re-sets the number of rows on all console table pages. However, each user can customize the rows-per-page for an individual page after the overall preference is set. The default setting is 25. This setting does not affect pages the use dynamic scrolling.</p>
Display Preferences: Default Starting Page	<p>Allows current user to choose (from a menu) which console page appears first upon login. Choices are:</p> <ul style="list-style-type: none"> • Home Page • Events • Computers • File Catalog • Policies • Find Files • Approval Requests
Display Preferences: Unsaved Changes Warning	<p>When checked, displays a warning dialog when this user attempts to navigate away from a page with unsaved changes. The dialog allows the user to leave the page as requested or cancel the navigation to stay on the current page. When unchecked, there is no warning when this user navigates away from a page with unsaved changes.</p>
Display Preferences: Click to Dismiss Feedback Banners	<p>Controls how long feedback banners display and how they are dismissed.</p> <ul style="list-style-type: none"> • If unchecked: The banner is automatically dismissed after a short period of time. • If checked: The banner is dismissed only after the user clicks the banner.

Panel:Field	Description
Unified Server Authentication	This panel appears only when Unified Management is activated for the current server and the logged in user has permission to user or configure Unified Management. It shows whether the currently logged in user has been authenticated to access clients of this Unified Management server, and provides an interface for authenticating this user for each client server. See Chapter 27, “Unified Management of Multiple Servers,” for more details.
Save/Cancel buttons	Save saves the user settings changes. Cancel returns to the previous page the user was on, without saving the changes.

Using Online Help

The console provides access to an online version of the *App Control User Guide*. In addition to a link to the Help home page, most pages include a context-sensitive link that takes you to information relevant to your current view, but from which you can also navigate to other topics. When you click a Help link or button, a new Help window opens, either as a new tab in your current browser or as a new, popup browser. If it displays as a tab, you can drag the tab off of the current browser to display Help in its own window.



If you are using Microsoft Internet Explorer and it has popup blocking enabled, you must allow popup displays from the App Control Server if you want to view Help as a popup. Also, you might see a certificate error the first time you open Help – see [“Console Access”](#) on page 53 for information on accepting the certificate.

To display online documentation from the console:

1. Launch Help either of the following ways:
 - Click the orange question mark button on the right side of the **main menu** to open the home page and table of contents for App Control Help.
 - Click the black question mark button on the right side of the page title to see the context-sensitive help topic relevant to that page.

VMware Carbon Black App Control help is displayed in a new window or tab. The controls on the help page, and their location, vary depending upon the size of the window, but all pages provide access to the index, table on contents, and search features.

2. To view the table of contents if it is not visible, either expand to browser window to display the contents on the left or click the table of contents button.

- In the table of contents, click a right arrow icon or the name next to it in the table of contents to expand the table to show more subtopics. Click the down arrow icon to collapse the items below it.
- To view an alphabetic listing of topics, click the Index button.
 - To search for topics using key words, enter the words in the Search box if it is visible, or if not, click the Search button to display the field.

Carbon Black App Control User Guide and Help Center

From here you can access the online User Guide for your version of VMware Carbon Black App Control. Use the index, table of contents, or the new list of tasks to navigate through the User Guide. You also can use the search facility in the help window to search for information about a specific subject.

New in 8.5.0

- See the respective Release Notes on the Carbon Black [User eXchange](#) for a complete list of changes.

Context-Sensitive Help

In the VMWare Carbon Black App Control console, if you click on the black question mark button to the right of a page title, you get information related to that page.

System Configuration

General Settings

Access Other Documentation on the Carbon Black Site

Didn't find what you were looking for in the User Guide? You can expand your search to the Carbon Black community through the [User eXchange](#) (login required). The User eXchange provides access to other documentation for VMWare Carbon Black App Control.

Ask the Community

Still need more help? The [User eXchange](#) also provides a forum for product discussions and solutions, with participation by other VMWare Carbon Black App Control users and Carbon Black employees.

Notes

- Unless you close the Help tab or browser, each requested Help topic displays in the same window. However, security measures in Internet Explorer and Firefox may prevent an open Help window from coming to the front when you load new topics. In this case, click on the tab or use desktop navigation tools such as **Alt - Tab** to bring Help to the front of your display.
- A navigation anomaly in Chrome causes context-sensitive help pages to display the content immediately *below* the topic heading you requested (for example, the first paragraph in the topic). If you are uncertain that you are in the correct topic, scroll up to the heading.

Chapter 3

Managing Console Login Accounts

This chapter describes how to create and manage login accounts for the App Control Console. It also describes how to define user roles that grant access to specified features, and in some cases limit this access by policy.

Sections

Topic	Page
Login Account Management	89
User Roles and Permissions	89
Enabling Console Access via AD Accounts	92
Creating Login Accounts in the Console	95
Logging In Using SAML	97
Changing Passwords and Other Account Details	98
Deleting Login Accounts	100
Disabling Login Accounts	101
Managing Console User Roles	101
Creating a New User Role	102
User Role Permissions	106
Mapping AD Groups to Roles	111
Editing a User Role	110
Disabling a Role	115
Deleting a Role	116

Login Account Management

Each App Control Console user logs in to the system with a user name and password. Login Accounts provide system-management professionals, security team members, and others who use the console the ability to access and manage App Control features.

There is one built-in login account for the App Control Console, the *admin* account. It provides a way to log in to the console before other accounts are created, and it cannot be deleted. By default, this account has administrative privileges for nearly all features, and it can modify its own privileges. It also has the ability to create new accounts, and to define their privileges.

The first thing you should do when you log in as *admin* is change the password (also *admin* by default). See [“Changing Passwords and Other Account Details”](#) on page 98.

To create additional App Control Console accounts, you have two choices:

- You can create accounts individually through the console. These accounts are managed through the console, and can be modified or deleted by users whose login accounts have the proper privileges.
- You can permit users to log in using Active Directory credentials and map different AD groups to different privileges. AD-based App Control Console logins appear as “External Accounts.” For environments requiring the best security practices, Carbon Black recommends using AD-based accounts.

Although you can have a mix of AD-based and console-created login accounts, consider your preferred account management strategy before beginning to create new accounts. It is less confusing to generate all of your App Control Console accounts in the same way, either as AD-based accounts or as accounts created in the App Control Console.

Otherwise, although there will not be literal duplication of full account names, you could have names that appear to be the same. For example, you could have a console-created account name “fred” and also an AD-based account “fred@somedomain.”

Beginning with release 8.1.0, you have the option of using SAML to authenticate console users. See [“Configuring SAML Logins”](#) on page 769 for information about enabling this feature.

User Roles and Permissions

Each user of the App Control Console has one or more *user roles*. A user role is a collection of permissions, each of which allows the user to view specified information or manage specified actions in the console. Usually, these permissions map to specific pages in the console.

For a permission that involves information about or actions affecting agent-managed computers, a role can be configured to restrict permissions to specified policies.

You can create as many role- and policy-specific permission sets as you need. Once roles are created, you can assign or remove them as needed, giving each user just the permissions they need at any time. You can make these assignments manually or use AD-mapping for automatic role assignment.

For example, your organization might divide responsibility for IT support or security so that different people support different types of computers (desktops, servers, point-of-sale systems, etc.). On the other hand, you might divide responsibilities by region. With role-based access and policy definitions, you can configure App Control user accounts so that their privileges apply only to those computers they are responsible for. In addition, you can define the level of access each user has to App Control features, limiting some user to

viewing information while allowing others to create and modify rules, configurations, and other App Control resources.

For user accounts created in the console, roles are assigned on the Add Login Account page and can be changed on the Edit Login Account page. [Table 7](#) summarizes the default privileges for the built-in User Roles:

Table 7: Built-in User Roles and their default capabilities

User Role	Capabilities Summary
Administrator (Unified Management)	Access to all features. This is the only role that has permission to configure Unified Management; this permission cannot be added to any other role.
Administrator	Access to almost all features; does not enable permission to: <ul style="list-style-type: none"> • Manage uploads (any type) or access uploaded files • Extend connectors through API • View process command lines • Use (or configure) Unified Management Can add or remove privileges from any user, including itself.
PowerUser	Access to most features; does not enable permission to: <ul style="list-style-type: none"> • View process command lines • View (or manage) file uploads (or access uploaded files) • Manage system configuration • Manage login accounts (can view their own account) • Manage user roles and mappings • Extend connectors through API • Use (or configure) Unified Management
ReadOnly	View-only access to information on most table, report, and details pages; does not enable permission to: <ul style="list-style-type: none"> • View process command lines • View approval requests • View file uploads • View system configuration • View login accounts and user roles • View system health indicators • View certain advanced details on Computer Details page (Policy Override tab, CLI command) ReadOnly users can make the following modifications: <ul style="list-style-type: none"> • Can create personal dashboards with existing portlets only. • Can modify their own password and page view defaults through the User Settings interface.
User (Unified Management)	All permissions for a ReadOnly user plus can use Unified Management features.

Built-in user roles cannot be deleted, but the privileges of the Administrator, PowerUser and ReadOnly roles can be edited to enable or disable access to features. In addition, the roles themselves can be disabled.

Administrators can create new user roles with custom privileges (including the ability to create accounts and roles). See [“Managing Console User Roles”](#) on page 101 for instructions on creating user roles and customizing account privileges.

Upgrades from Previous Versions

If you have upgraded to App Control from pre-8.0.0 versions of Bit9 Platform or Parity, whatever permissions you provided to console users remain in place, as do any AD role mappings. However, there are significant enhancements in your control of user privileges, and these have resulted in changes to the user interface for managing users.

- **Roles instead of Groups** – In this release, instead of assigning a user to one group, you can assign one or more *roles* to the user. Roles are defined using the same list of permissions previously used for groups, but because you can assign more than one role to a user, you can define role-based permission sets, and add or remove them from users as needed.
- **Policy-Specific Permissions** – Another change in this release is that you can restrict access to certain features according to the policy a computer is in. For example, you might want to assign some management tasks to a group of users but only for computers used by the sales team. By putting all of the sales team’s computers in one policy, you can then set up a user role specific to that policy.
- **Users without Roles** – In previous releases, there was a named user group called Unauthorized. Users could be assigned to this group to indicate that they had no access to the console, and users were automatically assigned to the group if AD mapping was enabled and the user did not match the mapping rule for any other group. In this release, instead of being assigned to the Unauthorized group, users with no console access have no assigned *roles* (they appear in the table of accounts with “<Unassigned>” in the User Roles column).
- **Conversion of Previous AD Mapping Rules** – Although there are no default AD mapping rules for new installations of App Control, upgrades from previous releases convert the old mappings into new, role-based mappings that assign the same privileges that users had before.
- **Stop Evaluation of Mapping Rules** – Another AD-mapping-related change is the addition of a “stop evaluation” rule. Because you can assign more than one role to a user, evaluation of a user against AD mapping rules does not necessarily stop when a match is found. However, you might want to assign only one role to users matching certain AD characteristics, and prevent those users from having access to other features. In that case, you can put that rule at the top of the mapping rules list and check the *Stop evaluation* box. This box is checked for all of your pre-8.0.0 group assignments since there could be only one group per user.
- **View Effective User Permissions** – So that you can see the permissions a user has without having to click back and forth between the Edit User Account page and the (possibly multiple) Edit User Role pages, the Edit User Account page for each user now shows both the roles a user has and the user’s effective permissions. The effective permissions list is the combination of all permissions provided by all of the roles the user has.
- **Unified Management** – Unified Management of multiple servers is a new feature introduced in version 8.0.0, and new user roles have been added for it. Only the built-in admin account has these privileges immediately after upgrade. You can log in as admin to assign these privileges to other users, either individually or by changing your AD mapping rules.

Enabling Console Access via AD Accounts

If you use Active Directory and the App Control Server has been joined to an Active Directory domain, you can use AD accounts to log in to the console.

When a user logs into the console with an AD-based account name, that account is added as a App Control Console account. Users attempting to login to the console with a legitimate AD account but who are not members of a group that is not mapped to any role will be added to the console accounts table, but without any privileges. As such, they will not be able to login to the console.

You can map an AD account to multiple App Control Console roles, either intentionally or because the account happens to match more than one mapping rule. If you choose, you can stop evaluation when an account matches the highest ranking mapping rule in the list. See [“Managing Console User Roles”](#) on page 101 for more details.

Note

Unless you are using a Windows 2000 domain controller, you can specify a security domain separate from the login domain of your user accounts. This allows you to create App Control Console user roles in the named security domain rather than in the domain for each of your users.

To enable use of AD logins on the App Control console:

1. For each AD user account that you want to give console access, make sure you have assigned the account to a mapped AD security group.
2. Log in to the console as `admin` or any other administrator account you have created.
3. In the console menu, click on the configuration (gear) icon and choose **System Configuration**. The System Configuration page opens.
4. On the System Configuration page, click on the **General** tab. Initially, the settings on this page are grayed out.


Active Directory / LDAP integration

AD-Based Logins: Disabled

AD Security Domain: (none)

AD-Based Policy: Disabled

Windows 2000 DCs:

Test AD Connectivity:  Test

5. Examine the Active Directory/LDAP integration box. If AD-based logins already shows as *Enabled*, you do not have to make any changes and you can skip the remaining steps.
6. If AD-based logins shows a value of *Disabled*, click the **Edit** button at the bottom of the page to make the settings editable.
7. In the dropdown menu for AD-based Logins, choose **Enabled**.
8. If you are using Windows 2000 domain controllers, check the Windows 2000 DCs box. This notifies the App Control Server that cross-domain membership features are not available.

9. If you created the AD Security Groups for App Control in a domain other than the login domain for the users who will log in to the console, enter that domain in the AD security domain field. (This feature is not available if you are using Windows 2000 domain controllers).
10. Click the **Update** button, and when the Confirmation dialog appears, click **Yes**. You can now use Active Directory login accounts (if from one of the mapped groups) to access the console.

You disable the use of AD-based logins with the same procedure, except that you choose *Disabled* for the AD-based logins setting. If you disable AD-based logins, users will no longer be able to use their AD account names and passwords to access the console.

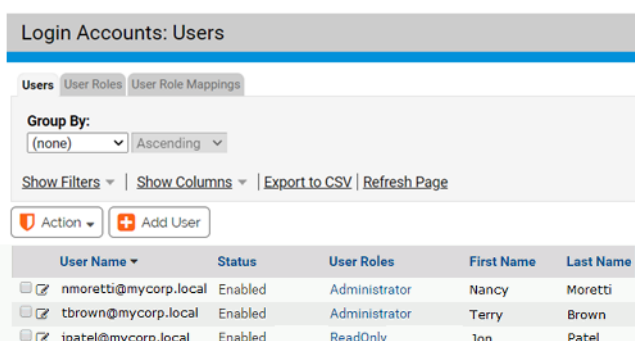
AD Login Account Format

The format for logging into the console with an Active Directory account name depends upon whether the account name is in the same domain as the App Control Server:

- AD accounts in a different domain must use a fully qualified version of their name (i.e., in the format *NTDOMAIN\Username* or *Username@dnsDomain*).
- AD accounts in the same domain as the App Control Server can log in either with a fully qualified username or their username only (provided the username is not the same as a login account created directly using the console).

There are several differences in the details for an AD-based account and an account created in the console:

- When a user with an AD-based account logs in to the console, the username on the Login Accounts page and the User Details page includes both the user and the domain name, in the form *user@dnsDomain*.



User Name	Status	User Roles	First Name	Last Name
<input checked="" type="checkbox"/> nmoretti@mycorp.local	Enabled	Administrator	Nancy	Moretti
<input checked="" type="checkbox"/> tbrown@mycorp.local	Enabled	Administrator	Terry	Brown
<input checked="" type="checkbox"/> jpatel@mycorp.local	Enabled	ReadOnly	Jon	Patel

- When you click on the View Details button to open the User Details page, the box at the top of the details panel is labeled “External Account” for AD users.
- There is no Save button on the Login Account Details page for AD users because their account details can’t be edited in the console.

Adding, Deleting, and Changing AD Login Accounts

The App Control Server stores user information for AD accounts that have logged in to the console, but re-validates that information for each login attempt. Any AD account changes that occur while that user is logged in to the console take place only after they log out and

log in again. Also, account updates depend upon how frequently the AD domain controllers on the network send out changes. Among the AD account changes that can affect console login accounts are:

- User accounts added to AD become available as console login accounts as long as they meet the security group and forest criteria.
- User accounts eliminated from AD can no longer be used to log in to the console.
- If there is a change in an AD-based user's security group assignment in AD, and if that AD change affects mapping of App Control user roles, the user's access level in the console changes when they next login.
- Other console User Details (contact information, etc.) for an AD-based user can be changed in AD and will appear when that user next logs in to the console.

Notes

- All of the AD-based login features depend on the App Control Server being able to communicate with the AD system and being in the Domain. If for some reason the App Control Server cannot communicate with the AD System (due to network setup change, network failure, AD system unavailable, etc.), AD-based Logins will stop working until the condition is corrected.
- AD-based login features also require that AD security groups are defined in each forest that contains users who will access the App Control Server; and that users you want to allow access to the App Control Server are added to the forest-specific security group.

Changing AD User Details Displayed in the Console

Whether an AD User has a console login account or not, anytime an AD user account appears in a table (other than the Login Accounts page) in the console, additional information can be displayed by clicking on that user name. For example, if you display the Events page, some events include the user associated with the event:

If the name is identified as an AD username, it is highlighted in blue, and when you click on it, a User Details window appears (note that this is not the same as the User Details page that appears when you click on a name on the Login Accounts page):

You can change, add, or remove fields from this page by editing the file `UserProps.txt`. This file is located in the "Scripts" subdirectory of the App Control Server installation directory. For example, if you accepted the default installation directory, it would be in `C:\Program Files\Bit9\Parity Server\Scripts`.

The file is a two-column, colon-separated list. The App Control label (for example, "Name") is on the left, and the AD property displayed for that field is on the right. Be sure to use actual AD object properties for the term on the right of the colon if you edit this file.

Similar customization can be done for AD-based computer details displayed in the console.

Creating Login Accounts in the Console

The following instructions are for creating individual login accounts through the console interface. If you want to use existing Active Directory accounts for Console access, see [“Enabling Console Access via AD Accounts”](#) on page 92.

Note

Login Accounts are for access to the App Control Console. A login account is not necessary (nor appropriate) for someone whose only relationship to App Control is as a user of a computer that has the App Control Agent installed.

Login Account creation privileges depend on user role:

- By default, Administrators can create any level of account.
- By default, PowerUsers and ReadOnly accounts cannot create new accounts.
- Custom user Roles have whatever account-creation privileges are shown for the *View login accounts and user roles*, *Manage login accounts* and *Manage user roles and mappings* settings on their Add/Edit User Role page.

To create a console login account:

1. From the console menu, click the configuration (gear) icon and choose **Login Accounts**. The Login Accounts page displays.
2. Click on the **Users** tab.
3. On the Login Accounts: Users page, click **Add User**.
4. From the Add Login Account page, enter information about the new account in the categories shown in [Table 8](#).
5. After you have filled out the form, click the **Add User** button at the bottom of the page.

Table 8: Login Account Details Fields


Field	Description
User Name (required)	<p>Name that the user enters to log in to the console.</p> <p>Enter any combination of letters, numbers, or English-keyboard characters fewer than 32 characters in length. User names are not case sensitive.</p> <p>Note: User names should use standard, Latin alphanumeric characters. Symbols and punctuation characters are not allowed. In particular, be aware that user names created in the console cannot contain the “\” or “@” characters. This helps avoid conflicts with AD-based user names using <i>user@domain</i> or <i>domain\user</i> format. If you attempt to create a user account with an illegal character, the console will display a warning dialog.</p>

Field	Description
Password (required)	Password that authenticates this user. Enter any combination of letters, numbers, or English-keyboard characters fewer than 32 characters in length. Passwords are case sensitive. This field changes to New Password when you are editing existing accounts.
Confirm password (required)	Confirm password. Retyping the password ensures that the password is the one you intended to use.
Email address	Email address for the user.
User Roles	System privileges to be accorded to this user, according to the user's expected responsibilities. There are four built-in roles. You also can create custom roles with detailed feature-based access control – see “Managing Console User Roles” on page 101 for details. The built-in role options are: <ul style="list-style-type: none"> • Administrator • PowerUser • ReadOnly • Administrator (Unified Management) • User (Unified Management) See Table 7, “Built-in User Roles and their default capabilities” on page 90 for details.
Salutation	Courtesy or professional title of the user (Mr., Ms., Dr., etc.)
First name	First name of the user.
Last name	Last name of the user.
Title	Job title of the user.
Department	Group within the organization to which this user belongs.
Home phone	The user's phone number at home.
Cell phone	Primary mobile phone number.
Cell phone #2	Secondary mobile phone number.
Pager	Primary pager number.
Pager #2	Secondary pager number.
Comments	Further descriptive information that the user can change or enter. This can be any text you would like to display as part of the login account.
Admin comments	Further administrative information about the user. This can be any text you would like to display as part of the login account.
Show API Token	If you check this box, an interface is exposed that allows generation of an API Token for the current user account. It is best to create a special user account for this purpose. See “API Authentication and Access Control” on page 840 for details.

Logging In Using SAML

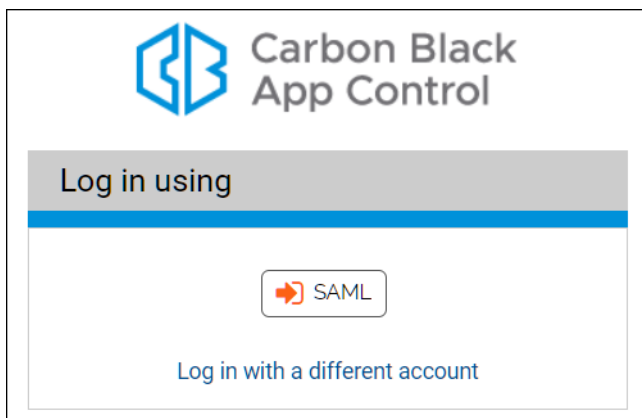
The App Control console can be integrated with identity providers (IdPs) that use the Security Assertion Markup Language (SAML). This allows you to require two-factor authentication (2FA) for logging in to the App Control console, for compliance purposes or to meet your own best practice standards. “[Configuring SAML Logins](#)” on page 769 provides instructions for setting up this integration.

Once you save a properly configured IdP in App Control, you will be able to use SAML to log in to an existing App Control account. A new button with your identity provider name appears on the App Control login page.



When a user clicks the button, they are directed to the login page of the identity provider. If they provide their correct credentials in the IdP, they are logged into the App Control Console as the user whose email address matches the one for the IdP account used.

When SAML is enabled, you can configure some user roles to allow local logins. Users that have one of these roles have the option of clicking **Log in with a different account** and entering their App Control credentials to access the console. You can enable local logins in either the System Configuration page for SAML Logins or in the Permissions table for a specific role.



A user who attempts a local login without having a user role that allows this will see a login error that requests a valid user name and password. The user can click the **Log in with SAML** button to go to their IdP for authentication.

Notes

- The IdP account you log in with must have the same email address as a user in App Control, provided either through the IdP NameID attribute or an attribute named EmailAddress (always used if provided). The identities are not matched by name.
- If the email address for the IdP account a user logs in with does not match the email address of any App Control user, the login will fail.

Changing Passwords and Other Account Details

When you initially log in to the console as “*admin*”, you should change the default password (also “*admin*”) to something unique. All users with login accounts, including *admin*, should change their passwords periodically.

For Active Directory-based accounts, password changes and other account information must be changed in Active Directory – they cannot be edited through the console.

For a login account *created in* the console:

- By default, accounts in the Administrators role may change passwords, contact information, and roles for any console-created account. Note that the role for the account *admin* may not be changed.
- By default, accounts in the PowerUsers role may change passwords and contact information for their own account.
- Account-editing privileges in custom roles vary.

Note

This section describes the Login Accounts administrative interface for changing account details. There is a more limited interface, the User Settings page, on which each account user, including ReadOnly users, can make certain changes to *their own* account only, including changing their password. See [“Preference Settings for Console Users”](#) on page 83 for details.

To change a console password and other login account details:

1. From the console menu bar, click the configuration (gear) icon and choose **Login Accounts**. The Login Accounts page appears:

User Name	Status	User Roles	First Name	Last Name
<input checked="" type="checkbox"/> nmoretti@mycorp.local	Enabled	Administrator	Nancy	Moretti
<input checked="" type="checkbox"/> tbrown@mycorp.local	Enabled	Administrator	Terry	Brown
<input checked="" type="checkbox"/> jpatel@mycorp.local	Enabled	ReadOnly	Jon	Patel

2. If the Login Accounts: Users page is not displayed, click on the **Users** tab.
3. On the Login Accounts page, locate the account of the user whose password you are changing, in the Login Accounts: Users table.
4. In the far left column next to the Username, click the View Details icon. The Edit Account Details page opens (see Table 8, “Login Account Details Fields”, for a description of the fields).
5. On the Edit Login Account Details page:
 - a. In the New Password field, enter the new password.
 - b. In the Confirm Password field, enter the password again to confirm it.
 - c. If the password is temporary and you want to force the user to reset it upon logging in, select, **Make Password Temporary**.
 - d. Optionally, change other Login Account Details.
 - e. Click the **Save** button.

Notes

- If the top box on the Login Account Details page for a user is labeled “External Account,” the user accesses the console with an Active Directory account and the details cannot be edited. If an account shows “Account” (without “External”) in the title for the top box, you can edit that account.
- If you change another user’s password, be sure to inform them of the change.
- If you are using SAML to authenticate users, the email address for a user in the Login Account Details must match the identity provider’s email address for that user. Keep this in mind before editing an email address.

Deleting Login Accounts

Login accounts can be removed from the system, for example, when an employee no longer needs access to the console or leaves the company. Console users can delete any account type they are allowed to create:

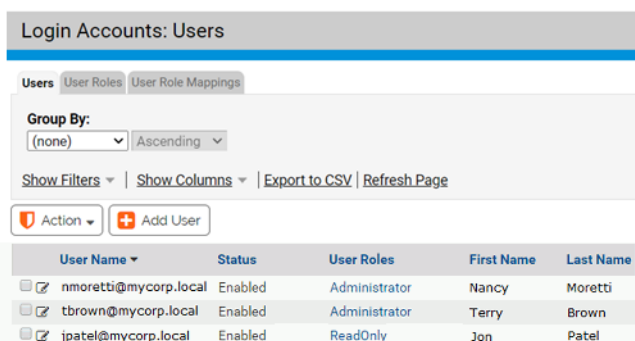
- By default, accounts in the Administrators role can delete any account except their own.
- By default, accounts in the PowerUsers role can delete ReadOnly accounts but not PowerUsers or Administrators.
- Account-deletion privileges of accounts in custom roles vary.

Note

You cannot delete the default *admin* administration account.

To delete a login account:

1. From the console menu bar, click the configuration (gear) icon and choose **Login Accounts**. The Login Accounts page appears:



Login Accounts: Users

Users | User Roles | User Role Mappings

Group By: (none) | Ascending

Show Filters | Show Columns | Export to CSV | Refresh Page

Action | Add User

User Name	Status	User Roles	First Name	Last Name
<input type="checkbox"/> nmoretti@mycorp.local	Enabled	Administrator	Nancy	Moretti
<input type="checkbox"/> tbrown@mycorp.local	Enabled	Administrator	Terry	Brown
<input type="checkbox"/> jpatel@mycorp.local	Enabled	ReadOnly	Jon	Patel

2. If the Login Accounts: Users page is not displayed, click on the **Users** tab.
3. In the Login Accounts: Users table, locate the username.
4. In the far left column next to the user name, click the Delete icon.
5. Respond to the confirmation prompt. To delete the account, click **OK**.

Disabling Login Accounts

When a user no longer needs access to the console you can remove access without deleting the login account. You do this by changing the account status to Disabled. Users permitted to *create* a particular login account can also disable that account:

- By default, accounts in the Administrators role can disable any account except their own.
- Account-disabling privileges of accounts in custom roles vary.

Note

Console login accounts created through AD mapping cannot be disabled directly. The only way to disable an AD account is to change the mapping rules for their AD security group so that they are mapped to user role that is disabled or has no privileges.

To disable a login account:

1. From the console menu bar, click on the configuration (gear) icon and choose **Login Accounts**. The Login Accounts page appears:

User Name	Status	User Roles	First Name	Last Name
nmoretti@mycorp.local	Enabled	Administrator	Nancy	Moretti
tbrown@mycorp.local	Enabled	Administrator	Terry	Brown
jpatel@mycorp.local	Enabled	ReadOnly	Jon	Patel

2. If the Login Accounts: Users page is not displayed, click on the **Users** tab.
3. In the Login Accounts: Users table, locate the username.
4. Click the View Details icon next to the username whose account you want to disable.
5. In the Status line on the Edit Login Account page, click the **Disabled** radio button.
6. Click the **Save** button at the bottom of the page.

Managing Console User Roles

The capabilities of a console login account are determined by its user roles. A user with permission to manage console user roles can perform the following tasks:

- Create new user roles with custom permissions.
- Restrict certain permissions to computers in specified policies.
- Disable a user role (except for the built-in Administrator role).
- Edit existing user roles.
- Delete any custom-created user role (but not any built-in role).

- Change the mapping of AD security groups to console user roles and the order in which mapping rules are evaluated.

One login account can have multiple user roles.

You can view the current user roles on the Login Accounts: User Roles page. This page is also the place from which you access other role management features.

To view the Login Account: User Roles page:

1. From the console menu bar, click on the configuration (gear) icon and choose **Login Accounts**.
2. If the Login Accounts: User Roles page is not displayed, click on the **User Roles** tab.

The screenshot shows the 'Login Accounts: User Roles' page. At the top, there are tabs for 'Users', 'User Roles', and 'User Role Mappings'. Below the tabs, there is a 'Group By' section with a dropdown menu set to '(none)' and a sort order dropdown set to 'Ascending'. There are also links for 'Show Filters', 'Show Columns', 'Export to CSV', and 'Refresh Page'. Below this is a button labeled '+ Add User Role'. The main content is a table with the following data:

Name	Status	Policy	AD Assigned	Manually Assigned
Administrator	Enabled	All Policies	0	1
Administrator (Unified Management)	Enabled	All Policies	0	1
PowerUser	Enabled	All Policies	0	0
ReadOnly	Enabled	All Policies	0	0
User (Unified Management)	Enabled	All Policies	0	0

Creating a New User Role

Although the built-in user roles provide options for different levels of feature access, users with sufficient permission can create new custom user roles and modify roles. You might want to have a special user role whose level of access falls between two of the built-in options. Creating a special user role can not only prevent unauthorized access to critical features but also might make it easier for users with limited roles to learn those roles without having to see features they will not use.

For example, you might want to allow members of a help desk team to view all information available through the console but only to be able to change policy for a computer, put a computer into local approval, or access debugging features. You can create a user role with these characteristics. [Table 9](#) shows the information used to define a user role.

Table 9: User Role Parameters

Field	Description
Name (required)	Name that will appear in the Login Accounts: User Roles list and will be used when assigning a role to a login account. Enter any combination of letters, numbers, or English-keyboard characters fewer than 32 characters in length. Role names are not case sensitive. Note: User names created in the console cannot contain the “\” or “@” characters. This helps avoid conflicts with AD-based user names using <i>user@domain</i> or <i>domain\user</i> format.
Description	Optional descriptive information about this role, such as who should be in it and perhaps a high-level summary of its permissions.
AD Mapping Name	If AD-based login mapping is enabled, the AD security group that you would like mapped to this role.
Status	Determines whether this role is Enabled or Disabled. Note that disabling a role disables it for accounts it is assigned to and prevents AD-mapping from matching it. If you disable the only role assigned to a user, that user loses console access.
Permissions	A table of checkboxes that determine what users with this role are allowed to do in the console. See Table 10, “Permissions Settings for User Roles,” on page 106 for a complete description.

To create a new console user role:

1. In the console menu, click the configuration (gear) icon and choose **Login Accounts**. The Login Accounts page appears.
2. Click on the **User Roles** tab.

Login Accounts: User Roles					
Users User Roles User Role Mappings					
Group By: (none) Ascending					
Show Filters Show Columns Export to CSV Refresh Page					
+ Add User Role					
Name	Status	Policy	AD Assigned	Manually Assigned	
Administrator	Enabled	All Policies	0	1	
Administrator (Unified Management)	Enabled	All Policies	0	1	
PowerUser	Enabled	All Policies	0	0	
ReadOnly	Enabled	All Policies	0	0	
User (Unified Management)	Enabled	All Policies	0	0	

3. On the Login Accounts: User Roles page, click the **Add User Role** button. The Add User Role page appears.

Add User Role ?

General

Copy Settings From: (none) v

Name:

Description:

Status: Enabled Disabled

Permissions

Asset	Permission	Scope	<input type="checkbox"/> Enabled
Computers	View computers	Policy	<input type="checkbox"/>
Computers	Temporary assign computers	Policy	<input type="checkbox"/>
Computers	Manage computers	Policy	<input type="checkbox"/>
Computers	Change advanced options	Policy	<input type="checkbox"/>

4. Enter a name for the new role, and optionally, a description to specify the purpose of the role, intended members, or any other information about the role.
5. Assuming you want this role to be available immediately for login accounts, leave the Status radio button set to Enabled.
6. Check the box next to each permission you want to enable for this role, and un-check any permissions you do not want this role to have. See [Table 10](#) for a complete list of permissions.

Note: If you are giving this role permission to perform most console activities, it might be more efficient to click the Enabled box in the table header, which checks all boxes, and then remove the few permissions you *don't* want to provide.

Add User Role ?

General

Copy Settings From: (none) v

Name: Help Desk

Description: Necessary permissions for all Help Desk personnel for all regions

Status: Enabled Disabled

Permissions

Asset	Permission	Scope	<input type="checkbox"/> Enabled
Computers	View computers	Policy	<input checked="" type="checkbox"/>
Computers	Temporary assign computers	Policy	<input checked="" type="checkbox"/>
Computers	Manage computers	Policy	<input checked="" type="checkbox"/>
Computers	Change advanced options	Policy	<input checked="" type="checkbox"/>

7. If you have AD account mapping enabled and want to automatically map members of an AD security group to this console role, put the name of the AD security group in the AD Mapping Name box.

8. If you want to limit this role so that it has access to computers in certain policies only, scroll to the bottom of the page, click the **Selected Policies** radio button, and check the box next to each policy you want this role to have access to.

Scope of Policy Permissions

All Current and Future policies
 Selected policies

<input type="checkbox"/>	Policy
<input checked="" type="checkbox"/>	APC-99986-ucdjfp
<input type="checkbox"/>	Default Policy

9. When you have finished configuring this role, click **Create & Exit** at the bottom of the page. The new role appears in the Login Accounts: User Roles table. Notice that it includes a delete button since, unlike a built-in role, a user-created role can be deleted.

Login Accounts: User Roles

Users **User Roles** User Role Mappings

Group By: (none) Ascending

[Show Filters](#) |
 [Show Columns](#) |
 [Export to CSV](#) |
 [Refresh Page](#)

	Name ▲	Status	Policy
	Administrator	Enabled	All Policies
	Administrator (Unified Management)	Enabled	All Policies
	Help Desk	Enabled	APC-99986-ucdjfp
	PowerUser	Enabled	All Policies
	ReadOnly	Enabled	All Policies
	User (Unified Management)	Enabled	All Policies

10. If you have AD mapping enabled, a new role is first in the mapping rank. Rank is significant if you have rules with *Stop evaluation* checked since no rules ranked lower than a Stop evaluation rule will be processed. If you want the new role to rank lower, use the arrow keys in the AD Rank column to move it down in rank, or to move another role up. Accounts that match multiple mappings are assigned all roles they match in rank order (beginning with number 1), until and unless they reach a Stop evaluation rule.
11. If you are not using AD mapping to assign console login accounts, manually assign this new role to user accounts you want to have its permissions.

User Role Permissions

On the Add/Edit Role page, the Permissions table shows the capabilities that can be enabled or disabled for members of the role – items that are checked are enabled and items that are not checked are disabled. You can customize permissions to achieve exactly the level of access you want for a role.

For the most part, permissions can be divided into two categories: *view* permissions that allow you to see a particular page or dialog in the console, and *manage* permissions that allow you to create, edit, and delete managed assets, rules, and console users. Some permissions depend on others – you cannot manage something unless you can see it. If you disable *View system configuration*, for example, *Manage system configuration* is automatically disabled as well.

Checkboxes for permissions that depend upon other permissions are gray (instead of white) when they are not enabled. In addition, permissions that depend upon other permissions are indented to make the relationship between them clearer.

The Scope column indicates whether a permission is global or policy-specific. Policy-specific permissions are affected by your choice in the Scope of Policy Permissions section of the Add/Edit User Role page.

Notes

- Carefully consider any permissions changes you make, especially to the built-in Administrator role. In particular, avoid removing permissions to view and manage user accounts and roles from the Administrator role since this will make it impossible to restore access to these features without the use of special recovery commands.
- The console user interface, including pages, menus and links, is documented based on users having the full administrative permissions. Any permissions that are turned off will remove related user interface elements. Consider making users with restricted permissions aware of this possibility so that they are not confused by the absence of features described in App Control help.

Table 10: Permissions Settings for User Roles

Asset	Permission Name	Scope	Description
Computers	View computers	Policy	Ability to view computer pages
Computers	Temporary assign computers	Policy	Ability to generate temporary Enforcement Level override codes. Requires View computers permission.
Computers	Manage computers	Policy	Ability to manually assign computers to policies and change Enforcement Level. Ability to manage template computers.

Asset	Permission Name	Scope	Description
Computers	Change advanced options	Policy	Ability to change advanced computer options such as collection of computer diagnostics and re-synchronizing.
Files	View files and applications	Policy	Ability to view files and applications pages.
Files	Manage files	Policy	Ability to approve, ban, and acknowledge files. Ability to mark files as installers. Note that this does not include the ability to directly change local file state.
Files	Change local state	Policy	Ability to change local state of files on computers.
Files	Delete files	Policy	Ability to delete files on computers.
Devices	View devices	Policy	Ability to view device pages.
Devices	Manage device rules	Policy	Ability to manage device rules.
Policies	View policies	Global	Ability to view Policies page.
Policies	Manage policies	Policy	Ability to manage policies (changing mode, Enforcement Level, etc.)
Policies	Manage policy mappings	Global	Ability to manage automatic policy mapping rules.
Software Rules	View software rules pages	Global	Ability to view Software Rules pages. Also allows viewing of Event Rules page for servers licensed for App Control Connectors for Network Security Devices.
Software Rules	Manage event rules	Global	Ability to manage event rules. Requires separate license for the App Control Connectors for Network Security Devices. Note: Some event rules require other permissions for the actions they specify, such as file upload and analysis and file approval.
Software Rules	Manage trusted directories	Global	Ability to manage trusted directories.
Software Rules	Manage publisher rules	Policy	Ability to manage trusted publishers.
Software Rules	Manage trusted users	Global	Ability to manage trusted users.

Asset	Permission Name	Scope	Description
Software Rules	Manage custom/registry/memory rules	Policy	Ability to manage custom, registry and memory rules.
Software Rules	Manage application updaters and configurations	Global	Ability to enable, disable, add, and view details of software updaters and configurations for applications, and to modify configurations.
Software Rules	Manage custom scripts	Global	Ability to manage custom definitions of what the App Control Server treats as scripts
Software Rules	Manage indicator sets	Policy	Ability to enable, disable, and create exceptions for indicator sets used in advanced detection
Reports	View events	Policy	Ability to view event pages.
Reports	View server events	Global	Ability to view server events.
Reports	View process command lines	Global	Ability to view process command lines for events. Important: Command lines may include confidential information such as passwords. This permission is not enabled by default, even for administrator accounts, and should be limited to those who require it.
Reports	Manage shared dashboards	Global	Ability to manage shared dashboards.
Reports	View drift reports and snapshots	Global	Ability to view snapshots, drift reports and drift report results.
Reports	Manage drift reports	Global	Ability to manage baseline drift reports.
Reports	Manage snapshots	Global	Ability to manage snapshots used in drift reports.
Reports	Manage saved views	Global	Ability to manage saved views on all pages.
Tools	View alerts	Global	Ability to view alert pages.
Tools	Manage alerts	Global	Ability to manage alerts.
Tools	View meters	Global	Ability to view meters and meter results.
Tools	Manage meters	Global	Ability to manage meters.

Asset	Permission Name	Scope	Description
Tools	View approval requests	Policy	Ability to view user-generated requests for approval of blocked files and justifications of files approved by users.
Tools	Manage approval requests	Policy	Ability to manage user-generated requests for approval of blocked files and justifications of files approved by users.
Tools	View file uploads	Global	Ability to view uploaded files on the Requested Files page.
Tools	Manage uploads of inventoried files	Global	Ability to initiate manual file uploads from agent computers, and to create event rules that upload files. This permission applies only to files considered "interesting" (i.e., executables and scripts) by App Control. Requires separate license for File Uploads.
Tools	Manage uploads of files by pathname	Global	Ability to initiate manual file uploads from agent computers, and to create event rules that upload files. This permission applies to <i>all</i> files on agent computers, even if not in the App Control inventory. Requires separate license for File Uploads.
Tools	Access uploaded files	Global	Ability to download files that are uploaded on the server. Requires separate license for File Uploads.
Tools	Submit files for analysis	Global	Ability to submit files for analysis by network security devices, either manually or through creation of an event rule. Requires separate license for the App Control Connectors for Network Security Devices, unless implemented through the API.
Notifiers	View notifiers	Global	Ability to view the details of blocked file notifiers.
Notifiers	Manage notifiers	Global	Ability to edit blocked file notifiers or create new ones.
Analytics	View external analytics reports	Global	Ability to view and use links from the console to external analytics reports (if external analytics is enabled and configured)

Asset	Permission Name	Scope	Description
Administration	View system configuration	Global	Ability to view system configuration pages.
Administration	Manage system configuration	Global	Ability to manage system configuration; this includes uploading agent and rule packages to the server.
Administration	View login accounts and user roles	Global	Ability to view login accounts and user roles for accounts.
Administration	Manage login accounts	Global	Ability to manage login accounts.
Administration	Manage user roles and mappings	Global	Ability to manage user roles.
Administration	Local login override	Global	Ability to login with a local (App Control) account when SAML logins are enabled. You can also enable this feature on the System Configuration page for SAML logins.
Administration	View System Health Indicators	Global	Ability to view the system health page and system health alerts.
Administration	Extend connectors through API	Global	Ability to register and unregister connectors with the App Control Server through APIs so that they can send notifications and (if part of their feature set) analyze files.
Administration	Use Unified Management	Global	Ability to use Unified Management features on multiple servers.
Administration	Configure Unified Management	Global	Ability to configure Unified Management (enable and disable, add and delete servers). This permission is built in to the Administrator (Unified Management) role, and cannot be added to any other role.

Editing a User Role

You can edit a console user role in the following ways:

- You can add and subtract permissions at the feature level for the built-in console user roles, and for any custom role shown on the Login Accounts: User Roles tab.
- If you have AD mapping enabled, you can change the AD security group that is mapped to a console login user role.
- You can change the policies this role provides permissions for.

- You can enable a user role, activating the ability of user accounts assigned this role to access the console, or you can disable the role, removing the permissions it provided to the user (unless another role assigned to the user provides the same permissions).
- You can edit the optional Description for a group.

To change permissions or other properties of a console user role:

1. In the console menu, click the configuration (gear) icon and choose **Login Accounts**. The Login Accounts page appears.
2. Click on the **User Roles** tab.
3. On the Login Accounts: User Roles page, click the View Details button for the user role whose privileges you want to change. The Edit user Role page appears.

Edit User Role ?

General

Name: PowerUser

Description:

Status: Enabled Disabled

Permissions

Asset	Permission	Scope	<input type="checkbox"/> Enabled
Computers	View computers	Policy	<input checked="" type="checkbox"/>
Computers	Temporary assign computers	Policy	<input checked="" type="checkbox"/>
Computers	Manage computers	Policy	<input checked="" type="checkbox"/>
Computers	Change advanced options	Policy	<input checked="" type="checkbox"/>

4. On the Edit User Role page, review the current permissions. Permissions with checkmarks in the right column are enabled; permissions with an empty checkbox are disabled. Click the checkbox for any capabilities whose status you want to change.
5. If you want to change the policies to which the permissions for this role are granted, make those changes in the Scope of Policy Permissions panel.
6. Review the Description field and make any changes to this information.
7. Click the **Save** button at the bottom of the page to save your changes.

Mapping AD Groups to Roles

When AD integration is enabled, the User Roles tab shows the AD mapping and AD Rank of console user roles. Rank determines the order in which AD mapping rules are evaluated, which is significant if an AD security group would match more than one mapping rule and one of the rules is configured to stop evaluation of other rules. You can change rank using the arrow keys on the Login Accounts: User Roles page.

Each time a console user logs in, App Control evaluates the user against the AD mapping rules. Mappings will be executed in the order they are ranked on the User Roles page. Each mapping rule can have one of two modes:

- **Assign role and continue evaluation** - If a user matches the mapping conditions, the mapped role is assigned to the user, and the next mapping rule will be evaluated.
- **Assign role and stop evaluation** - If a user matches the mapping conditions of a stop evaluation rule, that role will be assigned to the user and system will stop further evaluation. The rank of a rule is a significant factor in the effectiveness of a stop evaluation rule.

Mapping rules that stop evaluation are useful in two cases:

- If you upgrade from previous versions of App Control (Bit9 Server), there are stop evaluation mapping rules that convert each of the previous *group* mappings to a single *role* mapping with the same permissions.
- They allow you to deny access to lower ranked rules. For example, if you have temporary workers and assign them to an AD group called “contractors,” you can create a mapping to this group, enable Stop evaluation on that mapping, and rank the rule #1 so that it is evaluated first and stops any further role assignments. Without enabling Stop evaluation, a contractor might match some other, lower-ranked rule, and gain permissions that you do not want to grant.

AD Role Mapping Summary

To make use of AD-based user role assignment, you must:

- **Install App Control in an AD Domain** – Install the App Control Server on a computer that is a member of an Active Directory domain. By default, the App Control Server must be in the same AD forest as the computers and users you want to map. If you require cross-forest integration, contact your Carbon Black Support representative.
- **Enable the AD Mapping Interface** – You enable the AD-based user mapping interface in the Active Directory / LDAP integration panel on the General tab of the System Configuration page. See [“Enabling Console Access via AD Accounts”](#) on page 92 if you have not already done this.
- **Create AD-mappable Target Roles** – Create the user roles to which you want computers assigned by AD Mapping.
- **Create Mappings** – On the Mappings tab of the Login Accounts page, create AD Role Mapping rules that use AD data to assign computers to different security policies.

Platform Note: The App Control Server will do AD-mapping for any computer you have configured through your Active Directory server, including those on non-Windows platforms.

Creating AD Mapping Rules

After the AD-based User Roles interface is enabled, a new tab, “Mappings,” will be visible when you view the Login Accounts page. Clicking on this tab opens the Active Directory User Role Mappings page. This is where you create rules to map computers with specified AD data to certain roles.

You can create mapping rules that test for matching AD data including organizational units, domains, security groups, computer names, and user names. Keep the following in mind when creating mapping rules:

- App Control does not support policy mapping for AD object names that contain double quotes. Object names with double quotes cannot be handled properly by the directory object browser you use to create a mapping rule.

- In general, create as few rules as possible and use them to test for groups rather than individual objects.

Table 11 shows the rule parameters you provide for a mapping rule.

Table 11: AD User Role Mapping Rule Parameters

Parameter	Description
Object to Test	The object that will be tested to see whether it matches the rule. This is always "user".
Relationship	The relationship being evaluated between the Directory Object specified in the rule and the AD data from the user attempting to log in. The choices are: <ul style="list-style-type: none"> • is member of group • is in OU or domain • is • is not in any domain
Directory Object	The object in AD that the data from the tested object must match. Clicking the right end of this field opens an AD browser from which you can search for an object from your AD environment. The choices for the Directory object field change depending upon which Relationship you choose. If you choose "is not in any domain," no Directory object is necessary. The object browser for choosing the directory object is similar to the one for AD policy mapping. See "AD Object Browser Options" on page 128 for a more detailed description of this browser.
User Role to Apply	The role to apply to a computer if its tested object matches the rule. The dropdown menu shows all available roles.
Stop Evaluation	Checking this box causes evaluation of users against mappings to stop when this rule is reached. If this is the first rule that matches a user, the role it assigns will be the only role they have. Leaving the box unchecked allows evaluation to continue to the next ranked rule.

The result of providing these parameters is a rule that can be read like a sentence. The following is how you might set up one rule.

Parameter	Example (value in bold)
Computer Object to Test	If a User ...
Relationship	... is in OU or domain ...
Directory Object	...matching OU = Support,DC=hq,DC=xyzcorp,DC=local ...
User Role to Apply	... assign the Help Desk role to the user...
Stop Evaluation	... and do not assign roles from any lower-ranked rules, even if this user matches their conditions.

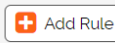
The procedure below shows how to configure a mapping rule. Although entry of most of the parameters are reasonably straightforward, pay particular attention to the Directory Object field, which requires use of a special AD browser.


To create an AD role mapping rule:

1. In the console menu, click the configuration (gear) icon and choose **Login Accounts**.
2. Click the **User Role Mappings** tab. The Active Directory Mappings page appears with the User Role Mappings table. If you have upgraded from a pre-8.0 version of App Control (Bit9 Server) and were using AD group mapping there, a series of default mappings appear. Otherwise the table only has a default mapping to no role.

Login Accounts: Active Directory Mappings


Users User Roles **User Role Mappings**

 Add Rule

Object	Relationship	Match	Action	User Role to Add
 [all others]			add User Role and stop evaluation	(None)

Note: If no Mapping tab appears, the AD mapping interface has not been enabled. Go to the General tab of the System Administration page and enable the feature.

3. On the Active Directory Mappings page, click **Add Rule**. This displays the User Role Mapping Rule panel in which you enter the rule parameters.

User Role Mapping Rule 

Rule Parameters


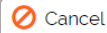
Object To Test: user

Relationship: is member of group ▾

Directory Object: - choose a Security Group - ▾

User Role To Apply: (None) ▾

Stop Evaluation:

 Save  Cancel

4. Choose the Relationship between the data of the user being tested and the Directory Object specified in the rule. The choice for this field changes the choices available in the other fields.

In this field, you can specify that objects must be in a OU or domain, a security group, in no domain, or that they exactly match the directory object you choose (the “is” choice on the Relationship menu).

5. Choose the Directory Object that the data from the tested computer must match.
 - a. Click in the *Directory Object* field to open the AD browser. The browser opens immediately below the Directory object field. The left panel is labeled “Search in,” and shows a tree of your AD domains.

- b. To expand the AD tree in the left panel, click on the plus button, next to the node you want to expand. To collapse the view on the left, click the minus button next to the node you want to collapse.
 - c. Click on the object in the left pane that defines the scope of your search. For example, if you have two domains, you might click on one of them, such as “DC=hq,DC=xycorp,DC=Local” in the example above.
 - d. If you see the object in the right panel that you want to use for this rule, double-click on it. The object, including full information about its location in the AD object tree, appears in the *Directory Object* field of the Rule Parameters panel and the browser will close.
 - e. If your actions did not automatically close the browser, click the ‘X’ button in the top right corner to close it.
- Note:** There are additional options for using the directory object browser. See [“AD Object Browser Options”](#) on page 128 for more information.
6. From the *User Role to Apply* dropdown menu, choose the role you want assigned to users whose AD details match this rule. Only existing roles appear on the dropdown.
 7. When you have entered all of the parameters for the rule, click **Save**. A newly created rule goes to the top of the table of AD rules.

Rolling the mouse cursor over the **i** button next to an object in the Match column provides a description of the object.

8. If necessary, use the up- and down-arrow buttons on the left side of each rule (or the drag-and-drop method) to change the order in which the rules are evaluated against a user.
9. Repeat this procedure beginning with step 3 for any other rules you need to create.

Disabling a Role

Any role can be disabled. If a role is disabled, all accounts with this role lose the permissions it provided. If an account has no enabled roles, it loses access to the console.

Important

Avoid disabling the Administrator and Administrator (Unified Management) roles if they are the only roles with permissions to view and manage user accounts and groups since this will make it impossible to restore access to these features without the use of special recovery commands.

To disable an *account*, see [“Disabling Login Accounts”](#) on page 101.

Deleting a Role

Custom user roles may be deleted *if* there are no accounts associated with them. Built-in user roles may not be deleted.

To delete a console user role:

1. From the console menu bar, click the configuration (gear) icon and choose **Login Accounts**. The Login Accounts page appears.
2. Click on the **User Roles** tab.
3. Click the Delete (trashcan) button next to the role you want to delete and confirm the deletion.

Chapter 4

Managing Computers

This chapter describes the steps necessary to install App Control agents on computers. It also describes how to upgrade or uninstall agents, view the details agents provide to the App Control Console and manage operating system updates on agent-managed systems.

Computer configuration tasks include choosing the method for assigning each computer to a security policy, adding installation packages for agents and rules files to the server, downloading the App Control Agent from a server to an endpoint, and installing the agent on client computers. Security policies are described in [Chapter 5, “Creating and Configuring Policies.”](#)

If you will be managing virtual machines, see [Chapter 6, “Managing Virtual Machines,”](#) in addition to this chapter.

Sections

Topic	Page
Computer Configuration Overview	118
Assigning Computers to a Policy	121
Uploading Agent Installers and Rules to the Server	131
Downloading Agent Installers	134
Installing App Control Agents	136
Upgrading App Control Agents	147
Uninstalling App Control Agents	154
Viewing the Table of Computers	156
Viewing Complete Details for One Computer	159
Moving Computers to Another Policy	170
Moving a Computer to Local Approval Mode	173
Adding Computers	173
Deleting Computers	173
Operating System Updates on Agents	175

Computer Configuration Overview

When you install and run the App Control Agent on a computer, the system become protected by rules defined on a App Control Server. After the agent is installed, an initialization process begins, and connected agents become visible to their App Control Server, delivering information about the endpoint and its files to the server.

Pre-Installation Activities

You make some key computer configuration decisions *before* installing agents on endpoints:

- **CLI Management** configuration options allow you to designate a user or group, or a password usable by anyone, to perform certain agent management activities in conjunction with Carbon Black Support. *Especially if you have systems that will be permanently offline*, it is best to choose one of these options *before* creating policies and distributing agent installation packages. See [“Advanced Configuration Options”](#) on page 744 for more details.
- **Rules file and agent installer packages** must be uploaded to the server from the User Exchange. Beginning with server v8.1.4, rules and agent installers have been separated from the server installation to allow for greater flexibility in updates.
 - For a new App Control Server, you must upload the rules file and agent package installers to the server before agents can be downloaded to endpoints.
 - For a server upgraded from a previous version, your previous rules and agent installers remain in place, but there might be new rule and/or agent updates.
- **Policies** determine the groups of security settings available to computers – every agent belongs to a policy. See [Chapter 5, “Creating and Configuring Policies,”](#) if you have not yet created policies.
- **Script Rules** are best created and enabled before you deploy agents. This ensures that all files matching those rules are in the inventory and can be approved or banned if you choose. Script rules created or enabled after an agent is deployed require that computers be rescanned before the files they identify are inventoried. See [Chapter 13, “Script Rules,”](#) for more details.
- **Review the expired certificate validation setting**, especially if you will be running offline systems. If you intend to allow file approval by certificates that have expired, make this choice before you download and install the agents on permanently offline systems. Otherwise, they will not be able to use expired certificates. See [“Approval with Expired Certificates”](#) on page 288 for more details.
- **Initial Policy assignment** to a computer can be determined by Active Directory data, as described in [“Assigning Policy by Active Directory Mapping”](#) on page 122; or by the agent installer used, as described in [“Downloading Agent Installers”](#) on page 134. Although you can change this decision later, determining how you want policies assigned before installing agents is recommended.
- **Preparing a reference computer for a “snapshot” of files** can give you a baseline for the files in your environment if you plan to closely monitor changes in your file inventory. Ideally, this is a clean computer onto which you install only the applications that you would like to run on some or all of your systems. Once the computer is prepared, you can install the agent and, after initialization is complete, use the Snapshot process as described in [Chapter 22, “Monitoring Change: Baseline Drift Reports.”](#)

Installation and Initialization

For each security policy you create, an *agent installer* is created for each supported platform (i.e., Windows, Mac, Linux) for which an initial installer package has been uploaded to the server. Each agent installer includes the policy assigned to the computer and the App Control Server address. If you do not use AD-based policy assignment, you choose the agent installer for each computer based on the computer's platform and the policy you want to control that computer.

Setting up your server so that it can create installers is described in [“Uploading Agent Installers and Rules to the Server”](#) on page 131. Installation of agents on endpoints is described in the sections [“Downloading Agent Installers”](#) on page 134 and [“Installing App Control Agents”](#) on page 136.

As soon as the agent software is installed, file initialization begins. The agent takes an inventory of all “interesting files” (executables and defined scripts) on the client computer's fixed drives (but not removable drives) and creates a hash of each file. When a computer first connects to the server, its agent sends these hashes to the App Control Server to update the server's file inventory.

Note

Virtual machines cloned from template computers can be configured to include or omit their initial (cloned) files in their inventory. See [“Configuring Clone Inventory”](#) on page 217 for more details.

App Control assigns files both a local and a global file state. Files on a computer at initialization receive a *local* state of Approved unless they have previously been identified and globally banned or banned by policy on the App Control Server.

Unless pre-banned or pre-approved by a App Control rule, files that the App Control Server has never seen before will get the *global* state of Unapproved and be added to the catalog. If a file was first seen on this agent *after* initialization, it will also get the *local* state of Unapproved on the agent. For more information on file state, see [“File State, Approving and Banning”](#) on page 47.

During initialization, the computer is protected by whatever security policy is assigned to it, and file activities are allowed or blocked according to that policy.

Post-Installation Activities

After you have installed the App Control Agent on endpoints and initialization is complete, there are many ways to monitor and manage your computers, including:

- **Viewing Computer Details** – App Control Server keeps details about each computer running a App Control Agent, including the computer's IP Address, whether it is currently connected to the server, the policy, mode and Enforcement Level it is assigned, computer model and system details, and its connection history. See [“Viewing the Table of Computers”](#) on page 156.
- **Viewing Computer-related Events** – You can monitor events related to a specific computer. See [“Event Reports”](#) on page 590.
- **Changing Policy** – You can change the security policy assigned to a computer if necessary. See [“Moving Computers to Another Policy”](#) on page 170 and [“Restoring Computers from the Default Policy”](#) on page 171.

- **Creating Clones** – If you plan to use a computer as the template for cloning other computers, see [Chapter 6, “Managing Virtual Machines.”](#)
- **Locally Approving Files** – You can temporarily put a computer into Local Approval mode so that files with a global state of Unapproved on the App Control Server can be installed locally and locally approved on this computer. See [“Moving a Computer to Local Approval Mode”](#) on page 173.
- **Viewing Details of Connected Devices** – You can track and manage fixed and removable storage devices on agent-managed Windows and Mac computers. See [“Viewing Devices on Computers”](#) on page 376 for more details.
- **Saving a Snapshot** – Once agent installation and initialization is complete, you can instruct the App Control Server to save a named snapshot of all files (by hash) on this computer currently inventoried by your server. This provides a reference point for analyzing changes in file inventory for that computer, other computers, or your whole network. See [“Creating and Modifying Snapshots”](#) on page 653 for more details.
- **Deleting Computers** – If a computer is going to be removed from your network or from App Control control, you can uninstall the agent and remove the computer from the table of computers on the server. This requires a specific series of actions detailed in [“Deleting Computers”](#) on page 173.

Permissions for Computer Management Features

Access to computer management features depends upon the Login Account Role Permissions for the user attempting access. The relevant permissions are:

- **View computers** – Ability to view computer pages
- **Temporary assign computers** – Ability to generate temporary policy override codes
- **Manage computers** – Ability to manually assign computer to policies and change Enforcement Level
- **Change advanced options** – Ability to change advanced computer options such as collection computer diagnostics and re-synchronizing
- **Manage system configuration** – Ability to upload new agent installer and rule packages.

The built-in user roles have the following computer management permissions:

- Administrator and PowerUser accounts (including Unified Management versions) with default permissions have full access to these features.
- ReadOnly users with default permissions can view the details of computers running App Control Agents but cannot add, delete, or change their configuration.
- The access level of users in custom login account roles depends on the role's permissions in the Computers asset rows on the Add/Edit Role page. Note that some features described here require additional permissions.

See [“User Role Permissions”](#) on page 106 for full details on viewing and changing login account role permissions.

In addition to standard computer management features, some or all users can be allowed to access agent management commands that can be used in special situations, usually in consultation with Carbon Black Support. See [“Configuring Agent Management Privileges”](#) on page 727 for more details.

Assigning Computers to a Policy

Every computer running a App Control Agent is assigned a security policy. There are three standard ways a computer can be assigned its policy:

- **By Agent installer** – Every policy you create generates a policy-specific App Control Agent installer for each supported platform, so when you install the agent on a computer, it is assigned a policy. When the agent contacts the App Control Server after agent installation, the computer is added to table of computers in the console. If you have not set up AD-based policy assignment, the agent remains in the policy embedded in its installer unless you manually reassign it.
You do not have to (nor should you) reinstall App Control Agent to make a policy change for a computer. You normally only need to install the agent once per computer.
- **Automatically, by Active Directory (AD) group mapping** – You can set up the App Control Server to run a script that assigns new and, if configured, existing computers to security policies according to the AD group information of the computer (or the user logged in on it). A computer's initial policy is defined by the agent installer. If that initial policy is configured to allow automatic policy assignment, this AD-based policy assignment takes precedence. Policy assignment by AD mapping is described later in this section.
- **Manually** – You can move any computer to a policy other than the one assigned by the installer or the AD-mapping facility. This might be useful if you discover that a particular computer used the wrong installer, or that its security policy should differ from other computers in the AD group used to map its policy. Manual assignment also might be used for a temporary situation that requires more or less restriction for a computer or its user. If you change a computer's policy manually, you can later restore it to its original policy (or to automatic assignment). Manual policy assignment is described in [“Moving Computers to Another Policy”](#) on page 170.

You can move computers from manual to automatic policy assignment and vice-versa.

Note

In certain cases, policy may be changed for reasons other than those listed above. For example:

- If you delete the policy an agent belongs to while the computer is offline, the agent moves to the Default policy group. See [“Restoring Computers from the Default Policy”](#) on page 171 for more detail.
- There is an Event Rule action that can move computers to a different policy when a specified event occurs. See [“Creating and Editing Event Rules”](#) on page 525 for more details.

If you are *not* using AD-based policy assignment, you can skip the next section and go directly to [“Downloading Agent Installers”](#) on page 134 for instructions on choosing a policy-specific installer.

Assigning Policy by Active Directory Mapping

You can create rules that map each computer to a certain policy based on its Active Directory (AD) data. AD-based policy assignment happens when an agent first contacts the App Control Server, and is checked again each time the server and agent re-establish contact or the logged-in user on the agent computer changes (see [“Computer Registration and AD Mapping”](#) on page 129 for more on when mapping can change).

AD Policy Mapping Summary

To make use of AD-based policy assignment, you must:

- **Install the App Control Server in an AD Domain** – Install the App Control Server on a computer that is a member of an Active Directory domain. By default, the App Control Server must be in the same AD forest as the computers and users you want to map. If you require cross-forest integration, contact your Carbon Black Support representative.
- **Enable the AD Mapping Interface** – You enable the AD-based policy mapping interface in the Active Directory / LDAP integration panel on the General tab of the System Configuration page.
- **Create AD-mappable Target Policies** – Create the security policies to which you want computers assigned by AD Mapping, and make sure these policies allow automatic policy assignment.
- **Create Mappings** – On the Mappings tab of the Policies page, create AD Policy Mapping rules that use AD data to assign computers to different security policies
- **Install or Move Agents to AD-mappable Policies** – For new agent installations, make sure the policy for the agent installation packages allows automatic policy assignment. For mapping to be successful, both the current policy of an agent and the policy to which will be mapped must have automatic policy assignment enabled. For existing agents, if necessary, you can change a policy from manual to automatic after installation or move the agent to an AD-mappable policy.

Platform Note: The App Control Server will do AD-mapping for any computer configured through your Active Directory server, including non-Windows platforms.

To enable the AD Mapping interface:

1. In the console menu, click the configuration (gear) icon and choose **System Configuration**. The System Configuration page appears.
2. If the General Settings view is not already displayed, click the **General** tab. The second panel on the General tab is Active Directory/LDAP integration.

The screenshot shows the 'System Configuration' page with the 'General' tab selected. The 'Active Directory / LDAP integration' section contains the following settings:

- AD-Based Logins: Disabled
- AD Security Domain: (none)
- AD-Based Policy: Disabled
- Windows 2000 DCs:
- Test AD Connectivity:

3. In the Active Directory/LDAP Integration panel, click the **Test** button next to Test AD connectivity. If you see a *Success* message, continue to the next step. If you see an *Error* message, your App Control Server is unable to access AD. AD Mapping will not work until you correct the problem.
4. If AD connectivity succeeds, click the **Edit** button at the bottom of the window.
5. In the *AD-based Policy* dropdown menu, choose **Enabled**.
6. To submit the changes, click the **Update** button and choose **Yes** on the confirmation dialog.

Creating AD Mapping Rules

After the AD-based Policy interface is enabled, a new tab, “Mappings,” is visible on the Policies page. Clicking on this tab opens the Active Directory Policy Mappings page. This is where you create rules to map computers with specified AD data to certain policies.

Before you begin setting up mapping rules, be sure you have created all of the policies to which you want computers mapped.

You can create mapping rules that test for matching AD data including organizational units, domains, security groups, computer names, and user names. Keep the following in mind when creating mapping rules:

- Although you can choose to match AD Security Group data for either users or computers, computer-based rules are recommended. With multiple users on a computer, sometimes simultaneously logged on, AD Mapping rules based on users could lead to unexpected results.
- App Control does not support policy mapping for AD object names that contain double quotes. Object names with double quotes cannot be handled properly by the directory object browser you use to create a mapping rule.
- Try to create as few rules as possible and test for groups rather than individual objects.

Table 12 shows the rule parameters you provide for a mapping rule.

Table 12: AD Mapping Rule Parameters

Parameter	Description
Computer Object to Test	The object that will be tested to see whether it matches the rule. The choices are Computer, User, and User or Computer.
Relationship	The relationship being evaluated between the Directory Object specified in the rule and the AD data from the computer being assigned a policy. The choices are: <ul style="list-style-type: none"> • is member of group • is in OU or domain • is • is not in any domain
Directory Object	The object in AD that the data from the tested object must match. Clicking the right end of this field opens a browser from which you can search for an object in your AD environment. The choices for the Directory object field change depending upon which Relationship you choose. If you choose "is not in any domain," no Directory object is necessary.
Policy to Apply	The policy to apply to a computer if its tested object matches the rule. The dropdown menu shows all available policies. Note: For policies created before implementation of Active Directory policy mapping, "Automatic policy assignment" is off by default. If you implement AD policy mapping and set up new mapping rules that apply to a pre-existing policy, you will need to change the setting on the policy itself for automatic mapping to take place. See " Creating Policies " on page 181 for more on automatic assignment choices.


The result of providing these parameters is a rule that can be read like a sentence. The following is how you might set up one rule.

Parameter	Example (value in bold)
Computer Object to Test	If a Computer ...
Relationship	... is in OU or domain ...
Directory Object	...matching OU = Marketing,DC=hq,DC=xyzcorp,DC=local ...
Policy to Apply	... assign that computer to the Standard Protection policy.

The procedure below shows how to configure a mapping rule. Although entry of most of the parameters are reasonably straightforward, pay particular attention to the Directory Object field, which requires use of a special AD browser.

To create an AD policy mapping rule:

1. In the console menu, choose **Rules > Policies**. The Policies page opens showing a list of all available policies.
2. Click the **Mappings** tab. The Active Directory Policy Mappings page appears with the Policy Mappings table, initially showing only the default rule.



Note: If no Mapping tab appears, the AD mapping interface has not been enabled. Go to the General tab of the System Administration page and enable the feature.

3. On the Active Directory Policy Mappings page, click **Add Rule**. This displays the Active Directory Policy Rule panel in which you enter the rule parameters.

4. Choose the Computer Object to Test (Computer, User, or Computer and User) from the dropdown menu. In most cases, **Computer** is the best choice.
5. Choose the Relationship between the data of the object tested and the Directory Object specified in the rule. The choice for this field changes the choices available in the other fields.

In this field, you can specify that objects must be in a OU or domain, a security group, in no domain, or that they exactly match the directory object you choose (the “is” choice on the Relationship menu). Generally it is best to choose a relationship that

maps multiple computers to a policy rather than one that singles out an individual computer or user.

Policy Mapping Rule

Rule Parameters

Computer Object To Test: Computer

Relationship: is member of group

Directory Object: is member of group

Policy To Apply: is in OU or domain

Save Cancel

6. Choose the Directory Object that the data from the tested computer must match.
 - a. Click in the *Directory Object* field to open the AD browser. The browser opens immediately below the Directory object field. The left panel is labeled “Search in,” and shows a tree of your AD domains

Policy Mapping Rule

Rule Parameters

Computer Object To Test: Computer

Relationship: is in OU or domain

Directory Object: - choose an OU or domain -

Search In: [All Domains]

- [All Domains]
 - DC=xyzcorp,DC=Local
 - DC=hq,DC=xyzcorp,DC=Local

For: all OUs and domains

select a DC or OU to search in

- b. To expand the AD tree in the left panel, click on the plus button, next to the node you want to expand. To collapse the view on the left, click the minus button next to the node you want to collapse.
 - c. Click on the object in the left pane that defines the scope of your search. For example, if you have two domains, you might click on one of them, such as “DC=hq,DC=xyzcorp,DC=Local” in the example above.

Policy Mapping Rule

Rule Parameters

Computer Object To Test:

Relationship:

Directory Object:

Search In: For:

[All Domains]

- DC=xyzcorp,DC=Local
 - OU =Administration
 - OU=Business Development
 - CN=Computers
 - OU=Domain Controllers
 - OU=Engineering

Administration
Business Development
Computers
Domain Controllers
Engineering
Marketing
Services
Resources
Sales

- d. If you see the object in the right panel that you want to use for this rule, double-click on it. The object, including full information about its location in the AD object tree, appears in the *Directory Object* field of the Rule Parameters panel and the browser will close.

Policy Mapping Rule

Rule Parameters

Computer Object To Test:

Relationship:

Directory Object:

Search In: For:

[All Domains]

- DC=xyzcorp,DC=Local
 - OU =Administration
 - OU=Business Development
 - CN=Computers
 - OU=Domain Controllers
 - OU=Engineering

Administration
Business Development
Computers
Domain Controllers
Engineering
Marketing
Services
Resources
Sales

- e. If your actions did not automatically close the browser, click the 'X' button in the top right corner to close it.

Note: There are additional options for using the directory object browser. See [“AD Object Browser Options”](#) on page 128 for more information.

7. From the *Policy to Apply* dropdown menu, choose the policy you want assigned to computers that meet the requirements of this rule. Only existing policies appear on the dropdown – if the policy for this rule has not been created yet, cancel the creation of this rule and go to the Policies page to create the new policy.

- When you have entered all of the parameters for the rule, click **Save**. A newly created rule goes to the bottom of the table of AD rules, just above the default rule, and all rules above it take precedence. In the example, the rule instructs the App Control Server to assign any computer belonging to the Engineering OU in the domain *hq.xyzcorp.local* to the Research Group policy.

Object	Relationship	Match	Action	Policy
<i>if Computer</i>	is in OU or domain	Engineering	apply policy from	Research Group
<i>[all others]</i>			apply policy from	Default Policy

Rolling the mouse cursor over the **i** button next to an object in the Match column provides a description of the object.

- Once you have addition rules, if necessary, use the up- and down-arrow buttons on the left side of each rule (or the drag-and-drop method) to change the order in which the rules are evaluated against a computer. Remember that the [all others] rule always is the last one in the table.
- Repeat this procedure beginning with step 3 for any other rules you need to create.

Mapping Rule Ranking

AD Mapping rules are scanned in top-to-bottom order on the Mappings page, and only the first match on the list is applied. You can rearrange the order of rules if you find that you would prefer a different policy assignment outcome than you are seeing.

There is a default AD Mapping rule that cannot be deleted, nor can it be moved from the bottom of the Policy Mappings rule table. It maps “[all others]”, that is, all computers that have not matched any of the other rules in the table, to the policy you choose. Because it remains at the bottom of the table, it assures that any automatically mapped computer is assigned to some policy. It is initially mapped to the Default Policy, but you can change this. Creation of an “AD Default Policy” is recommended so that computers not matching other rules have a policy that best reflects a default security level with settings you want.

AD Object Browser Options

This section describes the AD Object browser, which you use to select a Directory object when defining an AD Mapping rule, in more detail.


The left panel of the AD Object browser is where you determine the scope of your search. It displays an AD tree with “[All Domains]” at the top of the tree and then shows the

contents of the tree in standard browser format, with +/- buttons at each node that contains other objects so that you can collapse or expand the tree at that point.

The right panel has a description of what you are searching for, based on the “Relationship” value you entered in the Active Directory Policy Rule parameters. When you click on a node in the tree on the left, all objects immediately under that node matching the “Relationship” (e.g., “OUs and domains”) appear in the right panel. You click on an object in the right panel to select it and enter it in the Rule Parameters panel.

Object Search Depth

In the upper right area of the browser, there is a checkbox labeled “Deep”. When you check the Deep box and click **Go**, this results in a multi-level search that examines not just the immediate contents of the selected node but the contents of any nodes inside it, regardless of how many layers deep they are. Notice the greater number of results in the right panel of case B in the illustration below.




A. Results of a standard search in an AD domain



B. Results of a Deep search in the same domain

Object String Match

Another option in the AD Object browser is searching by string match. If you enter a string of characters in the box immediately to the left of the “Deep” checkbox, you can search for AD objects in the selected node that start with, end with, or contain the string. You make the choice of how to use the string via the dropdown menu to the left of the text box. For example, if you entered “eng” in the text box and then searched for “group names that contain” the string, you would match both “Engineering” and “System Engineering” groups if they existed in the node selected on the left.



Computer Registration and AD Mapping

Certain events trigger registration of the agent on a computer with its App Control Server. When this occurs, the following conditions may affect AD policy mapping:

- When the App Control Agent is first installed, the computer will register with the server for the first time, with the users that are logged on at the time. If no users have logged

- on since the last time this computer was started, the App Control Server shows an empty user list for that agent computer.
- When an agent computer is restarted, if the App Control Agent reconnects to the server before any user logs in, the user list for that registration will be empty.
 - All agent computers (whether or not they use automatic policy assignment) re-register whenever their list of user sessions changes.
- Platform Note:** Because of the way Windows handles sessions, a user's session on a Windows computer does not necessarily end upon logout. It persists until it is replaced by a different user's session.)
- Agent computers are disconnected by the server whenever the server restarts and re-registered when they reconnect to the server.
 - The server disconnects a computer (forcing re-registration) whenever the agent computer's policy assignment is changed manually, or if it is changed from manual to automatic.

Clearing the Server AD Cache

The AD information that is used to map agent computers to policies is cached on the App Control Server and updated every four hours. It is also updated whenever a App Control rule change occurs that is related to AD mapping.

If you make a change to this AD information on your AD server – for example, changing the group a computer or user is in, or adding a computer – this information normally does not become available to the App Control Server until the next scheduled cache upgrade. If you know you have made relevant changes or you see incorrect policy mapping, you can clear the server cache so that the App Control Server immediately begins updating AD information.

To clear the server cache and update AD information:

- On the Mappings tab of the Policies page, click **Clear Server Cache** in the Actions menu.

Viewing AD Computer Details in the Console

If you have integrated AD and App Control Server, anytime a computer name in an AD domain appears in a table in the App Control Console, additional information can be displayed by clicking on that computer name. For example, if you display the Events page, some events include the computer associated with the event.

If the name is identified as an AD computer name, it is highlighted in blue, and when you click on it, the Computer Details page appears. If you click the **AD Details** tab on this page the AD information that is available for that computer is displayed.

Similar information is displayed about a user when you click on a highlighted AD username in a console table.

Uploading Agent Installers and Rules to the Server

Beginning with App Control 8.1.4, agent installers and the rule file that determines their behavior are no longer included as part of a App Control Server installation. You upload rule and agent installer packages separately after you install the server. This allows Carbon Black more flexibility to make improved agents and new rules available to you independent of server releases.

As part of this enhancement, App Control includes a new drag-and-drop interface that you use to add new rule files and agent installers to your server as they become available. This eliminates complex and error-prone manual installation procedures.

A user must have *Manage system configuration* permission to upload and install agent installers and rule files. These files are available on the Carbon Black User Exchange. If you have enabled the Carbon Black File Reputation (CDC) connection from your server and the health indicators option within the CDC, a health indicator will inform you when agent installers or rule files newer than the ones you current have are available.

When new rule files and agent installers are uploaded and installed on the server, the server is restarted and agent installation package generation is enabled. However, automatic agent upgrades are also disabled when you upload a new agent installer (but not a rules file) so they must be re-enabled if you are planning to use automatic upgrades.

Notes

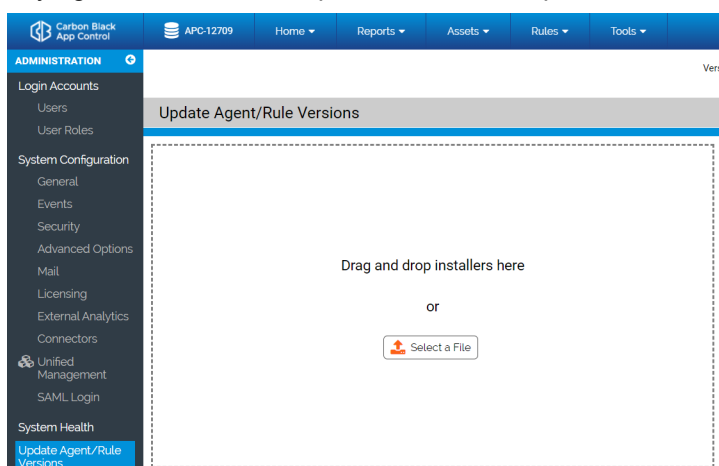
- The installers for rule files and agent packages found on the User Exchange are strictly for upload to the server. They cannot be used directly to install or update agents on endpoints.
- Rule file and agent package installers for upload to the server cannot be installed on pre-8.1.4 servers.
- Rules and agent installation packages must be uploaded to the server one file at a time. If you drag multiple files into the upload interface simultaneously, the uploads will fail.
- When you use the Update Agent/Rule Version page to upload agent package installers, generation of agent installer for endpoints is enabled on the server. However, if you use other methods to update or add agent installers, the installers will not automatically be generated on the server.

To upload installers for rule files and agent packages to a server:

1. Log in to the Carbon Black User Exchange and locate the new rules file and/or agent installer packages. Links to these packages are found on the [Documentation & Downloads](#) area for App Control on the User Exchange.
2. Download the rules file installer and the agent installer packages for each OS platform you have in your environment to a filesystem on or accessible to your App Control server. These files are named as shown below:
 - **Rules file installer** – RulesInstaller.exe
 - **Windows agent installer** – WindowsHostPackageInstaller.exe
 - **Mac agent installer** – MacHostPackageInstaller.exe
 - **Linux agent installer** – LinuxHostPackageInstaller.exe

3. Log in to your App Control Server using an account that has *Manage system configuration* permission.
4. In the console menu, click on the configuration (gear) icon and choose **Update Agent/Rule Versions**.
5. To install a new rules file on the server, drag the RulesInstaller.exe file from your download folder into the target zone on the Update Agent/Rule Versions page, or click **Select a file** to find the file via a browser.

Important: If you are updating the rules file, do not attempt to simultaneously upload any agent files. Each file upload must be complete before the next one is started.



When the upload begins, the server checks to see whether the package is correctly signed. If so, it is installed on the server. Messages report on each stage of the progress (or failure) of the upload and installation.

6. When the rules upload is complete, repeat the file drag-and-drop or selection process for each agent installer package you want to upload to your server.

Important: Do not simultaneously upload multiple agent files. Each file upload must be complete before the next one is started. The server restarts after each upload. A success message appears when the new agent installer package is available.

Important: Remain on the Update Agent/Rule Version page while uploads are proceeding. You can navigate to other pages, but since the server is restarted after the upload, activity on another page can be interrupted at an unpredictable point.

7. Once you have finished uploading rules files and agent installers to a server:
 - If you are setting up a new server, set up the policies you want to control your agents (see [Chapter 5, “Creating and Configuring Policies”](#)).
 - Choose a policy assignment method (“[Assigning Computers to a Policy](#)” on page 121).
 - Install agents on endpoints (see “[Installing App Control Agents](#)” on page 136).
 - If you are uploading new agent installers or rules files on an existing server, begin upgrading agents according to the upgrade plan appropriate to your site. See “[Upgrading App Control Agents](#)” on page 147.

Unified Management Note: If you are using Unified Management to manage multiple App Control Servers, you must upload new rule and agent packages to each server separately. The management server does not broadcast these packages to the managed servers.

Viewing Agent / Rule Versions and Package Generation Status

If you have System Health indicators enabled, you will be notified when your agent or rule installer versions are out of date. With or without System Health indicators enabled, you can view the current versions of agent and rule installers in the App Control console. This allows you to compare them to the latest versions listed on the User Exchange.

To view current agent and rule installer versions:

1. On the console menu, choose **Rules > Policies**. The Policies page appears.
2. On the Policies page, click the following link: **Click here to view available Carbon Black App Control Agent/Rules versions**. The Installer Versions page appears.

[Go back to the policies page](#)

Installer Versions		
Installer	Version Installed	Package Generation Status
Windows Agent	8.1.8.306	Enabled
Mac Agent	None	Disabled
Linux Agent	None	Disabled
Rules	1.4.35	N/A

3. Review the version numbers and any status messages for the rules and agent installers.
4. When you have finished reviewing the versions, either use the **Go back to the policies page** link or use the console menu to navigate elsewhere.

The Installer Version page shows two key pieces of information for the Windows Agent installer, Mac Agent installer, Linux Agent installer, and Rules installer currently on the server:

- **Version Installed** – This is either a version number for the installer in each category or “None” if there is no installer for that item.
- **Package Generation Status** – This indicates whether installation packages – that is, the installers that are used on endpoints – are being generated. Even if there is an agent installer for a platform (e.g. Linux) available on the server, generation of the installer that will be used on the endpoint might be disabled for that platform. The possible status messages are:
 - **Enabled** – This indicates that a Rules file is available and agent package generation is enabled for the agent on this platform.
 - **Disabled** – This indicates that agent package generation is disabled. If a version number appears in the Version Installed column, generation is disabled by a setting on a hidden page in the console. If Version Installed shows “None”, generation is disabled because the agent package for that platform was never installed.
 - **Disabled due to missing default rules** – This indicates that an installer version has been uploaded for the agent on this platform but installers for endpoints cannot be generated because no Rules file has been uploaded to the server from the User Exchange.
 - **N/A** – This appears for the Rules file because package generation is unnecessary for rules.

Downloading Agent Installers

When you create a new policy, the App Control Server generates a policy-specific agent installer for each agent platform and posts it to an agent download area. Each installer specifies the policy, policy settings, Enforcement Level, and the address of the server managing the agent.

When the App Control Server is upgraded, agent installers are also upgraded to the new version. Depending upon your upgrade plans, you might download and install the new agent version on the endpoint or allow the App Control Server to manage the upgrade. See “[Upgrading App Control Agents](#)” on page 147 for more details.

Note

If you are using Active Directory to assign policies to all computers, use any installer whose policy has the *Automatic Policy Assignment for New Computers* box checked. Once the agent is installed on a computer and makes contact with the App Control Server, the correct AD-based policy for the computer will be assigned automatically. If the computer is unable to contact the App Control Server to retrieve AD mapping rules, the policy specified in the agent installer you used remains in effect.

App Control Agent installers are created in a file format appropriate for each platform:

- MSI (Microsoft installer) packages for Windows
- DMG files for Mac OS X
- TGZ archives for Linux

The download page for these packages is accessible via a URL on the server. You can bookmark this URL and access the page without logging into the console.

To download an agent installer:

1. In the console menu, choose **Rules > Policies**. The Policies page appears:

Version 8.1.10.109

Users can download Carbon Black App Control Agent software from <https://APC-12709/hostpkg>

Click here to view available Carbon Black App Control Agent/Rules versions.

Policies

Policies Mappings

Group By: (none) Ascending

Show Filters | Show Columns | Export to CSV | Refresh Page

Action Add Policy

Policy	Connected Enforcement	Disconnected Enforcement	Total
APC-99986-ucdjfp	High (Block Unapproved)	High (Block Unapproved)	1
Default Policy	None (Visibility)	None (Visibility)	0
Help Desk	None (Visibility)	None (Visibility)	0
Local Approval Policy	Local Approval	Local Approval	0
Research Group	None (Visibility)	None (Visibility)	0
Template Policy	None (Visibility)	None (Visibility)	0

6 items Page 1/1

2. On the Policies page, click the download link at the top of the page. The publicly accessible URL for this page takes the following format:

`https://server_name/hostpkg`

The Download Install Packages page appears:

Download Carbon Black App Control Agent Install Packages

Carbon Black App Control protects your computer and the network from viruses, spyware, and other malicious applications.

Installing the Carbon Black App Control Agent software is simple:

1. Click the installation setup file for the policy assigned to you by your network administrator.
2. Download the installation setup file to a convenient location on your hard-drive.
3. From the download directory, double-click the newly downloaded file to install Carbon Black App Control Agent.

Carbon Black App Control Agent Installation Setup Files ?				
Refresh Page				
Policy Name	Install Package	Description	Date Created ▲	Date Modified
APC-99986-ucdjfp	Windows	simple install agent installation policy	Jun 23 2020 05:16:01 PM	Jun 25 2020 04:56:10 PM
Research Group	Windows		Jun 25 2020 05:14:21 PM	Jun 25 2020 05:14:22 PM
Help Desk	Windows		Jun 25 2020 05:14:39 PM	Jun 25 2020 05:14:39 PM
3 items			Page 1/1	

3. In the Agent Installation Setup Files table, locate the installer file by policy name.
4. To download the installer, click the platform name (e.g. Windows) for the computer on which you want to install the agent, and save the file.
5. When the download is complete and you are ready to install the agent, follow the instructions in the next section, "[Installing App Control Agents](#)".

Installing App Control Agents

The App Control Agent installation process is non-interactive; it requires no user input. As soon as installation is completed, the App Control Agent begins working – no additional configuration is needed, and in most cases a restart is unnecessary.

Preparing for New Agent Installation

Before installing a new App Control Agent on any platform, review the following considerations:

- The agent is a per-system application, not per-user.
- Installing App Control agents on containers is not supported.
- Make sure the computer and operating system on which you are installing the agent is supported. See the separate *Operating Environment Requirements* document for the agent hardware requirements and *Supported Carbon Black sensors and agents* on the User Exchange for the OS versions supported by the current agent.
- As soon as the agent is installed, the computer is protected by a security policy, and the agent connects to the server and begins initializing files. Because initialization can involve significant data flow between the server and its new clients, consider your network capacity and number of files when planning agent roll-out. Simultaneous agent installation on all computers on a large network is not recommended.
- If you are configuring your App Control Server for the first time, consider setting up a reference computer with files you know you want to globally approve; you can also use that computer as a baseline for measuring any file inventory drift. See [Chapter 22, “Monitoring Change: Baseline Drift Reports”](#).
- Decide how the agent will be installed on this system. You can choose among the following options:
 - Use an existing software deployment mechanism. Although new agent installations are normally done in non-interactive mode, you can optionally create an interactive end-user installation experience. If you use a third-party distribution system to install agents, follow all recommended procedures. For Windows installations, disable any possible MSI or MSP transformations inside your distribution system (such as SCCM).
 - Have a system administrator or other qualified employee install the agent software manually on each user’s computer.
 - Permit users to install the agent software themselves. Send e-mail to users associated with each policy and inform them to browse to the agent download URL or another shared location, download the specific installer file for their policy, and run the installation on their computers. No interaction is needed – the installation runs without prompts and then the agent begins to initialize files.
- The agent installer must be run by a user with the appropriate administrative rights. On Windows, this can be either by Local System or by a user account that has administrative rights and a loadable user profile. On Mac and Linux, the user must be able to use *sudo*.
- Make sure your server has the latest agents and rules; see [“Uploading Agent Installers and Rules to the Server”](#) on page 131.
- Be sure to download the correct installation package for your policy and platform; see [“Downloading Agent Installers”](#) on page 134. If you are using AD-based policy

- assignment, a platform-specific agent installer for any policy that allows automatic policy assignment may be used.
- Although the console prevents creation of policies whose names have generally known invalid characters, examine the policy name to see whether it contains characters that might require special handling (such as escaping in a command line) on your specific platform.
 - If Microsoft OneDrive™ is in use, only the default path is supported. (c:\users\\OneDrive)
Custom OneDrive paths are not supported.
During Initialization, the App Control agent will ignore the One Drive directory, thus leaving all of the files inside it as unknown.

NOTES:

- VMware Carbon Black does not recommend storing executables in the cloud. In the event that a file is executed from the cloud, the agent treats the file as unknown.
- Support for OneDrive is enabled by default. To disable OneDrive support, contact CB Support.

Considerations When Installing a App Control Agent using GPO

You can deploy a Carbon Black App Control Agent using a GPO, Group Policy Object.

Using a Group Policy to deploy software is a common task that is well documented. For example: [How to use Group Policy to remotely install software in Windows Server](#)

Some network latency issues have occasionally been reported when deploying agents using a GPO. If you intend to deploy your App Control agents using a GPO, please consider the following:

- Make sure that the MSI is copied to a local drive and then called with a fully-qualified path.
- Don't let the deployment get in the business of "monitoring" for the presence of the expected version.

This will break upgrades done via console because SCCM/etc will want to roll back any upgrades.

- Disable any possible MSI or MSP transformations inside your distribution system.
- To reduce network traffic, copy the installer first and then install the agent locally.

Finally, if you continue to experience network latency issues after implementing all the previously listed suggestions, consider wrapping the agent installer in another script. Most commonly, this is done via a batch file which copies the installer locally and then executes it. For example:

```
@ECHO OFF
COPY "\\network_location\folder\bit9_agent_install.msi"
"%WINDIR%\Temp\bit9_agent_install.msi"
msiexec /i "%WINDIR%\Temp\bit9_agent_install.msi" /qn /
l*v+"%WINDIR%\Temp\bit9_agent_install_log.txt"
```

If you continue to experience problems deploying the installer using a GPO or need further assistance, contact Customer Support.

Installing the Agent on a Windows Computer

As an MSI package, an agent installer for Windows can be customized as you choose, including modification of the installation directory. Please refer to the Microsoft MSI documentation for information about configuration options. The installer for Windows is named in the following way, varying by policy:

- *polycyname.msi*

Notes

- The use of Windows Installer Transform files (.mst) is *not supported* with the agent installer on Windows clients.
- Windows Installer Patch files (.msp) are no longer used for build-to-build agent upgrades. Be sure to update any scripts referring to these files.
- The agent for this release cannot be installed on systems running Windows 2000, Windows 2003 Server versions prior to SP1, or Windows XP versions prior to SP2.

The following procedure assumes you have already:

- uploaded agent and rule packages as described in [“Uploading Agent Installers and Rules to the Server”](#) on page 131
- created one or more security policies for your agents as described in [Chapter 5, “Creating and Configuring Policies”](#)
- downloaded the appropriate installer as described in [“Downloading Agent Installers”](#) on page 134
 - you can download the installer either to the computer on which you want install the agent or to a location accessible to a third-party software distribution system
 - for AD-based policy assignment, use an installer for any policy with automatic policy assignment enabled.

To install a new agent on a Windows computer:

1. On the client computer, run the Windows agent installer you have selected. You can use any standard means for installing from MSI files, with the following considerations:
 - a. The default agent application directory is **C:\Program Files\Bit9\Parity Agent** for 32-bit systems and **C:\Program Files (X86)\Bit9\Parity Agent** for 64-bit systems. To change the installation directory, perform the installation from the command line using the appropriate MSI command-line options.
 - b. If you plan to accept the default application directory, you can use any MSI installation method, including simply double-clicking on the MSI filename.
 - c. If you are installing the agent manually or via a third-party distribution system and want to specify a non-default *data* directory, **do not** choose a data directory that is underneath the main program installation directory. Putting the data directory under the installation directory will cause the agent to malfunction and disconnect.
2. During Windows agent installation, the installer displays a message box that closes automatically when installation is complete. This box includes a Cancel button so you can end the installation before it completes, if necessary.

3. If you run anti-virus software, exclude the App Control Agent installation directory from anti-virus scanning. For enhanced security, App Control self-protects its application directory. To avoid performance issues, configure your AV software so that the following files and directories are not scanned or blocked:
 - **Parity.exe** – the agent process
 - **Program Files\Bit9** – the default agent program directory on 32-bit systems; if you did not take the default, use the directory you chose
 - **Program Files (x86)\Bit9** – the default agent program directory on 64-bit systems; if you did not take the default, use the directory you chose
 - **ProgramData\Bit9\Parity Agent** – the default agent data directory on Vista, Windows 7, 8 and 10, and Windows Server 2008 through 2016 systems; if you did not take the default, use the directory you chose
 - **\Documents and Settings\All Users\Application Data\Bit9\Parity Agent** – default agent data directory for supported OSs not listed in the previous item
4. Personal firewalls such as Zone Alarm may recognize the agent as a new application and block access to the network. Instruct users running the agent to permanently allow it access on their computers.

See [Chapter 20, “Endpoint Notifiers and Approval Requests,”](#) for a description of what the user sees on a system protected by the agent.

Conditions Requiring Reboot after Installation

Normally, the VMware Carbon Black App Control Agent does not require a reboot after it is installed. However, a reboot is required under the following conditions:

- If you are using DFS and have installed an agent on a Windows 2003 or XP system, you must reboot the agent system to get full enforcement of App Control file rules. Because of an operating system limitation, DFS operations (including file executions) cannot be detected by the agent until the system has been rebooted. In this case, the Upgrade Status column on the Computers page shows Reboot Required for the affected computer.
- On any version of Windows, if a file is in use by another application when the App Control installer tries to write that file, the system schedules the file to be replaced on next reboot, and the console shows Reboot Required for the affected computer.

Command Line Installations of Windows Agents

You can install or upgrade agents using MSIEXEC commands, either manually or with third-party software distribution tools. In some cases, you might do this using an installer created in the App Control Console for a specific Policy, in which case the Policy and Server information is built in to the installer. In other cases, you might use the “unbranded” agent installer, **ParityHostAgent.msi**, which does not include this information. With the unbranded installer, you can provide custom parameters.

If you are going to create custom MSIEXEC commands for agent installation, you should be aware of the standard MSIEXEC parameters you might want to use, such `/quiet` or `/qn` for automatic installations without prompts. See <https://technet.microsoft.com/en-us/library/bb490936.aspx> for a list of those parameters.

In addition to standard MSIEXEC syntax, you will need to use parameters specific to App Control agent installation. [Table 13](#) shows these parameters. They can be used to modify an agent installation in various ways. For example, the following example installs an agent and overrides the server defaults to connect to a specific server:

```
msiexec /i ParityHostAgent.msi B9_SERVER_PORT=41002
B9_SERVER_ID={b9}Fkmv+XIVXwjg7654AB2oxgxh/qxs8tsPGbX1Dabi19xs
B9_SERVER_IP= newserver.mycorp.local
```

Table 13: Specific Parameters for Agent Installation and Upgrade

Parameter	Description & Example
B9_CONFIG	<p>This allows you to specify the location of the configlist.xml file when installing or upgrading an agent using the unbranded package. The configlist contains all of the App Control rules, such as file approvals and bans, that are created on the server and applied to the agent. The agent can download the configlist from the server once it is connected, but since there will be a delay before it can complete the download, it is typically best to import all the rules immediately during agent installation.</p> <p>This option requires an additional URL argument (or a local path) added to the MSIEXEC command to indicate of the location of the configlist.xml file that the installer should use.</p> <p>Example: B9_CONFIG=https://<serveraddress>/hostpkg/pkg.php?pkg=configlist.xml</p> <p>Important: This setting is not for use with a “branded” agent installation package (i.e., one that is specific to a policy).</p>
B9_NOCONFIG	<p>This specifies that you don't want to download all of the App Control rule information at the same time as agent installation or upgrade. In that case, you must rely on the agent connecting to the server later and downloading any rule changes.</p> <p>Important: This option is reasonably safe to use for upgrades, which should already have nearly current rules, but is not recommended for new installations since it can result in agents not enforcing rules properly until they can download them all from the server, an unpredictable period of time. Also, it is unnecessary for branded (policy-specific) installation packages.</p>
B9_SERVER_PORT	<p>For unbranded package installations, this allows you to set the port for communication from the agent to the server if the unbranded installer is used.</p> <p>Example: B9_SERVER_PORT=41002</p> <p>Important: This setting is for use with unbranded installation packages or if you need to change this parameter as part of repair or upgrade. It is not for use with installations of a “branded” agent (i.e., one that is specific to a policy).</p>

Parameter	Description & Example
B9_SERVER_ID	<p>This sets the value for the App Control Server ID. Set during installation to manually establish the ID setting if the msi package used is unbranded. This value is the serverIDString property on the shepherd_config.php page in the console.</p> <p>Example: B9_SERVER_ID={b9}cu+ox2O9/EvVtKe+eMlkwVqpiy+kJsgs+opq8jjFWZw=</p> <p>Important: This setting is for use with unbranded installation packages or if you need to change this parameter as part of repair or upgrade. It is not for use with installations of a “branded” agent (i.e., one that is specific to a policy).</p>
B9_SERVER_IP	<p>This sets the address of the App ControlApp Control Server. You can use it to manually establish the location setting for unbranded agent installer packages.</p> <p>Example: B9_SERVER_IP=server2.mycorp.local</p> <p>Important: This setting is for use with unbranded installation packages or if you need to change this parameter as part of repair or upgrade. It is not for use with installations of a “branded” agent (i.e., one that is specific to a policy).</p>
B9_HOSTGROUP	<p>This sets the desired Policy for the agent. You can use it to manually establish the Policy setting for unbranded agent installer packages that are unbranded.</p> <p>Example: B9_HOSTGROUP=Monitor</p> <p>Important: This setting is for use with the new installations of the unbranded installation package. It does not change policy during an upgrade, and if AD policy assignment is enabled for this agent, the policy will be changed according to your AD mapping rules.</p>

Installing the Agent on a Mac Computer

Important

App Control supports installation of agents only on systems listed for this release in the *Supported Carbon Black sensors and agents* on the [User Exchange](#). Also, see also the *Release Notes* for your version of App Control for any special considerations.

For Mac computers, you install the App Control Agent by using the appropriate installer DMG file. Installers for Mac are named as follows, varying by policy:

- *policyname-mac.dmg*

The following procedure assumes you have already:

- uploaded agent and rule packages as described in [“Uploading Agent Installers and Rules to the Server”](#) on page 131
- created one or more security policies for your agents as described in [Chapter 5, “Creating and Configuring Policies”](#)
- downloaded the appropriate installer as described in [“Downloading Agent Installers”](#) on page 134
 - for AD-based policy assignment, use an installer for any policy with automatic policy assignment enabled.
 - the same downloaded agent installer can be used on multiple endpoints, and can also be distributed to endpoints via SSH or distribution mechanisms like Casper
- for systems running 10.13 High Sierra or later, allowed the agent kernel extension as described in [“Allowing the Agent Kernel Extension \(High Sierra or later\)”](#) on page 143

To install a new agent on a Mac computer:

1. Open a Terminal window and change directory to the location where the installer was downloaded (by default, the user-specific Download directory).

```
cd ~/Downloads
```

2. To begin the installation, double-click on the agent installation file you downloaded, *policyname-mac.dmg*. A standard package installation dialog begins.
3. Respond to the installation dialog prompts, and when the dialog indicates the installation was successful, click **Close**. The agent begins operating immediately.
4. If you run anti-virus software, exclude the App Control Agent installation directory from anti-virus scanning. For enhanced security, App Control self-protects its application directory. To avoid performance problems, use whatever mechanism is provided by your anti-virus software vendor to specify that the following directories are not scanned:
 - **/Applications/Bit9/Daemon/b9daemon** – the App Control Agent process
 - **/Applications/Bit9** – the App Control program directory
 - **/Library/Application Support/com.bit9.agent** – the App Control data directory
 - **/Library/Extensions/b9kernel.kext** – the App Control driver location for OS X versions 10.9 (Mavericks) and later

-or-

/System/Library/Extensions/b9kernel.kext – the App Control driver location for OS X versions prior to 10.9

5. The Mac firewall may detect the agent as a new application and block access to the network. Instruct users to permanently allow incoming connections to **b9daemon**.
6. Enable the Mac System Updates updater, which allows minor updates to the OS to be approved for installation. Be sure you are running at least version 9 of this updater. You enable updaters on the **Software Rules > Updaters** page.

See [Chapter 20, “Endpoint Notifiers and Approval Requests,”](#) for a description of what the user sees on a system protected by the agent.

Allowing the Agent Kernel Extension (High Sierra or later)

Note

If you are using a version of Mac prior to 10.13 the steps described in this section are not necessary.

The macOS version 10.13 (High Sierra) introduced changes in the way system extensions are handled. If you are installing or upgrading the App Control Agent on any version of 10.13 (or later) additional steps are needed to approve the ‘Carbon Black, Inc.’ system extension, which is required for proper operation of the agent. This is true for manual agent installations and upgrades as well as those initiated from the App Control Console.

To allow the kernel extension during agent installation or upgrade:

1. Run the App Control Agent installer. For non-MDM installations on High Sierra (or later), while you are running the App Control Agent installer, macOS will report that a system extension signed by ‘Carbon Black, Inc.’ was blocked. This will happen even if the extension is already approved on the system.
2. When this message appears, go to **System Preferences > Security & Privacy** on the Mac system and click the **Allow** button for “Carbon Black, Inc.”. (This was ‘Bit9, Inc.’ for agents prior to 7.2.3 Patch 12.)

Important

It is possible that you delay or are unable to allow the kernel extension immediately after agent installation or upgrade. For example, you might automatically upgrade unattended endpoints.

If you do not allow the kernel extension, agent installation continues, and the upgraded agent will connect to the server, but it will not enforce rules until you allow the extension to load. On the App Control console, this agent will show a status of Unprotected, Reboot Required. If this is the case, complete the steps in the following section.

To allow the Carbon Black, Inc. kernel extension after agent installation or upgrade:

1. Navigate to **System Preferences > Security & Privacy** on the endpoint and click the **Allow** button for ‘Carbon Black, Inc.’.
2. Do one of the following to restart the agent:
 - Reboot the agent system.

- or -

- Manually stop and start the agent using the following steps in a terminal:

```
cd /opt/bit9/bin
./b9cli -password <password>
./b9cli -tamperprotect 0
./b9cli -shutdown
sudo ./b9cli -startup
```

Note that you **must** run the startup step as root or using sudo.

Installing the Agent on a Linux Computer

Important

App Control supports installation of agents only on those Linux versions and kernels listed in *Supported Carbon Black sensors and agents* on the [User Exchange](#). Also, see also the *Release Notes* for your version of both App Control Server and the App Control Linux Agent for any changes and special considerations.

For Linux computers, you install the App Control Agent by running a script after extracting the appropriate TGZ archive. App Control Server supports agents on Linux computers running Red Hat and CentOS versions, both of which use the same installation file. The installation files are tarballs named by policy and operating system, such as:

- *polycyname-redhat.tgz*

Carbon Black recommends disabling Prelinking on RedHat and CentOS computers before installing agents. Prelinking has negative impacts on performance and App Control features (see the Release Notes). However, if you must enable Prelinking on your RedHat and CentOS systems, enable the RedHat Prelinking updater before installing agents. See [“Approving by Updater”](#) on page 266 for instructions on enabling updaters.

Note

Although not required for the initial agent installation, **gawk** and **unzip** are required for Linux agent upgrades initiated by the App Control Server. If necessary, update the Linux distribution to include them before installing the agent.

The agent is normally installed with a GUI-based blocked file notifier. This notifier appears when a user attempts to take an action that is either totally blocked by the agent or that requires a user decision about allowing it to proceed. For Linux systems that are not running a graphic interface package or prefer to eliminate user interaction for some other reason, the agent for Linux can be installed without the notifier. This **-n** option may be

added as a flag on the installation script command for the agent, and is shown in the procedure below.

Note

On a system that you choose to run without the notifier, install an agent with a Low or High Enforcement policy. Agents in Medium Enforcement policies prompt users to allow or block many actions, and this prompt will not be available without a notifier.

The following procedure assumes you have already:

- uploaded agent and rule packages as described in [“Uploading Agent Installers and Rules to the Server”](#) on page 131
- created one or more security policies for your agents as described in [Chapter 5, “Creating and Configuring Policies”](#)
- downloaded the appropriate installer as described in [“Downloading Agent Installers”](#) on page 134
 - for AD-based policy assignment, use an installer for any policy with automatic policy assignment enabled
 - the same downloaded agent installer can be used on multiple endpoints, and can also be distributed to endpoints via SSH or other distribution mechanisms

To install a new agent on a Linux computer:

1. Make sure the account for the user installing the agent has administrative rights, or that the user can use **sudo**.
2. Extract and uncompress the agent tarball archive for the policy you have chosen for this computer. If the policy name contains characters not accepted in command arguments, such as spaces or parentheses, escape these characters with a backslash:

```
tar -xvzf <policyname>-redhat.tgz
```

3. Change to directory matching the download tarball name.

```
cd <policyname>-redhat
```

4. Use **sudo** to run the agent installation shell script using whatever shell you choose, adding the **-n** option if you do not want the blocked file notifier installed. For example, to use the Bourne shell to install an agent:

```
sudo sh ./b9install.sh
-or for installation without the notifier-
sudo sh ./b9install.sh -n
```

5. If you run anti-virus software, exclude the App Control Agent installation directory from anti-virus scanning. For enhanced security, App Control self-protects its own application directory. To avoid performance problems, use whatever mechanism is provided by your anti-virus software vendor to specify that the following directories or files are not scanned:
 - **/opt/bit9/bin** – the agent application and uninstall script
 - **/srv/bit9/data** – the agent database and diagnostics logs
 - **/lib/modules/kernelversion/kernel/lib/b9kernel.ko** – the agent kernel
 - **/etc/rc*/*b9daemon** and **/etc/init.d/b9daemon** – the agent startup script

- `/etc/X11/xinit/xinitrc.d/90b9notifier.sh` – the App Control blocked file notifier
6. Firewalls may recognize App Control software as a new application and block access to the network. Instruct users running the agent to permanently allow it access.

See [Chapter 20, “Endpoint Notifiers and Approval Requests,”](#) for a description of what the user sees on a system protected by the agent.

Verifying the Installation

To verify connected computer is running the agent and visible to the server:

1. On the console menu, choose **Assets > Computers**.
2. Examine the Computers page, which lists all computers running agent software, for the name or IP address of each system you want to confirm. You can use the Search box to find each computer of interest.

The screenshot shows the 'Computers' page in the VMware Carbon Black App Control console. At the top, it indicates 'Computers connected: 1' and 'Total computers: 1'. Below this, there are filter options for 'Saved Views', 'Group By', and 'Days Disconnected'. A search bar is present with 'Go' and 'Clear' buttons. The main table has the following data:

Computer Name	Connected	Policy Status	Upgrade Status	Connected Enforcement	Disconnected Enforcement	IP Address	Policy
WORKGROUP\APC-99986	●	Up to date	Up to date	High (Block Unapproved)	High (Block Unapproved)	10.210.65.144	APC-99986-ucdfjp

3. Note the computer's policy. If it was assigned by Active Directory, the policy will have dashes at the beginning and end of its name. Also note the Connected and Policy Status columns to determine whether the machine is up to date.

Note

During file initialization for a newly installed agent, the computer is already protected at the Enforcement Level associated with its policy.

Verifying Installation on the Agent Computer

You also can verify the presence of the App Control Agent locally on the agent computer:

- On Windows computers, open the **Task Manager** and click on the **Services** tab. You should see **B9Daemon** running.
- On Mac computers, run **Activity Monitor** and view **All Processes**. You should see **b9daemon** running.
- On Linux computers, use `ps aux | grep b9` in a command window. You should see **b9daemon** running.

Upgrading App Control Agents

Beginning with App Control 8.1.4, agent installers and the rule file that determines their behavior are no longer included as part of a App Control Server installation. You upload the rule file and agent installer packages separately after you install the server.

Once new agents are available on the server, there are several ways to upgrade the agent on endpoints:

- Enable automatic agent upgrades on a per-policy basis, allowing the server to manage the upgrade process.
- Initiate agent upgrades on one or more specific computers from the console.
- Manually upgrade agents on the agent machine.
- Use your standard software distribution system to manage upgrades.

Feature Limitations for Non-Upgraded Agents

You can continue to run some older agents as long as they are supported versions and are fully patched. However, it is best to upgrade your agents as soon as possible. In addition to including new features, each agent release generally includes performance and security enhancements.

The console displays a message when the presence of older agents affects the data shown or actions possible on a particular page.

Enabling Automatic Agent Upgrades

When new agent installers are added to the server, the flag that triggers the automatic agent upgrade process is set to “Disabled”. After you have upgraded the server, follow these steps to enable automatic upgrade of agents on systems connected to the server:

- For each policy whose agents you *do not* want to upgrade now, make sure the **Allow upgrades** box in the Options section of the Add/Edit Policy page is *not checked*.
- For each policy whose member agents you do want to upgrade, check the **Allow upgrades** box in the Options section of the Add Policy or Edit Policy page. Avoid doing this all at once for a large number of agents (see the note below).
- On the System Configuration/Advanced Options tab, check **Automatic Agent Upgrades**.

Important

- Before you re-enable system-wide agent upgrades, be sure you *disable* upgrades for policies you don't want upgraded immediately.
- Simultaneous upgrade of a large number of agents may impact system performance. Contact Carbon Black Support for best practices for bulk agent upgrades.
- When a App Control Server is upgraded, ongoing enhancements to “interesting” file identification make it necessary to rescan the fixed drives on all agent-managed computers. These upgrades also require a new inventory of files in any trusted directories to determine whether there are previously ignored files that are now considered interesting. This process can cause considerable input/output activity, which can require between minutes and many hours, depending upon the number of agents and the number of files.
For both upgrades managed by the App Control Server and those using third-party distribution methods, Carbon Black recommends a gradual upgrade of agents to avoid an unacceptable impact on network and server performance.

Upgrading Immediately from the Console

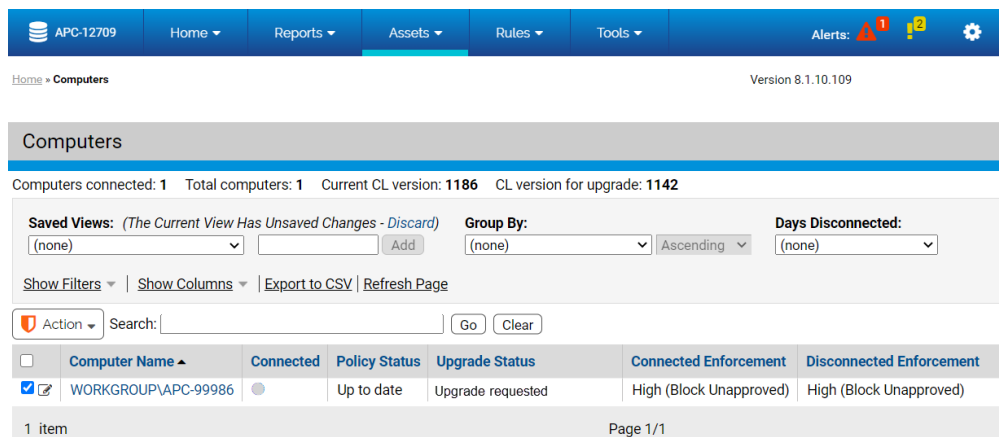
From the console, you can *enable* automatic agent upgrades to happen as part of the App Control Server's regular maintenance of computers, but you can also *force* upgrade of an agent through the console. This has the same effect as running the upgrade from the installer file. Use of this feature requires the following:

- Automatic Agent Upgrades must be Enabled on the Advanced Options tab of the System Administration page. The Upgrade Computers choice does not appear on the menu unless this is enabled.
- The agent(s) must be at least at version 7.0.0 – upgrades from older agents are not supported.

To immediately upgrade one or more agents from the console:

1. On the console, click the configuration (gear) icon and choose **System Configuration** and then click on the **Advanced Options** tab.
2. On the Advanced Options tab, if the Automatic Agent Upgrades field is *Disabled*, click the **Edit** button, choose **Enabled** from the Automatic Agent Upgrades menu, and then click **Update** to make the change.
3. On the console menu, choose **Assets > Computers**.

- Find the computer(s) you want to upgrade and check the checkboxes next to their names. Check the Upgrade Status to make sure the computers are capable of upgrade and not already up to date.



APC-12709 Home Reports Assets Rules Tools Alerts: 1 2

Home » Computers Version 8.1.10.109

Computers

Computers connected: 1 Total computers: 1 Current CL version: 1186 CL version for upgrade: 1142

Saved Views: (The Current View Has Unsaved Changes - Discard) (none) Add Group By: (none) Ascending Days Disconnected: (none)

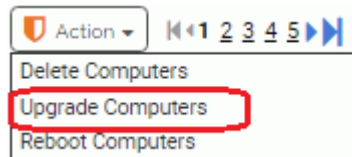
Show Filters Show Columns Export to CSV Refresh Page

Action Search: Go Clear

<input type="checkbox"/>	Computer Name	Connected	Policy Status	Upgrade Status	Connected Enforcement	Disconnected Enforcement
<input checked="" type="checkbox"/>	WORKGROUP\APC-99986	<input type="radio"/>	Up to date	Upgrade requested	High (Block Unapproved)	High (Block Unapproved)

1 item Page 1/1

- In the Action menu, select the **Upgrade Computers** command.



Action

Delete Computers

Upgrade Computers

Reboot Computers

- In the confirmation dialog, click **OK** to trigger the upgrade. Watch the description of the computer in the table to see when the change is completed.

Note

Agents disconnected from their server at the time of a console-based “immediate” upgrade will be upgraded the next time they are connected.

Manually Upgrading Agents

For disconnected systems or if you are using a software distribution system such as SCCM or Altiris to distribute upgrades, you will have to distribute App Control Agent installation files to the endpoints or distribution server.

Installation files for agent are located on the App Control Server in **Program Files\Bit9\Parity Server\hostpkg** on 32-bit systems and **Program Files (x86)\Bit9\Parity Server\hostpkg** on 64-bit systems.

Manually Upgrading Windows Agents

In this release, you use **ParityHostAgent.msi** for *all* manual Windows agent upgrades.

Important

- Manual upgrades must be run either by Local System or by a user account that has administrative rights and a loadable user profile.
- Manual upgrades must use a full path to the installer in the MSIEXEC command.

When a App Control Server manages upgrades to 8.5.0 agents, the agents receive a new list of rules. For manual agent upgrades and upgrades using a third-party distribution method, major upgrades require that the file containing the new rules, **configlist.xml** be copied to a location accessible to the agent installer. On the App Control Server, this file is located in the same **hostpkg** folder as the agent installer but it does not have a link on the Downloads page. It must be manually copied or referenced with a URL or path in the installer.

The following procedure assumes that you take the default choices for parameters that would be used by an automatic upgrade run via the server. [Table 13](#) on page 140 shows parameters that can allow non-default configurations of an installation using MSIEXEC.

To do a Windows agent upgrade manually or using third-party tools:

1. Log in to the console on the computer to which you want to download the installer.
2. On the console menu, choose **Rules > Policies** and click on the download agent software link at the top of the Policies page
3. Download the agent upgrade installer file **ParityHostAgent.msi** to the location from which you want to run or distribute the upgrade

Note that this is not listed on the Downloads page in the console. You can do the download by substituting the file name at the end of the path, using a URL, UNC path, or any other standard means of getting to the file.

For example, to use a URL, you can choose **Rules > Policies** in the console, click on the Download link at the top of the page, and edit the URL for the download page as follows:

https://<your server name>/hostpkg/pkg.php?pkg=<installerfile>

4. Choose the Save option provided by your browser.
5. Follow the same procedure to download the new App Control rules list **configlist.xml** to a location accessible to the agent installer, or make sure the agent installer system can access the *hostpkg* folder on the App Control Server. To use a URL, you would enter the following in a browser on the computer to which you want to download the file:

https://<your server name>/hostpkg/pkg.php?pkg=configlist.xml

Note: If you are using a command line argument to upgrade the agent, you do not necessarily have to *download* configlist.xml. You can use the URL above as an argument in the command line. See [Step 7](#).

6. If you are upgrading a single computer manually, move the configlist.xml file to the agent data folder, usually **C:\ProgramData\Bit9\Parity Agent**, and then run the installer, for example, **ParityHostAgent.msi**.

- If you are preparing to upgrade agents via a third-party distribution system, you can use that system to distribute the configlist.xml file to the agent folder on all agents, or you can use command line arguments in MSIEXEC to include the new rules file in the upgrade installations. A command line for such an upgrade using ParityHostAgent.msi might look like the following:

```
msiexec /fvamus <path>\ParityHostAgent.msi B9_CONFIG=
https://<serverIP>/hostpkg/pkg.php?pkg=
configlist.xml /L*v+ c:\ParityHostAgentUpgrade.log
```

You can use a URL, a UNC path, or a full local path in the command to specify the location of configlist.xml. You cannot use a relative path or a file name without a path.

Important

Certain agent releases may come with special instructions that supersede or supplement the standard installation instructions. If in doubt about how to install an upgrade, consult the [User Exchange](#) or contact Carbon Black Technical Support.

Manually Upgrading Mac Agents

Important

Before manually upgrading Mac agents, you must first install the up-to-date Host Package and rules, as described in “[Uploading Agent Installers and Rules to the Server](#)” on page 131. The following procedure assumes those steps have been completed.

To manually upgrade a Mac agent (agent version 7.2.0 or later):

- Log in to the App Control Console as an administrator, navigate to the **Assets > Computers** page, and click on the computer name or View Details link for the computer you intend to manually upgrade.
- On the Computer Details page, either:
 - Disable tamper protection by clicking **Disable Tamper Protection**, located on the far right under the Advanced section. It may take a few minutes before tamper protection is disabled on the agent.
 - or-
 - Move the agent to a Disabled mode policy.
- Download the upgrade installer for Mac agents, which is **Bit9MacInstall.bsx**. You can do this by using a URL, UNC path, or any other standard means of getting to the file. Note that this installer is not listed on the Downloads page in the console. To use a URL, you can choose **Rules > Policies** in the console, click on the Download link at the top of the page, and edit the URL for the download page as follows:

```
https://<serveraddress>/hostpkg/pkg.php?pkg=Bit9MacInstall.bsx
```

- Open a Terminal window and change directory to the location where the installer was downloaded (by default, the user-specific Download directory).

```
cd ~/Downloads
```


5. Enter the following command to install the agent:

```
sudo bash Bit9MacInstall.bsx
```
6. If you are not using MDM and are currently on any version of High Sierra, see the section [“Allowing the Agent Kernel Extension \(High Sierra or later\)”](#) on page 143 for additional, mandatory steps to allow the Bit9, Inc. kernel extension.
7. If you have not already done so, enable the Mac System Updates updater, which allows minor updates to the OS to be approved for installation. Be sure you are running at least version 9 of this updater. You enable updaters on the **Software Rules > Updaters** page.

Manually Upgrading Linux Agents

Important

Before manually upgrading Linux agents, you must first install the up-to-date Host Package and rules, as described in [“Uploading Agent Installers and Rules to the Server”](#) on page 131. The following procedure assumes those steps have been completed.

To manually upgrade a Linux agent:

1. Log in to the App Control Console as an administrator.
2. Navigate to the **Assets > Computers** page, and in the row for the computer you intend to manually upgrade, click on the computer name or the View Details link.
3. On the Computer Details page, click **Disable Tamper Protection**, located on the far right under the Advanced section. It may take a few minutes before tamper protection is disabled on the agent.
4. Download the upgrade installer for your Linux version to the system on which you plan to upgrade the agent (or a point from which you can copy it):
 - `Bit9Redhat6Install.bsx` – for 6.x versions of RHEL, CentOS or Oracle RHCK
 - `Bit9Redhat7Install.bsx` – for 7.x versions of RHEL, CentOS or Oracle RHCK
 - `Bit9Redhat8Install.bsx` – for 8.x versions of RHEL, CentOS or Oracle RHCK

You can use a URL, UNC path, or any other standard means of downloading the file. Note that this installer is not listed on the Agent Downloads page in the console.

To use a URL, choose **Rules > Policies** in the console, click on the Download link at the top of the page, and edit the URL for the download page as follows:

```
https://<serveraddress>/hostpkg/pkg.php?pkg=Bit9Redhat{6,7 or 8}Install.bsx
```

5. If necessary, copy the downloaded BSX file to the system where you are upgrading the agent.
6. Open a Terminal window and change directory to the location where the installer was downloaded or copied. For example:

```
cd ~/Downloads
```
7. Execute the following command with the appropriate version of the BSX file:

```
sudo bash Bit9Redhat{6,7,8}Install.bsx
```

Agent Upgrade Status

To make the upgrade process easier to manage, the Computers page in the console provides an Upgrade Status column and also visually differentiates between computers running up-to-date agents and those running previous versions. Also on this page, the Connected column uses different color dots to indicate different agent conditions. Hovering the mouse over the dot provides a text description of the condition. See [Table 16, “Computer Details \(Details page and Computers table\)”](#) on page 161 for more information on these indicators.

Computer Name	Connected	Policy Status	Upgrade Status	IP Address	Policy
MYCORP\DESKTOP-3	●	Approvals out of date	Up to date	10.38.90.101	--Administration--
MYCORP\DESKTOP-7	●	Up to date	Not requested	10.38.90.123	--IT Group--
MYCORP\LAPTOP-5	●	Up to date	Upgrade requested	10.38.90.167	--R&D Group--
MYCORP\LAPTOP-2	●	Up to date	Up to date	10.38.90.145	--Sales Group--
MYCORP\SERVER-1	●	Up to date	Up to date	10.38.90.189	--IT Group--

In addition, the Upgrade Status column in the Computers table shows a more detailed description of agent status as each agent goes through the upgrade process. Clients will transition to an Upgrade Status and Policy Status of “Up to Date” when all their upgrade processing has been completed. [Table 14](#) shows the possible Upgrade Status values.

An upgraded agent begins running immediately. You usually do not need to reboot the agent computer, but there are cases in which you may see an Upgrade Status is “Reboot Required”:

- Some Windows XP/2003 systems must be rebooted after upgrade to assure proper ordering of processes and enforcement of rules on systems using DFS.
- On any Windows version, if a file is in use by another process when the agent installer attempts to write that file, you must reboot the computer to allow the system to replace the old file with the current version.

Table 14: Upgrade Status Messages

Upgrade Status	Description
Not Requested	Agent can be upgraded but upgrades are not enabled for the policy, or they are turned off globally.
Upgrade waiting	Agent can be upgraded and is in a policy that allows upgrade. Waiting to be scheduled by server.

Upgrade Status	Description
Upgrade scheduled	Agent has been scheduled for upgrade, or computer has downloaded the upgrade package and not run it yet. Note that the server does not track <i>when</i> the agent upgrade package is downloaded and run.
Upgrade requested	An agent upgrade for this computer was requested from the console.
Reboot required	Agent is waiting for a reboot after upgrade. Reboot is required only under certain conditions (see note above).
Not supported	Agent cannot be upgraded because the computer is running Windows 2000 or another operating system not supported for 7.2.
Upgrade blocked	Agent configuration list is not up-to-date and is missing one or more values required for a successful upgrade. One example of this is use of an out-of-date port number for communication with the App Control Server. Agent cannot upgrade through the server until the configuration is up-to-date, but can be upgraded through other means. In most cases, a connected agent will eventually reach the required configuration list version without intervention. Prioritizing the agent for updates (on the Computer Details page Action menu) expedites configuration list updates. If an agent still remains in "Upgrade blocked" for an extended period, contact Carbon Black Support.
Up to date	Agent upgrade (or new installation) has been completed.
Agent uninstalled	Agent was on this computer but has been uninstalled.

Uninstalling App Control Agents

Standard un-installation procedures delete all App Control files, including the notifier program and drivers. Computer users are not permitted to uninstall an enabled agent unless they have special agent administrative access as described in [“Configuring Agent Management Privileges”](#) on page 727.

To uninstall, you must disable the agent by placing the computer in a policy that is in Disabled mode, which can be done on the Computers page. If you have not already done so, log in to the App Control Console and create a policy with its Mode set to **Disabled** before attempting to uninstall any agents. You might name this “Agent Disabled” or “Ready to Uninstall”. When you create the policy, the server automatically creates an agent installer for it and adds the installer to the Download Install Packages page.

Uninstalling the Agent from a Windows Computer

To uninstall the agent:

1. From the console, find the computer on the Computers page and move it into the agent disabled policy.
2. On the client computer, shut down all other applications.

3. On the client computer, run the standard program removal procedure from the Windows Control Panel:
 - a. On the Windows Control Panel, choose **Add or Remove Programs**, or for Vista or Windows 7 systems, **Programs and Features**.
 - b. From the list of programs, select **App Control Agent**.
 - c. Click the **Remove** button or **Uninstall** button (depending upon your operating system) and wait for the uninstall to complete.
4. Delete the computer from the Computers page. This tells the App Control Server that the computer is no longer in service (rather than temporarily disconnected from the network) and removes its name from the table of active computers.

Uninstalling the Agent from a Mac Computer

1. From the App Control Console, move the computer into the agent disabled policy.
2. In a Terminal or another shell interface, run the following command:

```
sudo /Applications/Bit9/uninstall.sh
```

The agent and its data are removed.
3. Delete the computer from the Computers page. This indicates to the App Control Server that the computer is no longer in service (rather than temporarily disconnected from the network) and removes its name from the table of active computers.

Uninstalling the Agent from a Linux Computer

1. From the App Control Console, move the computer into the agent disabled policy.
2. On the client computer, login with administrator privileges or an account that can run sudo.
3. In a shell window, change to the agent application directory:

```
- cd /opt/bit9/bin
```
4. Run the uninstall script:
 - To remove the agent and all of its data:

```
sudo sh ./b9uninstall.sh
```
 - To remove the agent but preserve agent data in **/srv/bit9**:

```
sudo sh ./b9uninstall.sh -d
```
5. Delete the computer from the Computers page. This indicates to the App Control Server that the computer is no longer in service (rather than temporarily disconnected from the network) and removes its name from the table of active computers.

Viewing the Table of Computers

The Computers page provides a table of computers and information about them, including their platform, policies, Enforcement Levels, and whether they are currently connected to the server. As with most console tables, you can add or remove details shown in the table using the Columns button. You also can use the Filters panel or Search field to limit the computers shown to those you are most interested in. For more about customizing your view, see [Console Tables](#) in [Chapter 2, “Using the Console.”](#)

In addition to the table of agent-managed computers, the Computers page shows the following information:

- **Computers connected** – Shows the number of computers running the agent that are currently connected to the server.
- **Total computers** – Shows the total number of computers that are currently members of security policies managed by the server.
- **Current CL version** – Shows the version number of the latest Configuration List (CL) available from the server. This can be used to help determine whether the CL for a particular agent is out of date. Note, however, that some CL versions are agent-specific, so the fact that the CL version for an agent doesn't exactly match the CL version shown here does not automatically mean the agent is out of date.
- **CL version for upgrade** – Shows the minimum CL version that an agent must have before it can be upgraded. Earlier versions might be missing rules, such as approvals needed for upgrade packages, required for upgrade.

To view the table of computers managed by your App Control Server:

1. In the console menu, choose **Assets > Computers**. The Computers page appears:

2. The Search field provides a way to search for computers by name (or partial name), IP Address, or Policy to reduce the length of the Computers table and help you find the systems you want. Enter the string you want to match against computer names and then click **Go**. Click **Clear** to restore the list that appeared prior to the search.
3. Saved Views provide another way to limit the Computers table to systems matching certain characteristics:
 - Choose **Active Computers** to see currently active (not deleted) computers.
 - Choose **Carbon Black EDR Deployments** to see computers grouped by whether they have had a Carbon Black EDR sensor installed on them.
 - Choose **Cloned Computers** to see computers that have been cloned from a template computer. See [Chapter 6, “Managing Virtual Machines,”](#) for details.
 - Choose **Computers in Local Approval** to see previously locked down computers that have received approval from the server to install software in Local Approval mode.

- Choose **Computers Requiring Upgrade** to see computers running agents that are not up to the current version.
 - Choose **Connected Computers** to see only computers running agents that are currently connected to the server.
 - Choose **Disconnected Computers** to see computers running agents that are not currently connected to the server.
 - Choose **Duplicate Computers** to see computers that have the same name as other computers in your App Control database. See [“Duplicate Computers”](#) on page 174 for details.
 - Choose **Template Computers** to see computers that are templates for cloned computers. See [Chapter 6, “Managing Virtual Machines,”](#) for details.
 - Choose **(none)** to return to the complete list of computers managed by this server.
 - Other Saved Views may be available if you or another console user created them.
4. You can click on **Show Filter** and/or **Show Columns** to open the Filters and Column Settings interface, which let you further customize your view of the Computers table.

[Table 16](#) provides descriptions of the fields available on the Computer Details page, most of which are also available in the Computers table, either by default or by customization.

Agent Policy Status

The Computers table includes a column called “Policy Status,” which indicates whether the agent for each listed computer is up to date with the App Control Server rules that apply to it. Note that this field does not appear on the Computer Details page.

Note

During system initialization, the computer is already protected at the Enforcement Level associated with its security policy.

Computers						
Computers connected: 61 Total computers: 94 Current CL version: 824814 CL version for upgrade: 821882						
Saved Views: (The Current View Has Unsaved Changes - Discard)		Group By:		Days Disconnected:		
(none) Add		(none) Ascending		(none)		
Show Filters Show Columns Export to CSV Refresh Page						
Action 1 2 3 4 5 Search: Go Clear						
Computer Name	Connected	Policy Status	Upgrade Status	IP Address	Policy	
MYCORP\DESKTOP-3	●	Approvals out of date	Up to date	10.38.90.101	--Administration--	
MYCORP\DESKTOP-7	●	Up to date	Not requested	10.38.90.123	--IT Group--	
MYCORP\LAPTOP-5	●	Up to date	Upgrade requested	10.38.90.167	--R&D Group--	
MYCORP\LAPTOP-2	●	Up to date	Up to date	10.38.90.145	--Sales Group--	
MYCORP\SERVER-1	●	Up to date	Up to date	10.38.90.189	--IT Group--	

[Table 15](#) shows the possible values of Policy Status.

Table 15: Policy Status Messages

Policy Status	Description
Up to date	Agent Enforcement Level, policy, and rules are all up to date.
Policy out of date	Agent is not up to date on changes to its policy.
Approvals out of date	Agent rules (including file approvals or bans, trusted users, publisher rules, updater rules, device rules, memory rules and registry rules) are out of date.
Enforcement Level out of date	Agent Enforcement Level is out of date.
Out of date	Agent is out of date on more than one of these: Enforcement Level, policy, or other rules.
Unprotected	Agent is unprotected because of an upgrade failure.
Yara rules out of date	Agent does not have up-to-date Yara rules. Yara rules help identify malicious files.

Actions on Selected Computers

The Action menu on the Computers page has commands that can be applied to one or more computers. You select a computer for action by checking the box to the left of its row.



The actions include deleting, upgrading (if enabled), tagging, and rebooting computers, and moving the computers to different policies. Additional commands available on this menu are described in the Actions and Advanced sections of [Table 18, “Computer Details page menus”](#), on page 166.

Viewing Complete Details for One Computer

There are several ways to locate a computer and display its details. You can use the Find Computer portlet on the Home Page to locate the computer and then drill down to its details. The following procedure describes how you can locate and get details for a computer through the Computers page.

Note

If you request details for a Template Computer, clicking the View Details button shows a Template Details page, not a Computer Details page. See [Chapter 6, “Managing Virtual Machines,”](#) for more information.

To view the Computer Details page for a computer:

1. In the console menu bar, choose **Assets > Computers**. The Computers Page appears.
2. In the Computers table, locate the computer for which you want complete details (for example, using the Computer filters panel).
3. In the table, click either the name of the computer or the View Details button next to its name. The Computer Details page appears:

Computer Details

General

Computer Name: WORKGROUP\APC-99986
 IP Address: 10.210.65.144
 Connection Status: Connected
 Health Check: Passed
 Platform: Windows
 Description:
 Computer Tag:

Policy

Policy: APC-99986-ucdjfp
 Policy Mode: Control
 Connected Enforcement: High (Block Unapproved)
 Disconnected Enforcement: High (Block Unapproved)

Carbon Black App Control Agent

CL Version: 1186
 Debug Level: None (default)
 Carbon Black App Control Agent Version: 8.1.8.306
 Enabled Trusted Directories: 0
 Tamper Protect: Enabled

Save Cancel

Related Views

- Recent Events
- Health Check Event
- Files on this Computer

Actions

- Change Policy
- Delete Computer
- Prioritize Updates
- Add Files to Snapshot

Advanced

- Convert to Template
- Set Debug Level
- Configure Agent Dumps
- Disable Tamper Protection
- Change Local State
- Perform Cache Consistency Check
- Other Actions

4. The General and Policy sections of the Computer Details page appear in all views. The bottom panel on the page varies depending upon the tab you click:
 - Click **App Control Agent** (the default, shown above) to view version and other configuration information for the agent on the Computer whose details you are viewing.

- Click **Connection History** to see the status of the agent's communication with the App Control Server, including whether it fully initialized and synchronized with the server ("Synchronized" appears only after initialization is complete).

Carbon Black App Control Agent	Connection History	Policy Override	System Details	AD Details	Carbon Black EDR
First Registered:	Jun 23 2020 05:33:54 PM				
Last Polled:	Jun 29 2020 10:56:46 AM				
Last Register Date:	Jun 25 2020 04:52:56 PM				
Initialization:	Complete				
Synchronization:	100%				
Server Backlog:	0 files				
Last Logged In User(s):	WORKGROUP\APC-99986\$ APC-99986\bit9qa				

- Click **Policy Override** to generate an override code that can be used to temporarily reassign the agent to a different Enforcement Level.

Carbon Black App Control Agent	Connection History	Policy Override	System Details	AD Details	Carbon Black EDR
Temporary Enforcement:	Local Approval				
Enforcement Level Active For:	30 Minute(s) (up to 500)				
Code Valid For:	5 Minute(s)				
Generate Code					

- Click **System Details** to get any available information about the CPU, memory, and operating system of the computer.

Carbon Black App Control Agent	Connection History	Policy Override	System Details	AD Details	Carbon Black EDR
Computer Model:	VMware7,1				
Processor:	Intel(R) Xeon(R) CPU E5-2698 v4 @ 2.20GHz, 6 CPUs, 2.19 GHz				
Installed Memory:	15.75 GB				
Operating System:	Microsoft Windows 10 x64 Enterprise (10.0.19041)				
Virtualized:	Yes (VMware)				

- Click **AD Details** to see any information Active Directory provides about this computer (only available if you have AD integration activated).

Carbon Black App Control Agent	Connection History	Policy Override	System Details	AD Details	Carbon Black EDR
Distinguished Name:	CN=SERVER-4,OU=Servers,DC=mycorp,DC=local				

- Click **Carbon Black EDR** to see details reported about this computer by the server configured on the App Control System Configuration page Licensing tab. If a Carbon Black EDR server is not configured or the computer is not running a Carbon Black EDR sensor, this tab shows only a status of *Not installed*. By default, the App Control Server checks Carbon Black EDR status every 30 minutes.

Carbon Black App Control Agent	Connection History	Policy Override	System Details	AD Details	Carbon Black EDR
<p> Sensor Version: 5.2.060922 Last Status: Running Uptime: 93 hour(s) Registration Time: Feb 19 2016 10:29:17 AM Last Checkin: Dec 05 2016 6:26:19 PM Next Checkin: Dec 05 2016 6:26:19 PM </p>					

Table 16: Computer Details (Details page and Computers table)

Field	Description
Computer name	Network name for the computer.
IP address	IP address for the computer. This may be an IPv4 or IPv6 address – if the App Control Server is configured for IPv6, agents will attempt to connect via IPv6 first.
Identifier	MAC address for the computer. (Option in table only)
Connection status	<p>Status of computer's communication with the App Control Server:</p> <ul style="list-style-type: none"> • Connected – in communication with the App Control Server. • Disconnected – not communicating with the App Control Server. <p>The Computers table also includes a circle icon in the Connection status field that indicates connection and agent status:</p> <ul style="list-style-type: none"> ● (Blue) – Connected, up to date ● (Gray) – Disconnected, up to date ● (Yellow) – Connected, unsupported (agent out of date, requires reboot, or other reasons) ○ (Clear with Gray Border) – Template computer ● (Red) – Connected, health check failed; indicates that the agent needs immediate attention. Collect the Health Check Events for this computer and contact Carbon Black Support. Red may also appear if the agent is unprotected because a reboot is required or the kernel on the system is unsupported.
Health Check	<p>Agent health status. The health check includes a series of tests to see whether the agent is working properly. If the value is Passed, there are no known health issues with the agent on this computer. If the value is Failed, there is an issue with at least one aspect of agent health. In this case, click Health Check Events on the Computers Details page and contact Carbon Black Support.</p> <p>Note: Health checks run automatically, but if you have addressed an agent problem and want to be sure the agent is healthy, you can force a health check using the Run health check command on the Other Actions menu of the Computer Details page.</p>
Platform	The basic operating system platform of this computer. Possible values are Windows, Mac, and Linux. The System Details tab of the Computer Details page shows additional detail.

Field	Description
Days Offline	If a computer is disconnected, adding this column to the Computers table shows how long it has been disconnected, and allows filtering by number of days.
Upgrade status	Agent upgrade status of this computer. See “Agent Upgrade Status” on page 153 for status options. On the Computer Details page, only appears for computers requiring upgrade.
Upgrade error time	If an error occurred on agent upgrade, the time of that error. On the Computer Details page, only appears for computers on which an upgrade was attempted.
Policy status	Status (up-to-date or not, etc.) for the policy protection of this computer. See “Agent Policy Status” on page 157 for details.
Description	Optional information about this computer, displayed on the Computer Details page. When entering or editing this text on the details page, click the Update Computer button to save.
Computer tag	Optional text string you can add to identify groups of computers that you might want to get reports about or treat in a particular way. A tag offers an alternative to policies as a way to identify groups of computers. For example, you might want to apply a Low (Monitor Unapproved) policy to all computers in your office but be able to track file activity in more specific reports for computers in tagged subgroups such as sales or accounting. Tags may be set on the Computer Details page for one computer or on the Computers page Action menu for multiple computers.
Policy	Currently assigned policy for the computer.
Policy Mode	Security mode in which this policy is operating. The choices are Visibility, Control, and Disabled.
Connected Enforcement	Assigned Enforcement Level while the computer is in communication with the App Control Server. To change this setting for this computer and its fellow policy members, edit the policy. If the Enforcement Level is not up to date with changes to the policy on the server, “(out of date)” will be appended.
Virtualized	Indicates whether this computer is a virtual machine (Yes, No). On the Computer Details page, this is combined with Virtual Platform into a single field on the System Details tab.
Virtual Platform	If this is a virtual machine, the virtualization platform used to generate it. Current values are blank, VMware, and Unknown. On the Computer Details page, this is combined with Virtualized into a single field on the System Details tab.
Clone Inventory	For a template computer, shows whether the inventory for clones created from this template includes All Files (including those from the template image) or just New and Modified Files (since creation of each clone). Blank for non-template computers. See Chapter 6, “Managing Virtual Machines,” for more details.

Field	Description
Inventory	If this is a virtual machine, shows whether the inventory for this clone includes All Files (including those from the template image) or just New and Modified Files (since creation of this clone). Field is blank for non-clone computers. See Chapter 6, “Managing Virtual Machines,” for more details.
Save (button)	Applies changes made to the Description and Computer tag in the General panel of the Computer Details page.
Cancel (button)	Clears unsaved changes made to the Description and Computer tag if you click it before you click the Save button. Page reverts to the settings in effect before you began editing.

Table 17: Computer Details page: Tabbed sections

Field	Description
App Control Agent tab	
CL Version	Configuration List version number indicates synchronization of a computer with server rules. If not current, “(out of date)” appears with the number. Compare the computer’s CL version with the current server CL version on the Computers page. Details pages for many rules also shows the CL version in which the current rule definition was introduced. For use with Carbon Black Support. Note: Rarely, you may see this message next to the CL Version: <i>Agent did receive but is not enforcing all the rules yet.</i> This means the agent is still processing the rules it received, and some rules may not be fully functional. The message (and the state it represents) disappears within a few minutes.
Debug Level (Agent Debug Level in table)	Shows current debug level for this agent, indicating the amount of debugging information collected from it. This can be changed on the Advanced menu. For use with Carbon Black Support.
App Control Agent Version	Version number of the agent installed on this computer.
Enabled Trusted Directories	Number of Trusted Directories now enabled on this computer. See “Approving by Trusted Directory” on page 271 for details.
Tamper Protect	Status of agent tamper protection features (Enabled or Disabled).
Connection History tab	
First Registered	Date and time this computer first registered with its server.
Last Polled	Date and time this agent last polled the App Control Server for updated information and provided updated file information to the server. Agents may poll every 30 seconds, or as seldom as every 10 minutes if the agent is in “sleep” state because the server has no new information about policy changes, approvals, etc.
Last Register Date	Date and time the agent last connected to the server.

Field	Description
Synchronization (%Synchronization in table)	Percent of file information synchronization between this agent and its App Control Server. Appears only after initialization is complete.
Initialization (% Initialization in table)	During initialization, shows the percent of initialization that is complete. Shows as "Complete" after initialization reaches 100%.
Server Backlog	The number of files received from this computer but not yet fully processed on the server. Backlogged files appear in the File Catalog but not in the Files on Computers tab or Find Files page.
Last logged in user(s)	User(s) logged in when the computer last connected to the App Control Server. If AD integration is enabled, click this field for more information about the user.
Policy Override tab	
	Allows generation of a code to temporarily change the Enforcement Level of a disconnected computer. See "Using Timed Policy Overrides" on page 297.
System Details tab	
Computer Model	Model of this computer. Also identifies virtual machines.
Processor	Model, speed, and number of processors for this computer.
Installed Memory	Amount of memory installed on this computer.
Operating System/ Operating System Details	<p>Operating system version on this computer.</p> <p>In the Computers table:</p> <ul style="list-style-type: none"> Operating System shows the basic OS (e.g., Windows 7) Operating System Details includes the full name, the build and service pack level. <p>On the Computer Details page, the Operating System field shows full details.</p>
Virtualized	Indicates whether the computer is a virtual machine, and if so, its platform. Possible values are: No, Yes (VMware), Yes (Unknown)
AD Details tab	
	Clicking this tab shows any additional computer details available through Active Directory. No information is added if AD integration is not enabled or the AD server is unavailable.
Carbon Black EDR tab	
Sensor Version (Carbon Black EDR Version in table)	The version of the Carbon Black EDR sensor installed on this computer.

Field	Description
Carbon Black EDR Status (in table) Last Status (on Details page)	<p>This field shows the last Carbon Black EDR sensor status for this computer, as reported by the App Control Agent to the App Control Server. The App Control Server checks Carbon Black EDR sensor status every 30 minutes, and so status changes may be out of sync for up to that amount of time.</p> <p>The possible values for Carbon Black EDR Status in the table are:</p> <ul style="list-style-type: none"> • Unknown • Installed, initializing – sensor is installed but not fully initialized • Installed, running • Installed, not running • Not installed • Stopped <p>On the Details page, the Last Status field on the Carbon Black EDR tab is similar to Carbon Black EDR Status in the table. However, it does not appear if sensor status is Unknown. Values are:</p> <ul style="list-style-type: none"> • Running • Service not running • Kernel not running • Stopped <p>Notes: In addition to up to a 30-minute gap between sensor installation and App Control polling of Carbon Black EDR sensor status, status will continue to report as <i>Not installed</i> until the Carbon Black EDR sensor connects to the Carbon Black EDR server and receives a sensor id. Also, if the App Control Agent is offline or uninstalled from a computer, the last Carbon Black EDR sensor status reported by the agent is displayed in the console, even if sensor status changes.</p>
Uptime	Number of minutes and hours that the Carbon Black EDR sensor has been running since it was last started.
Computer Status	The status of this computer reported by the Carbon Black EDR server.
Registration Time	The date and time the Carbon Black EDR sensor on this computer registered with its server.
Last Checkin	The date and time the Carbon Black EDR sensor on this computer last checked in with its server.
Next Checkin	The date and time of the next scheduled server checkin for the Carbon Black EDR sensor on this computer.
More Information	<p>Connects to the login page of the Carbon Black EDR server configured on the System Configuration page Licensing tab. Logging in takes you to the Sensors page on the Carbon Black EDR console so you can view additional details about this computer.</p> <p>Note: You must have valid login credentials for the Carbon Black EDR server to successfully open its console.</p>

Table 18: Computer Details page menus

Menu/Options	Description
Related Views menu	
Recent Events	Opens the Events page and shows recent events (if any) for which this computer was the source.
Health Check Events	Opens the Events page and shows health check events for this computer. Use this information for troubleshooting an agent health check failure with Carbon Black Support. You can save the resulting events using the Export to CSV link on the events page.
Files on this Computer	Opens the Find Files page to list all tracked files on this computer.
Carbon Black EDR Details	<p>Opens a new browser window or tab showing the login page of the Carbon Black EDR server configured on the System Configuration page Licensing tab. Logging in takes you to the Sensors page in Carbon Black EDR so you can view additional details about this computer. Link appears only if Carbon Black EDR server is configured.</p> <p>Note: You must have valid login credentials for the Carbon Black EDR server to successfully open its console.</p>
Actions menu	
Change Policy	<p>The dropdown menu provides an alternate way to move the computer into another policy. One of the policies available on this menu is Local Approval, which you can use to temporarily place this computer in Local Approval mode.</p> <p>Click the Go button to apply the change.</p> <p>If this computer had its policy assigned automatically, <i>Automatic</i> shows next to the Go button and the menu is not active. You can un-check the Automatic checkbox to remove automatic assignment and then choose a policy from the menu.</p>
Prioritize Updates/ Remove Prioritization of Updates	<p>Temporarily increases the priority of this computer for receiving upgrades to the agent and configuration lists from the App Control Server. A disconnected host can be prioritized while disconnected and the state will be respected when agent comes online next time.</p> <p>Once a computer has been prioritized, this link changes to <i>Remove prioritization of updates</i>. You also can click <i>Remove prioritization...</i> to downgrade a prioritized computer immediately. Once it is up-to-date in all respects, an agent that had Prioritize Updates applied to it automatically returns to normal priority.</p> <p>An agent may also be assigned permanent prioritization status. This is done automatically for computers hosting Trusted Directories. Permanent prioritization also may be assigned through a command on the Advanced/Other Actions menu. The <i>Remove prioritization...</i> command removes both permanent and one-time prioritization.</p>

Menu/Options	Description
Request Agent Upgrade/Remove Agent Upgrade Request	<p><i>Request Agent Upgrade</i> schedules this agent for an immediate upgrade. Appears only if the agent is eligible for upgrade.</p> <p><i>Remove Agent Upgrade Request</i> removes the upgrade request and so the agent is not forced to upgrade. This appears only if you have previously scheduled an immediate upgrade request.</p> <p>The options apply only to policies with automatic agent upgrades enabled (See “Advanced Configuration Options” on page 744).</p>
Add files to Snapshot	<p>Adds the list of files on this computer (as stored in the App Control Server database) to a <i>snapshot</i> of files. You can use a snapshot to determine how far each of the computers on your App Control Server network have drifted from a baseline of known files. Files in a snapshot can have a variety of statuses; if the snapshot contains banned files, they remain banned. See “Managing Snapshots” on page 653 for more detail.</p> <p>There are two options on this menu:</p> <p>Choose existing snapshot – Adds the list of files on this computer to the snapshot you choose from a menu.</p> <p>Create a new snapshot – Prompts for a new snapshot name and saves the file list of this computer to that snapshot.</p>
Advanced menu	
Convert to Template	<p>Converts the current computer to a App Control computer <i>template</i>, after which clone computers created from the template's image (using third-party virtualization/imaging solutions) can be better managed. See Chapter 6, “Managing Virtual Machines,” for more details.</p>
Set Debug Level	<p>Changes the amount of debugging information collected from the agent on this computer. For use with Carbon Black Support.</p>
Configure Agent Dumps	<p>Changes the amount of information included in file dumps from the agent on this computer. For use with Carbon Black Support.</p>
Disable/Enable Tamper Protection	<p>If agent tamper protection is enabled, clicking Disable Tamper Protection disables it. If protection is disabled, clicking Enable Tamper Protection enables it. Disabling tamper protection is not recommended unless required to solve a particular problem, and the feature should be re-enabled as soon as possible.</p>

Menu/Options	Description
Other Actions submenu	Less frequently needed agent management features, often for use in conjunction with Carbon Black Support. The options are: <ul style="list-style-type: none">• Reboot computer• Upload diagnostic files• Delete diagnostic files on computer• Restore database• Delete database• Restart service• Make local copy of agent cache• Rescan installed applications• Resend all policy rules• Resynchronize all file information• Upload statistics• Run health check• Permanently prioritize updates
Change Local State	This menu allows you to locally approve all unapproved files on the computer. You might choose to do this if you have added a large number of known-good files to a computer after initialization.

Menu/Options	Description
Perform Cache Consistency Check	<p>A cache consistency check ensures that the agent on this computer has accurate information about the files actually present. It is necessary only if the agent was not running during a time when files were written to the computer. If the agent requires updating due to the consistency check, any differences are also sent to the server.</p> <p>Changes in the file cache may affect whether or not a file is approved. You can choose one of three levels of cache consistency checking from the menu:</p> <ul style="list-style-type: none"> • Quick Verification: Confirms that each file in the agent's cache exists, verifies that it is still an executable file that should be tracked, and compares the size of each file on disk to the size stored in its cache the last time the file was analyzed. If a file no longer exists, it is removed from the cache. If any of the other checks fail, the file is re-analyzed. • Rescan Known Files: Does everything in the Quick Verification, plus compares the hash of each file on disk to the same file's hash in the agent cache. If the hash does not match, the file is re-analyzed. • Full Scan for New Files: Does everything in the previous two levels, plus rescans the entire disk, looking for files that should be in the agent cache, but are not. Analyzes any file found. <p>In addition to the menu options, there are three checkboxes that can modify the consistency check:</p> <ul style="list-style-type: none"> • Preserve state of changed files: If the agent does not have a record of a hash in its cache, it will look up the file by name. If that is found, the file state from this record will be used for the current file. • Re-evaluate publishers: Re-examines each file to ensure that its certificate information is accurate and the certificate is not expired or revoked. Also reevaluates trusted publisher approvals. • Approve new files: Locally approve new files found during a full scan. <p>Note: This consistency check is a troubleshooting feature normally used in consultation with Carbon Black Support. Depending upon the option you choose, a cache consistency check could be a time-consuming operation.</p>

Moving Computers to Another Policy

Moving a computer into a different policy is a convenient way to change its protection without creating a new policy. From the Computers table, you can select and move computers into different policies. If you have enabled AD-based policy assignment, you can move computers from manual to automatic policy assignment, and vice versa.

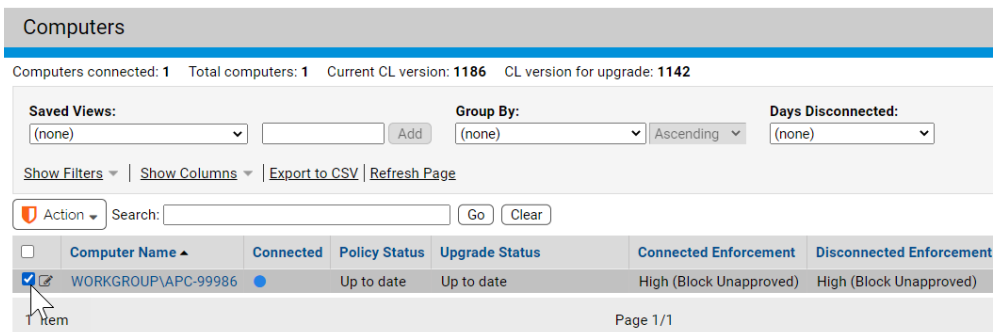
Notes

Changing AD mapping rules *does not* immediately change the policy for an affected computer. The change takes place the next time that computer re-registers with the App Control Server. The section “[Assigning Computers to a Policy](#)” on page 121 lists events that trigger agent computer registration.

In addition to the methods described in this section, you can use the Change Policy portlet on the console Home Page.

To move a computer to another policy:

1. In the console menu, choose **Assets > Computers**. The Computers Page appears:
2. In the Computers table, locate the computer(s) you want to move (using filters or Saved Views, if helpful) and check the associated checkbox for each computer.



Computers

Computers connected: 1 Total computers: 1 Current CL version: 1186 CL version for upgrade: 1142

Saved Views: (none) Add Group By: (none) Ascending Days Disconnected: (none)

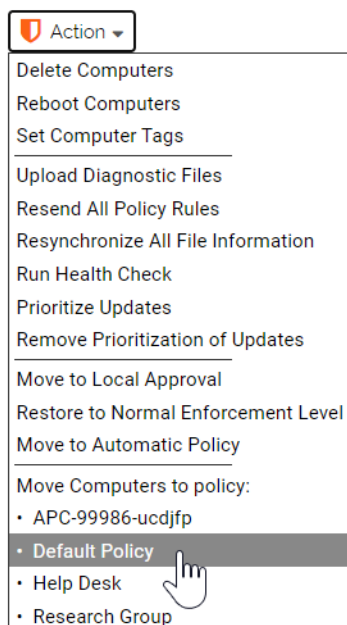
Show Filters Show Columns Export to CSV Refresh Page

Action Search: Go Clear

<input type="checkbox"/>	Computer Name ▲	Connected	Policy Status	Upgrade Status	Connected Enforcement	Disconnected Enforcement
<input checked="" type="checkbox"/>	WORKGROUP\APC-99986	●	Up to date	Up to date	High (Block Unapproved)	High (Block Unapproved)

1 item Page 1/1

3. Click the **Action** button to see the Action menu. The move options begin approximately halfway down the menu.



4. On the Action menu, choose the option that shows the move you want to make. In the confirmation dialog, choose **OK** to reassign the computer to the selected policy. The computer moves to the policy you selected, and if you moved it from Automatic, the policy assignment becomes manual.

Notes

You also can change a computer's policy by clicking on the computer name in the table and using the Change Policy menu on the Computer Details page.

In addition, Event Rules may be created that will automatically change a computer's policy when certain events occur.

Restoring Computers from the Default Policy

The Default policy is for computers that report to the App Control Server but cannot be associated with any other policy. Causes for this include:

- AD mapping is enabled, the default AD mapping rule (the last rule on the list) maps policies to Default Policy, and an agent does not match any other rule.
- An old installer associated with a deleted policy might be used for the initial App Control Agent installation on a computer.
- The last agent in a policy disconnects from the App Control Server and then is deleted from the Computers table on the console; because the policy now has no computers, a console operator decides to delete it. The agent later reconnects to the App Control Server.

In any of these cases, the computer is automatically moved into the Default Policy. Carbon Black recommends that you set the Enforcement Level for the Default policy to the appropriate protection level for your site. If you set the Default Policy to Visibility Mode,

which tracks but does not block file executions, any computers that appear in the Default Policy should be moved as soon as possible to a policy with the settings and Enforcement Level protection you want.

Notes

- If you do not have any full Suite licenses (Visibility and Control), your only Enforcement Level choices for the Default policy are Visibility and Disabled.
- Because the Default Policy is reserved by the system, you cannot delete it.

The procedure for restoring computers from the Default policy is essentially the same as that for moving computers to another policy, with additional filtering instructions.

To move a computer in the Default policy to another policy:

1. In the console menu, choose **Assets > Computers**. The Computers Page appears.
2. If it is not the current choice, choose **(none)** as the Saved View.
3. Click the **Show Filters** link, and on the Add filter menu, choose **Policy**.
4. In the Policy filter, choose **is** as the operator, choose **Default Policy** from the right menu, and click the **Apply** button to apply your filter. All computers in the Default policy appear.

The screenshot displays the 'Computers' management interface. At the top, it indicates 'Computers connected: 1' and 'Total computers: 1'. Below this, there are controls for 'Saved Views' (set to '(none)'), 'Group By' (set to '(none) Ascending'), and 'Days Disconnected' (set to '(none)'). A 'Filters' section shows an active filter: 'Policy is Default Policy'. An 'Action' dropdown menu is open, with 'Default Policy' selected. The table below has columns for 'Computer Name', 'Policy Status', 'Upgrade Status', 'Connected Enforcement', and 'Disconnected Enforcement'. The first row shows a computer named 'APC-99986-ucdjfp' with a 'Policy Status' of 'Default Policy'.

5. From the Computers table, check the checkbox(es) for the computer(s) to be moved. You can check multiple computers if you want to move them from the Default policy to the same non-Default policy.
6. On the Action menu, select the policy to which the checked computers are to be moved. If you are using AD-based policy assignment and you are certain this computer matches one of your mapping rules, choose **Move to Automatic Policy**.
7. In the confirmation dialog, click **OK** to reassign the selected computer to the new policy. This temporarily disconnects the App Control Server from the agents of any computers checked and causes them to reconnect. When reconnected, the computers are associated with the policy you moved them to.

Moving a Computer to Local Approval Mode

When computer users need to install new software and App Control trusted-approval methods (directory, user/group, publisher and updater) are inappropriate, you can temporarily put the user's computer into Local Approval mode, which is a special policy that permits software installation. Executable files introduced to a computer while it is in Local Approval mode become locally approved on that computer unless already banned. Files already on the computer before you enabled Local Approval mode are not locally approved, although there are other methods to approve them.

You move a computer to Local Approval mode either by checking the box next to its name on the Computers page and choosing **Move to Local Approval** on the Action menu, or by choosing **Local Approval** on the Change Policy menu on the Computer Details page. See [“Moving Computers to Local Approval Mode”](#) on page 294 for complete instructions.

Adding Computers

Computers are added to the Computers table on the App Control Console when you install the agent on them and they contact the App Control Server – there is no special “Add Computer” operation required. If you are using AD-based policy assignment, a new computer is assigned a policy based on the rules you set for mapping AD data for a computer (or its users) to security policies. Otherwise, the computer is assigned the policy specified in the agent installation package chosen for it.

Deleting Computers

Computers that are no longer in service or that you choose not manage with an agent may be deleted from the App Control Server. Before you delete a computer from the Computers table in the console, you first change the computer's Enforcement Level to Disabled and then uninstall the agent. See [“Uninstalling App Control Agents”](#) on page 154 for more detail.

If you *do not* uninstall the agent before you delete a computer and that computer remains connected to the same network as your App Control Server, the computer will reappear in the computer table as soon as it polls the App Control Server. If connected to the network, computers immediately return to the table; if off-line, computers return upon reconnection. Deleted computers that continue to run the agent return to their last recorded policy. If you have deleted the policy applied to the computer by its agent installer, the server moves the computer to the Default Policy.


Note

If a computer running the agent cannot connect to the App Control Server and you want to remove its agent, contact Carbon Black Support.


Files on deleted computers remain in the Files on Computers inventory for a short period of time, 24 hours by default. See [“Files on Deleted Computers”](#) on page 719 for an illustration of how these files appear in search results.

To delete a computer from a App Control Server:


1. In the console menu, choose **Assets > Computers**. The Computers page appears.
2. Find the computer you want to delete and check the checkbox next to its name.



3. In the Action menu, select the Move command for your agent disabled policy from the menu (it is shown as “Agent Disabled” below but you can call it anything you want; it must have an Enforcement Level/Mode of *Disabled*).



4. In the confirmation dialog, click **OK** to trigger the policy change. Watch the description of the computer in the table to see when the change is completed.
5. Once the agent for this computer is in the agent disabled policy and displays an Enforcement Level of *Disabled*, delete the agent software from the computer itself.
6. After the agent software is uninstalled, on the Computers page, locate the name of the computer whose agent you removed and check the box next to its name.
7. On the Action menu choose **Delete Computers**.



8. On the confirmation dialog, click **OK** to complete the deletion.

Duplicate Computers

In some cases, duplicate computer names can appear in the Computers table. This can happen when an agent-managed computer is taken offline, reconfigured or repaired, and then has the agent re-installed without having its previous agent uninstalled and its entry deleted from the table. This presents an asset management problem, one that can become much greater in larger organizations with many computers being reconfigured on a regular basis.

To make it easier to identify and eliminate duplicate computer names, the Computers view includes a Saved View called *Duplicate Computers*. This view lists every agent-managed computer whose name is the same as the name of another agent-managed computer, where neither computer has been deleted from the server.

Computers connected: 1 Total computers: 1 Current CL version: 1186 CL version for upgrade: 1142

Saved Views: Duplicate Computers Add Group By: (none) Ascending Days Disconnected: (none)

Show Filters Show Columns Export to CSV Refresh Page

Computer Name	Connected	Connected Enforcement	Disconnected Enforcement ...	IP Address	Policy	Last Poll
Computer Name: MYCORPTEST-4 2 item						
MYCORPTEST-4	●	High (Block Unapproved)	High (Block Unapproved)	10.32.60.32	Research	Dec 6 2016 11:49:56 A
MYCORPTEST-4	●	High (Block Unapproved)	High (Block Unapproved)	10.32.60.27	Research	Nov 4 2016 06:38:30 P

The Duplicate Computers view shows computer grouped by name and includes a *Last Poll* (the date and time when the agent and server last communicated so that you can decide which computer entry represents the currently active agent).

Note

You also can add the column *Duplicate* to any Computers table view to identify which computers are duplicates (the value is Yes) and which are not (the value is No).

Operating System Updates on Agents

Operating System Updates on Windows Agents

Support for operating system changes on systems running the agent depends upon the Windows release you are upgrading from and to:

- **Pre-Windows 10 versions, service packs** – Service pack upgrades for pre-Windows 10 releases are supported with the agent installed, and do not cause health check failures.
- **Pre-Windows 10 versions, major or minor version changes** – Changing the major or minor version of Windows on systems with an agent installed is not supported for upgrades to any pre-Windows-10 version. For example, upgrading from Windows 8.0 to Windows 8.1 without uninstalling and reinstalling the agent after the upgrade is not supported, and doing so will produce health check failures, and in some cases failure of the Windows upgrade.
- **Windows 10 versions prior to the Anniversary Edition**– Upgrading from another major or minor version (such as Windows 8 or 8.1) to a version of Windows 10 prior to the Anniversary Update is not supported with an agent installed.
- **Windows 10, Anniversary Update and later releases** – Upgrading to Windows 10 Anniversary Update (August 2016) and later Windows 10 releases can be done with an agent installed, if you meet certain requirements.

Please note these important requirements, recommendations, and open issues for operating system updates with an agent in place:

- You must enable Trusted Directory approval of WIM files for in-place agent upgrades to succeed, as described in the next section.
- If you do not have another anti-virus product installed, Windows Defender is enabled by default when you install Windows 10. Consider enabling the Windows Defender *updater* on the App Control Console (**Rules > Software Rules > Updaters**) to make sure update files for this application are not blocked.
- Updating the operating system to a version with a different major, minor or build number (for example, if any of the numbers in 10.0.35 change) with the agent in place requires that a cache consistency check be run. This happens automatically after the update. Until this cache consistency check is finished, the file inventory on the agent system will be incomplete.
- Before updating the OS, see *Supported Carbon Black sensors and agents* on the [User Exchange](#) to check current OS support information for your agent version.

If you meet the requirements described in this section, you may leave the v8.0.0 agent installed and upgrade to Windows 10 Anniversary Update (August 2016) or later Windows 10 releases, from the following versions (and their server equivalents):

- Windows 7, 8 and 8.1
- Windows 10 (previous versions)

Enabling Trusted Directory Approval of WIM Files

You can enable “crawling” and approval of the contents of Windows Image (WIM) files in trusted directories. Addition of support for WIM crawling will help increase approval coverage of updates you receive via Windows Server Update Service (WSUS).

One important use of this feature will be to enable App Control to support updates to the upcoming Windows 10 Anniversary Update and later Windows 10 releases on your endpoints without removal of the agent. Although full crawl of the ISO contents ahead of time is not necessary required for Windows 10 updates with the agent in place, it is recommended to avoid upgrade failures due to unexpected blocks of installation files.

See “[Approving by Trusted Directory](#)” on page 271 you have not already set up a trusted directory. If you use multiple trusted directories, the procedure must be repeated for each one in which you want WIM files to be approved.

To allow trusted directory approval of WIM files:

1. Choose or create the trusted directory in which you want to approve the content of WIM files (including those in ISOs). On the system where the trusted directory is located, download and install the Microsoft Windows Automated Installation Kit (AIK) or for Windows 10 the Assessment and Deployment Kit (ADK). The installation image and instructions can be found in the following locations:

For Windows XP and Vista (and server equivalents):

<https://www.microsoft.com/en-us/download/details.aspx?id=10333>

For Windows 7 and 8 (and server equivalents):

<https://www.microsoft.com/en-us/download/details.aspx?id=5753>

For Windows 10 (and server equivalents):

<https://docs.microsoft.com/en-us/windows-hardware/get-started/adk-install>

The AIK (or ADK) includes **ImageX.exe**, which is required for WIM approval. The version of ImageX.exe is the same for both downloads – using the version specified for your OS makes the installation simpler.

2. Disable tamper protection on the agent running on the trusted directory server.
 - a. On the console menu, choose **Assets > Computers**.
 - b. In the Computers table, find the name of the computer hosting the trusted directory, and click on the name or View Details button.
 - c. On the Computer Details page, click on **Disable Tamper Protection** in Advanced section of the right menu bar.
3. From the download location, copy **ImageX.exe** into the agent installation directory (typically C:\Program Files (x86)\Bit9\Parity Agent).

Note: Once you have copied ImageX.exe into the agent directory, you may uninstall the AIK (or ADK) software. It is not required for agent operation.
4. In the App Control Console, approve the ImageX.exe file on the agent hosting the trusted directory:
 - a. On the console menu, choose **Tools > Find Files** and search for **ImageX.exe**.
 - b. In the Find File results, check the box next to ImageX.exe and choose **Approve Locally** or **Approve Globally** on the Action menu.
5. On the Computer Details page for the agent on which you placed ImageX.exe, re-enable tamper protection.
6. Add WIM files to the file types that the agent can “crawl” in a trusted directory:
 - a. Navigate to the Support page in the App Control Console by manually entering the URL: **https://<yourserveraddress>/support.php**
 - b. Click on the Advanced Configuration tab, and in the Agent Configuration panel, check the box for **Enable Deep Crawl**.
 - c. In the next line, **Deep Crawl Files**, add ***.wim** (be sure to include the asterisk) to the end of the list of file extensions if it is not already there. Use a comma to separate the new extension from the previous one in the list. Click **Update** when you are finished.
7. On the console, choose **Assets > Computers**, and locate the computer that has the trusted directory. You must wait until this computer shows **Up to date** in the Policy Status column of the Computers page before proceeding.
8. Copy or move any ISO files and/or separate WIMs you want approved into the trusted directory. The inventory and approval of the contents of these files begins. Completion of this process could take several hours and consume considerable system resources, depending upon your hardware.

For additional info on ImageX.exe, see:

[https://technet.microsoft.com/en-us/library/cc722145\(v=ws.10\).aspx](https://technet.microsoft.com/en-us/library/cc722145(v=ws.10).aspx)

For additional info on WIM, see:

<https://www.microsoft.com/en-us/download/confirmation.aspx?id=13096>

Operating System Updates on Mac Agents

If you plan to update the operating system on a Mac endpoint to a *major* release, the agent on that endpoint should be put into disabled mode before the OS update. This would be true, for example, for upgrading from macOS 10.13 to 10.14. Also keep in mind

that your current agent might not be supported on some versions of macOS. Before updating the OS, see the release notes and also *Supported Carbon Black sensors and agents* on the [User Exchange](#) to check current OS support information for your agent version.

Once the macOS update is completed, move the agent back into to its previous policy. It will begin to reinitialize its files.

To make sure other updates to macOS are allowed, be sure you are running at least version 9 (the current version for the 8.1.4 server) of the Mac Systems Updates updater (see the **Software Rules > Updaters** page).

Upgrading to Mojave with an Agent Installed

Important

For systems running existing agents prior to 7.2.3 Patch 11, do not upgrade to Mojave 10.14 before upgrading to a newer agent (preferably the latest agent but at least 7.2.3 Patch 11). Previous agents are not compatible with 10.14 Mojave.

Updating from one major OS version to another with an agent in place requires disabling the agent during the upgrade. If you are upgrading to Mojave, take the following steps on each system:

1. Complete the agent upgrade to the latest macOS agent (or at least 7.2.3 Patch 11).
2. Place the agent in Disabled mode.
3. Upgrade the operating system to 10.14 Mojave.
4. Once the macOS update is completed, move the agent back into to its previous policy. It will begin to reinitialize its files.

Operating System Updates on Linux Agents

If you plan to update the operating system on a Linux endpoint to a *major* release, the agent on that endpoint should be put into disabled mode before the OS update. This would be true, for example, for upgrading from RHEL 6.8 to 7.3.

For *minor* Linux OS updates, as long as both the current and upgrade version are supported by the App Control agent running on the system, you may leave the agent in its current Enforcement Level. This would be true, for example, for upgrading from RHEL 6.7 to 6.8.

Also keep in mind that your current agent might not be supported on some versions of Linux. Before updating the OS, see the release notes and also *Supported Carbon Black sensors and agents* on the [User Exchange](#) to check current OS support information for your agent version.

Chapter 5

Creating and Configuring Policies

This chapter explains how to create policies and change their settings, including Enforcement Levels.

Sections

Topic	Page
Policy and Enforcement Level Overview	180
Creating Policies	181
Policy Settings	186
Rules Affecting Policies	191
Editing a Policy	194
Related Views in Policy Details	196
Enforcement Levels	197
Locking Down all Computers	201
Deleting Policies	205

Policy and Enforcement Level Overview

Each computer running a App Control Agent is assigned to a App Control *policy*. A policy creates a common file control definition for all of its computers. Each policy consists of a group of settings and an overall Enforcement Level.

Enforcement Level defines how strictly actions defined by the policy settings are controlled, especially for control of file writing and execution. The choices are:

- High (Block Unapproved)
- Medium (Prompt Unapproved)
- Low (Monitor Unapproved)
- None (Visibility)
- None (Disabled)

High, Medium, and Low Enforcement are available only if you have the full App Control license with both Visibility and Control features. Sites whose licenses are all for Visibility Only operation are limited to Visibility and Agent Disabled modes with no enforcement.

In Visibility mode, you can still choose settings that would block activity if you were operating another Enforcement Level, but these settings do not enforce the block or ban.

Policy settings specify the types of files or operations that App Control Agents will control as well as other choices such as how policies are assigned and whether agents on computers in the policy upgrade automatically.

Rules defined on other pages can be applied to specific policies. The details page for each policy includes a tabbed panel showing which rules are applied to that policy.

If you choose, you can restrict the ability of console users to perform certain functions so that it only applies to computers in certain policies. For example, you might want to allow one group of administrators to create rules for your sales team but not for the senior management. If you assigned the computers for your sales team to one policy, you can define a *user role* that grants permission to create and modify rules only for that policy. See [Chapter 3, “Managing Console Login Accounts,”](#) for more information on creating console user accounts and defining user roles for those accounts.

Unified Management Note: If you are using Unified Management to manage multiple App Control Servers, you can apply rules to specific policies on specific computers. See [Chapter 27, “Unified Management of Multiple Servers,”](#) for information on management of unified rules by policy.

Creating Policies

Policies enable you to organize computers running the App Control Agent into groups with common security requirements. For example, you can create policies based on departmental affiliations like sales, marketing, or other organizational relationships. You might also create policies specific to a computer's purpose, such as a special domain controller policy. A single policy may be appropriate if you want a single, company-wide operating standard for all computers, but typically you will create multiple policies.

Policies normally are assigned to computers, not users, although Active Directory data can be used to assign policy by user. Each computer has only one policy at a time, regardless of the number of users currently logged on.

Once a policy is created, you can assign computers to it through a variety of methods, including automatic assignment based on Active Directory group. See [Chapter 4, "Managing Computers,"](#) for more details on policy assignment.

When you create a policy, App Control attempts to create an agent installer that assigns the policy to computers that use the installer. If you have not yet uploaded agent installer packages and a rules file to your server, or if agent installer creation is disabled for all operating systems, creating a policy generates error events indicating that the agent installers for that policy cannot be created. You can still create the policy, but to avoid populating the Events log with errors each time you create a policy, the best practice is to upload agent and rule installers before creating policies. See ["Uploading Agent Installers and Rules to the Server"](#) on page 131 for more information.

Important

Policy names can use alphanumeric characters and certain symbols in the ISO-8559-1 set. Characters in the 32-127 range in the ISO-8559-1 set are legal, with the following exceptions: < > : , " / \ | ? * # @

If you enter Unicode characters or reserved symbols in a policy name, the console displays a warning dialog. You must remove the illegal characters from the name before you can save the policy.

Some characters that are allowable in policy names might cause problems when running the agent installer for the policy. For policies that will be applied to Mac computers, avoid parentheses and spaces in the name, or be prepared to "escape" these characters when you run the installer.

To create a policy:

1. On the console menu, choose **Rules > Policies**. The Policies page appears:

Policy	Connected Enforcement	Disconnected Enforcement	Total	Connected
APC-99986-ucdfjp	High (Block Unapproved)	High (Block Unapproved)	1	1
Default Policy	None (Visibility)	None (Visibility)	0	0
Help Desk	None (Visibility)	None (Visibility)	0	0
Local Approval Policy	Local Approval	Local Approval	0	0
Research Group	None (Visibility)	None (Visibility)	0	0
Template Policy	None (Visibility)	None (Visibility)	0	0

2. On the Policies page, click the **Add Policy** button. The Add Policy page appears (shown below for a Control policy):

Policy Name:
Description:
Mode: Visibility Control Disabled
Initial Settings:
Automatic Policy Assignment For New Computers:
Options: Allow Upgrades Track File Changes
 Load Agent in Safe Mode Suppress Logo In Notifier
Total Computers: 0
Connected Computers: 0

3. On the Add Policy page, enter a policy name and define the other policy parameters as you choose (see [Table 19](#)) – the parameters you see may vary depending upon other policy settings and configuration choices:

Table 19: Policy Definitions: Main Panel

Field	Description
Policy name	<p>Name of the policy.</p> <p>Choose a name that indicates the security level, function, or other common factor for computers or users you want to use this policy.</p> <p>Note: If you change the Policy Name, the new name will be reflected immediately in the console, but the name of the agent installer (the <i>polycyname.msi</i> file) requires approximately one minute to update. Keep this in mind if you intend to download agents immediately after a policy name change.</p>
Description	Optional information about the policy. This can be any text you choose to enter.
Mode	<p>The mode in which the App Control Server interacts with the computers in this policy:</p> <p>Visibility specifies file-tracking only. The App Control Server tracks file activity and events, but file execution and writing is not effected by policy settings or file bans. Enforcement Level menus do not appear in Visibility mode.</p> <p>If you have not purchased Control licenses, Visibility is the only mode choice other than Disabled.</p> <p>You might use Visibility when security features could interfere with operational functions for computers. For example, you might it for a computer on which you plan to configure a Trusted Directory for files you will allow to be installed on all computers.</p>
Mode (cont.)	<p>Control activates the Enforcement Level menus, from which you can choose the level of control over execution of Unapproved and Banned files.</p> <p>Disabled specifies pass-through mode; the agent neither blocks file activity nor reports it to the server. Executables run as if the agent were not installed. Use this setting for uninstalling the agent.</p> <p>File inventory for computers in Disabled mode will not be kept up to date on the server. Some operations are monitored (but not reported to the server) to avoid gaps in file and process information if the agent is later activated.</p>

Field	Description
Connected Enforcement Level	<p>The protection level for computers in this policy while they are connected to the network (menu only appears in Control mode):</p> <p>High (Block Unapproved) is the highest protection level you can set —no Unapproved or Banned files in categories tracked by App Control are allowed to run. Blocked file executions are recorded in the event log.</p> <p>Medium (Prompt Unapproved) blocks Unapproved executables on agent computers but displays a dialog box that gives users the option to permit or block the file execution. Users cannot permit execution of explicitly Banned files.</p> <p>Low (Monitor Unapproved) permits Unapproved executables to run but tracks them. Files allowed to run include running non-executables (such as dlls, com objects and loadable resources), unapproved scripts, and unapproved executables. Events are recorded for the first instance of a permitted file execution and for all blocked executions.</p> <p>At High, Medium or Low Enforcement Levels, determination of which files are blocked also depends on the Advanced Settings within each policy.</p> <p>Visibility and Disabled, for which the Enforcement Level is None, are set from the Mode line.</p>
Disconnected Enforcement Level	<p>The protection level for computers in this policy while they are out of communication with the App Control Server. If the Connected Enforcement Level is Low (or None) the Disconnected Enforcement Level is identical to the Online, and cannot be modified directly. If the Connected Enforcement Level is High or Medium, you can choose an Disconnected Enforcement Level of High or Medium, and it may differ from the Connected Enforcement Level.</p>
Initial Settings	<p>Existing policy that you would like to use as a template for the new policy. Although not visible when you create a policy, the Device and Advanced Settings (only) of the chosen policy are transferred to the new policy. See “Template Policy” on page 192 for more information.</p>
Automatic Policy Assignment for New Computers	<p>When this box is checked, if AD-based policy assignment is enabled and configured, new computers that used the installer for this policy get their policy according to the AD-mapping rules, regardless of the policy embedded in the installation package used to install their agent. When not checked, the install package determines the policy and AD mappings have no effect. See “Assigning Policy by Active Directory Mapping” on page 122 for more details.</p>
Set automatic policy for existing computers	<p>This checkbox appears only if the <i>Automatic policy assignment for new computers</i> box is checked. When checked, if any computers were manually (non-automatically) assigned to the current policy, they are changed to automatic policy assignment.</p>
Set manual policy for existing computers	<p>This checkbox only appears if the <i>Automatic policy assignment for new computers</i> box is checked. When checked, if any computers were automatically assigned to the policy, they are changed to have this policy manually assigned.</p>

Field	Description
Options: Allow Upgrades	If the App Control Server is configured for Automatic App Control Agent upgrades, checking this box causes computers in the policy to be notified of and scheduled for App Control Agent upgrades. Computers moved into this policy (either manually or by Active Directory mapping) also will be upgraded. See “Advanced Configuration Options” on page 744 and the upgrade sections of <i>VMware Carbon Black App Control Installation Guide</i> for more information. For use only during App Control Server upgrades.
Options: Track File Changes	<p>When checked (the default) file changes (files added, deleted, or changed) on a computer are tracked and added to the database for this App Control Server.</p> <p>You might deselect this option to remediate performance issues, perhaps while waiting to upgrade from SQL Express to a full version of SQL Server, or in a special policy for computers whose file activity you don't want to track.</p> <p>Important: If you turn off this feature, the App Control Server deletes the file inventory information for the agents in this policy after one day. The Files on Computers table, Find Files, and Baseline Drift reports will not provide accurate information about these computers. Also, if you turn this feature on after it has been off, this forces re-synchronization of the affected agents to update the file database, and this can have a performance impact.</p>
Load Agent in Safe Mode	<p>Loads the App Control Agent in Safe Mode on computers in this policy if the computer is booted in Safe Mode. In this case, the agent performs all enforcement activities, even with the system in Safe Mode. Full protection requires the agent kernel, which loads at boot time, and the agent itself, which runs as a service after boot time.</p> <p>Caution: Since the agent can interfere with Safe Mode recovery operations, use this option only if you have other means of recovery (other than Safe Mode). If you have questions about enabling the agent to run in Safe Mode, contact Carbon Black Support.</p>
Suppress Logo in Notifier	When App Control rule enforcement causes a notifier to be displayed on an agent system in this policy, do not show a logo, even if the rule's notifier definition includes a logo.
Total/Connected Computers	<p>Total Computers - The total number of computers managed by this policy on the App Control Server. Computers by platform is shown in parentheses.</p> <p>Connected Computers - The number of computers managed by this policy currently connected to the App Control Server. Computers by platform is shown in parentheses.</p>

- After you have provided the policy configuration parameters on this page, click the **Save** button. The new policy appears in the table on the Policies page.
- To modify the Device Settings or Advanced Settings for this policy, click the View Details button next to the new policy name, make your modifications, and click **Save**. See [“To edit a policy:”](#) on page 195 for detailed instructions on editing these settings. Note that Device and Advanced Settings do not appear on the Add Policy page – you must save the policy first to see them.

Notes

For more information about the Device Settings and other device monitoring and control features in App Control, see [Chapter 12, “Managing Devices.”](#)

For information about customizing the notifier displayed on a client computer when policy and ban settings are enforced, see [Chapter 20, “Endpoint Notifiers and Approval Requests.”](#)

Policy Settings

The Enforcement Level for a policy sets the overall security level and determines whether the policy is configured to block or permit execution of Unapproved files. More specific behavior is controlled by detailed policy settings and by rules that apply to a policy.

Detailed policy settings are divided into Device Settings and Advanced Settings. These two settings groups are shown and may be edited on a tabbed panel at the bottom of the Edit Policy page. Advanced Settings are described in the next section of this chapter. Device Settings are described in [Chapter 12, “Managing Devices.”](#)

In addition to the *settings* for a policy, *rules* managed on other pages may be applied to selected policies or all policies. So that you can see which rules apply to a particular policy, a set of tabbed rule views is included on the Edit Policy page. These are described in [“Rules Affecting Policies”](#) on page 191.

Important

Visibility mode allows you to activate settings that block files, but these settings have no effect while a computer is in Visibility mode. To enable file blocking and other control features, a policy must be in Control mode. You still might activate these settings in Visibility mode for information purposes, or if you plan a change to Control mode in the future.

Advanced Settings

When active, advanced settings block specified file activities and enforce other rules.

Name	Status	Notifiers
Block unanalyzed scripts and executables	Active	<default>: Block unanalyzed scripts and executables
Block unapproved scripts	Active	<default>: Block unapproved scripts
Block unapproved executables	Active	<default>: Block unapproved executables
Block banned file names	Active	<default>: Block banned file names
Block banned file hashes	Active	<default>: Block banned file hashes
Block executables run from a network drive	Off	<default>: Block executables run from a network drive
Block files with banned publishers or certificates	Active	<default>: Block files with banned publishers or certificates
Enforce memory rules	Active	<default>: Enforce memory rules
Enforce registry rules	Active	<default>: Enforce registry rules
Enforce custom (file and path) rules	Active	<default>: Enforce custom (file and path) rules
Enforce tamper protection	Active	<default>: Enforce tamper protection
Terminate processes with banned images	Report Only	<default>: Terminate processes with banned images

Locally approve unapproved files on transition from Visibility or Low Enforcement Level to Medium or High

Save & Exit Save Cancel Reset Policy

Because any file or activity is usually affected by more than one rule, turning a setting off can have varying results. There are three possible options for advanced settings:

Table 20: Policy Advanced Setting Options

Setting Options	Description
Active	Setting is enabled. Files are blocked or permitted according to the specified Enforcement Level.
Off	Setting is disabled and not enforced under any Enforcement Level. Files matching the setting continue to be tracked but are not blocked.
Report Only	A test state that permits actions that would have been blocked if the setting were active and records a <i>would-have-blocked</i> event in the Events table. You can use it to verify that settings and Enforcement Level in a policy work as intended, without actually blocking any files.

Turning off one setting that *blocks* an action or file does not necessarily mean the action or file is *permitted*; similarly, turning off one setting that *permits* an action does not necessarily mean that the action or file is *blocked*. The Events page might provide an explanation of why a file you expected to be permitted was blocked.

Table 21 shows the Advanced Settings and the effect of setting them to “Active” or “Off”. Some settings cannot be turned off, but are included so you can change or disable the *Notifier* that appears when they block a file execution.

Notes

- There are different settings for “executables” and “scripts”. App Control determines whether a file is executable based on content, not file extension alone, while scripts are identified by file extension. After examining a file, the App Control Agent applies the appropriate policy setting based on the file’s content. See [Chapter 13, “Script Rules,”](#) for information about how scripts are defined.
- Each setting has a Notifiers menu from which you can choose the notifier that appears on an agent computer when that setting in this policy blocks an action. See [Chapter 20, “Endpoint Notifiers and Approval Requests,”](#) for details about choosing and defining notifiers.
- For more about banning software, see [“Approving and Banning Software”](#) on page 260. For more about creating custom rules for special treatment of files at certain paths, see [Chapter 14, “Custom Software Rules.”](#)

Table 21: Advanced Setting Behavior

Setting	Active	Off
Block unanalyzed scripts and executables	<p>Tracks executables (for example, .exe, .dll, and .com) and script files (for example, .bat, .vbs) that have not yet been analyzed and blocks them for systems in High, Medium, and Low Enforcement Levels, and in Local Approval mode.</p> <p>Scripts and executables are reported as unanalyzed if a user or process tries to execute them and App Control cannot finish its run-time checks of file state in the expected time. This usually happens when the root certificate for a file is out of date or otherwise not verifiable.</p>	Permits unanalyzed executables and script files to execute if no other settings prevent execution. Not recommended.
Block unapproved scripts	<p>Tracks script files (for example, .bat, .vbs) that have an unapproved status and blocks them according to Enforcement Level:</p> <ul style="list-style-type: none"> • High Enforcement Level blocks unapproved scripts. • Medium Enforcement Level blocks unapproved scripts but presents a dialog that identifies the file and gives users the option to run it. • Low Enforcement Level permits files to execute; records an event the first time the executable runs. <p>Table 49 in Chapter 13, “Script Rules,” shows the file types considered scripts by App Control.</p>	Permits script files not explicitly banned to execute if no other settings prevent execution.
Block unapproved executables	<p>Tracks executable files, for example, .exe, .dll, and .com, that have an unapproved status and blocks or permits them according to Enforcement Level:</p> <ul style="list-style-type: none"> • High Enforcement Level blocks all unapproved executables. • Medium Enforcement Level blocks unapproved executables but presents a dialog that identifies the file and gives users the option to run it. • Low Enforcement Level permits files to execute; records an event the first time the file runs. 	Permits unapproved files not explicitly banned to execute if no other settings prevent execution.
Block banned file names	Blocks execution of files banned by file name on computers in Control mode.	Cannot be disabled on the policy page, but individual bans can be made policy-specific.

Setting	Active	Off
Block banned file hashes	Blocks all banned hashes on computers in Control mode.	Disables the Banned Hashes setting and permits banned hashes to execute if no other settings prevent it.
Block executables run from a network drive	Blocks execution of files (including Approved files) run over the network on computers in Control mode. Platform Note: This setting is effective for Windows agents only.	Permits network executable files not unapproved or explicitly banned to execute if no other settings prevent it.
Block files with banned publishers or certificates	Blocks execution of files with banned publishers (or certificates) in Control mode.	Permits files with banned publishers/certificates to execute if no other settings prevent it.
Enforce memory rules	Apply all enabled memory access, control, and reporting rules. Platform Note: This setting is effective for Windows agents only.	Cannot be disabled on the policy page, but individual rules can be made policy-specific.
Enforce registry rules	Apply all enabled registry access and reporting rules to this policy. Platform Note: This setting is effective for Windows agents only.	Cannot be disabled on the policy page, but individual rules can be made policy-specific.
Enforce custom (file and path) rules	Apply all enabled custom rules (special treatment of files at defined paths) to this policy. You configure custom rules on the Custom Rules tab of the Software Rules page.	Cannot be disabled on the policy page, but individual rules can be made policy-specific.
Enforce tamper protection	Apply rules to prevent tampering with a App Control Agent.	Cannot be disabled for a policy. Use the Computer Details page to disable for a specific computer.
Terminate processes with banned images	When a file is banned, terminate currently running processes that match the file.	Permits a file banned while already running to continue running.
Locally approve unapproved files on transition from Visibility or Low Enforcement Level to Medium or High	When checked, causes certain Unapproved files to be locally approved when the policy Enforcement Level changes from Low (or None) to Medium or High. This only applies to files that first appeared on the computer as Unapproved when the computer was in a Low (or None) Enforcement Level policy. These files have Local State Details of "Unapproved". See "Locally Approving Files" on page 289 for details.	When not checked, Enforcement Level changes do not affect local file state in this policy.

Rules Affecting Policies

App Control rules that block or allow actions may be applied to selected policies or all policies. Although these rules are managed on other pages, a set of tabbed rule views is included on the Edit Policy page so that you can see which rules apply to a particular policy. These views are on the same tabbed panel as the Device Control Settings and Advanced settings. The rule views on the Edit Policy page are:

- File Rules (bans and approvals) – Described in [“File-Specific Rules: Approvals and Bans”](#) on page 301.
- Custom Rules – Described in [Chapter 14, “Custom Software Rules.”](#)
- Memory Rules – Described in [Chapter 16, “Memory Rules.”](#)
- Registry Rules – Described in [Chapter 15, “Registry Rules.”](#)
- Publisher Rules – Described in [“Approving or Banning by Publisher”](#) on page 281.
- Rapid Configs – Described in [Chapter 18, “Rapid Configs.”](#)

To view the rules that affect a policy:

- On the Edit Policy page, go to the tab panel at the bottom of the page and click on the tab for the type of rule you want to view.

Advanced | File Rules | **Custom Rules** | Memory Rules | Registry Rules | Publisher Rules | Rapid Configs | Computers | Device Control Settings

Show Filters ▾ | Show Columns ▾ | Export to CSV | Refresh Table | Show Disabled Rules | Show Rules That Apply To All Policies

Search: [] Automatically apply Showing 0 out of 0 item(s)

Rank ▲	Status	Platform	Rule Type	Name	Action	Operation	Path	Process	User or Group	Date Modified	Las
No data found.											

Showing 0 out of 0 item(s) Showing all data

Save & Exit Save Cancel Reset Policy

Notice that no rules are shown in the previous screen. The default view on the rules tabs shows only rules that are enabled, and that are specified by policy. Your site might always apply rules to *all policies*, in which case, you will see no rules by default. However, there are checkboxes on the tab pages that can change the rules viewed on the tab:

- **Show Disabled Rules** – On tabs with this box, if you check the box, both Enabled and Disabled rules are shown. If this box is unchecked (the default), Disabled rules are not shown.
- **Show Rules that Apply to All Policies** – On tabs with this box, if you check the box, all rules that apply to this policy are shown, even if they apply to *all policies*. If this box is unchecked (the default), only rules that specify this policy by name are shown.

Advanced | File Rules | **Custom Rules** | Memory Rules | Registry Rules | Publisher Rules | Rapid Configs | Computers | Device Control Settings

Show Filters ▾ | Show Columns ▾ | Export to CSV | Refresh Table | Show Disabled Rules | Show Rules That Apply To All Policies

Search: [] Automatically apply Showing 42 out of 42 item(s)

Rank ▲	Status	Platform	Rule Type	Name
3	Enabled	Windows	Expert	Examine powershell script contents
4	Enabled	Windows	Expert	Block powershell scripts that execute memory
6	Enabled	Windows	Expert	Do not treat these processes as .NET applications

When you have a tab that shows rules, you can click the View Details button next to a rule and (if you have permission to manage that type of rule) edit the rule. You cannot change the rank of rules through this view – go to the rules page itself to change rank.

Template Policy and Default Policy

Default Policy

App Control includes a built-in policy named Default Policy. This is the policy to which computers are assigned in the following situations:

- If you are using AD Mapping to assign policies, App Control is initially configured to assign a computer that does not match any other mapping rules to the Default Policy. You can, however, change the policy to which unmatched computers are assigned, and it is generally advisable to create a separate "AD Default" policy for this purpose. See [“Assigning Policy by Active Directory Mapping”](#) on page 122 for more information.
- When computers in a non-existent (deleted) policy report to the App Control Server, they are automatically moved into the Default Policy and subject to enforcement based on that policy’s settings. See [“Restoring Computers from the Default Policy”](#) on page 171 for information about how you might deal with this situation.

If you are licensed for Control features, you can set the Default Policy Enforcement Level to High (Block Unapproved) to make sure that if a computer is switched to the Default Policy, neither Banned nor Unapproved files are allowed to run. If you are less concerned about Unapproved files but still do not want to allow them to execute without user interaction, you can set the Enforcement Level to Medium. You also can edit any of the other settings for the Default Policy.

Note

Computers can be assigned to the Default Policy unexpectedly. Because of this, the initial policy setting for “Locally approve unapproved files on transition from Visibility or Low Enforcement Level to Medium or High” is *off* (un-checked). Otherwise an unexpected transition to the Default Policy could locally approve many files without you wanting that to happen. See [“Automatic Local Approval on Enforcement Level Change”](#) on page 290 for more details about this setting.

Template Policy

The built-in Template Policy is intended as a “template” for creating other policies. By default, the initial Device and Advanced settings of the first policy you create are based on the settings of this Template Policy, although you can base the initial settings on any other existing policy, including the Default Policy.

Note

Policies inherit only the *Device Settings* and *Advanced Settings* from their template policy. Settings on the top panel of the Add/Edit Policy page, including Enforcement Level, are not inherited. Device Settings and Advanced Settings appear on the Edit Policy page once you save a new policy.

You can edit the Template Policy to include the Device and Advanced settings you expect to want most of the time, simplifying policy creation. Once you create a policy, there is no ongoing linkage to its template policy, so you can change any setting in the new policy.

One important part of policy configuration is assigning notifiers to each setting in the policy that could block an action. Initially, each policy setting has a notifier assigned to it, and the message from each can differ depending on the setting that caused the block. If you want to change the messages from their defaults, it is best to alter the Template Policy *before* you create other policies. See “[Customizing and Creating Notifiers](#)” on page 548 for more information. You also can let a setting block actions without displaying any notifier.

A key difference between the Template Policy and the Default Policy is the Advanced Setting called “Locally approve unapproved files on transition from Visibility or Low Enforcement Level to Medium or High”. Activating this setting usually makes sense for a newly created policy, so it is activated by default (and not shown) for the Template Policy.

The Template Policy has the following special characteristics:


- it appears only on the Policies page and its own Edit page
- it cannot be assigned to any computer
- no AD mapping rules can be created that point to the Template Policy
- there is no agent installation package corresponding to the Template Policy
- like the Default Policy, the Template Policy cannot be deleted
- the setting “Locally approve unapproved files on transition from Visibility or Low Enforcement Level to Medium or High” is not shown but is automatically activated

Important

When you create a new policy, be sure to verify or, if needed, change the setting values you inherited from the existing policy you based it on.

Resetting a Policy to Template Policy Settings

The Edit Policy page for each policy includes a **Reset Policy** button at the bottom of the page.



When you press this button and choose **OK** on the confirmation dialog, the Device and Advanced settings are reset to the *current* settings of the Template Policy. Other policy settings remain unchanged.

Important

Once you click **OK** in the reset dialog box, the Device and Advanced policy settings are reset without requiring that you click **Save**. To prevent the reset, you must cancel in the *confirmation dialog box*. You cannot prevent the changes by clicking **Cancel** on the Edit Policy page.

Tamper-Protection Setting

A tamper-protection setting blocks attempts to write to the App Control application directory or change App Control Agent files on client computers. Tamper-protection cannot be disabled on a per-policy basis, although you can use the Advanced menu on the Computer Details page to disable it for an individual system – consult with Carbon Black Support before changing this setting.

Agents normally cannot be uninstalled unless they are in Agent Disabled mode. However, if you need to uninstall an agent without disabling it, Carbon Black Support can provide special commands to accomplish this task.

Note

You can specify your own directory-protection policies. See [Chapter 14, “Custom Software Rules.”](#)

For more information about removing App Control Agent from a computer, see [“Uninstalling App Control Agents”](#) on page 154.

Editing a Policy

You can edit the basic definitions of a policy, including its description, and Enforcement Level, in the upper panel of the Edit Policy page. Beginning with v8.0.0, you also can change the Policy Name.

For most Device and Advanced Settings, you can:

- turn them on or off
- place them in report-only state, in which they report what they would have done if they had been activated
- choose a different (or no) *notifier*, which is the dialog box that is displayed when an action is blocked as a result of an active policy setting; this is covered in [Chapter 20, “Endpoint Notifiers and Approval Requests.”](#)

Certain settings have fewer choices or choices other than those on this list.

Notes

Although you can deactivate policy settings, you cannot create or delete them. The setting name (e.g., *Block unapproved scripts*), which is standard for all policies, cannot be changed.

To edit a policy:

1. On the console menu, choose **Rules > Policies**. The Policies page appears:

Policies

Policies Mappings

Group By: (none) Ascending

Show Filters Show Columns Export to CSV Refresh Page

Action Add Policy

	Policy	Connected Enforcement	Disconnected Enforcement	Total
<input type="checkbox"/>	APC-99986-ucdjfp	High (Block Unapproved)	High (Block Unapproved)	1
<input type="checkbox"/>	Default Policy	None (Visibility)	None (Visibility)	0
<input type="checkbox"/>	Help Desk	None (Visibility)	None (Visibility)	0
<input type="checkbox"/>	Local Approval Policy	Local Approval	Local Approval	0
<input type="checkbox"/>	Research Group	None (Visibility)	None (Visibility)	0
<input type="checkbox"/>	Template Policy	None (Visibility)	None (Visibility)	0

6 items Page 1/1

2. On the Policies page, click the View Details button next to the name of the policy you want to edit. The Edit Policy page appears:

Edit Policy Help Desk

Policy Name: Help Desk

Description:

Mode: Visibility Control Disabled

Automatic Policy Assignment For New Computers:

Set Automatic Policy For Existing Computers: There are currently no computers in this policy.

Options: Allow Upgrades Track File Changes Load Agent in Safe Mode Suppress Logo In Notifier

Total Computers: 0

Connected Computers: 0

Advanced File Rules Custom Rules Memory Rules Registry Rules Publisher Rules Rapid Configs Computers Device Control Settings

Name	Status	Notifiers
Block writes to unapproved removable devices	Off	<default>: Block writes to unapproved removab
Block writes to banned removable devices	Active	<default>: Block writes to banned removable de
Report reads from unapproved removable devices	Off	<none>
Report reads from banned removable devices	Off	<none>
Block executions from unapproved removable devices	Off	<default>: Block executions from unapproved r
Block executions from banned removable devices	Active	<default>: Block executions from banned remo

Save & Exit Save Cancel Reset Policy

3. Edit any of the details in the main panel by checking or un-checking the appropriate box, entering text, choosing a different mode and/or choosing a different Enforcement Level. Visible parameters may vary depending upon other policy settings and configuration choices. See [Table 19, "Policy Definitions: Main Panel,"](#) on page 183 for detail on these settings.

Note: If you change the Policy Name, that name will be reflected immediately in the console, but the name of the agent installer (the *polycyname.msi* file) requires approximately one minute to update. Keep this in mind if you intend to download agents immediately after a policy name change.

- From the Edit Policy page, click the **Device Control Settings** tab to see the Device Control settings for this policy.
- In the Device Control Settings table, use the dropdown menu to select one of the following states for any setting you want to change: **Off**, **Active**, and **Report Only** (*Active* is not a choice for the Read settings). See [Table 46, “Device Control Setting Behavior,”](#) on page 363 for information about these settings.
Platform Note: Visibility and control features for devices are not available for Linux computers.
- From the Edit Policy page, click the **Advanced** tab to see the Advanced settings for this policy.

Advanced			
File Rules			
Custom Rules			
Memory Rules			
Registry Rules			
Publisher Rules			
Rapid Configs			
Computers			
Device Control Settings			
Name	Status	Notifiers	
Block unanalyzed scripts and executables	Active	<default>: Block unanalyzed scripts and executables	Add Edit
Block unapproved scripts	Active	<default>: Block unapproved scripts	Add Edit
Block unapproved executables	Active	<default>: Block unapproved executables	Add Edit
Block banned file names	Active	<default>: Block banned file names	Add Edit
Block banned file hashes	Active	<default>: Block banned file hashes	Add Edit
Block executables run from a network drive	Off	<default>: Block executables run from a network drive	Add Edit
Block files with banned publishers or certificates	Active	<default>: Block files with banned publishers or certificates	Add Edit
Enforce memory rules	Active	<default>: Enforce memory rules	Add Edit
Enforce registry rules	Active	<default>: Enforce registry rules	Add Edit
Enforce custom (file and path) rules	Active	<default>: Enforce custom (file and path) rules	Add Edit
Enforce tamper protection	Active	<default>: Enforce tamper protection	Add Edit
Terminate processes with banned images	Report Only	<default>: Terminate processes with banned images	Add Edit
<input checked="" type="checkbox"/> Locally approve unapproved files on transition from Visibility or Low Enforcement Level to Medium or High			
<input type="button" value="Save & Exit"/> <input type="button" value="Save"/> <input type="button" value="Cancel"/> <input type="button" value="Reset Policy"/>			

- In the Advanced Settings table, use the menu to select one of the following states for settings you want to change: **Active** (on), **Report Only** (on, but not enforced), or **Off**. See [Table 21, “Advanced Setting Behavior,”](#) on page 189 for more on these settings.
Note: Some Advanced settings cannot be changed. Fixed settings show their value in a grayed-out menu box.
- To change the setting for *Locally approve unapproved files on transition from Visibility or Low Enforcement Level to Medium or High*, check or un-check the box.
- To customize the notifier shown by a Device or Advanced setting when it blocks actions on an agent computer, you can choose a different notifier from the Notifiers menu next to the setting, **Edit** the notifier (which affects all places in which this notifier is used), or **Add** and define a new notifier. See [“Customizing and Creating Notifiers”](#) on page 548 for more information.
- When you have finished changing policy settings, click **Save**. Your changes are saved and the Policies table is re-displayed.

Related Views in Policy Details

The Edit Policy page has a Related Views menu with links that provide information about computers in this the policy and the files on those computers:

- All files on computers in this policy* opens a Find Files page with all instances of tracked files on the computers assigned to the policy.
- Unapproved files on computers in this policy* opens a Find Files page with all file instances with a *Local State* of Unapproved on the computers assigned to this policy.

This helps show how the policy settings affect the files actually on these computers. You can add another filter to the results to show only files with *Local State Details* of Unapproved – these would be approved by an Enforcement Level change from Low to either Medium or High if the automatic approval box is checked for this policy.

- *Computers manually assigned to this policy* opens a filtered view of the Computers page, showing computers that have been manually assigned to the policy (i.e., were not assigned by AD mapping).

Enforcement Levels

Enforcement Level is the protection level applied to computers running the App Control Agent, specified on a per-policy basis. Enforcement Levels, which vary in restrictiveness, affect how file actions are controlled for policy settings. File-blocking and other control functions in App Control depend on both the Enforcement Level and on more specific policy settings in effect, including policy-specific bans.

In Control mode, you choose High (Block Unapproved), Low (Monitor Unapproved), or Medium (Prompt Unapproved) Enforcement Level from a menu. The other modes, None (Visibility) and None (Disabled), automatically designate the Enforcement Level as None.

Table 22: Enforcement Levels

Enforcement Level	Use when:
High (Block Unapproved)	<p>For the highest protection level, and when it is practical to pre-approve the applications you need and want to run on computers in the policy, use High enforcement.</p> <p>High enforcement permits only explicitly approved files to run. Computers on which the application configuration seldom changes – servers or single-purpose systems, for example – are good candidates for High enforcement. For computers with more dynamic application configurations, High enforcement might be usable <i>if</i> you also pre-approve files via trusted directories, trusted users, approved publishers, enabled updaters, or reputation approvals.</p> <p>Except for files already identified and banned on the App Control Server, all files that exist on computers before you install the App Control Agent are locally approved and permitted to run on that computer under High enforcement.</p> <p>High enforcement is available to policies in Control mode.</p>

Enforcement Level	Use when:
Medium (Prompt Unapproved)	<p>To operate in a condition that prevents unchallenged execution of unapproved files but does not completely block them, use Medium enforcement.</p> <p>Medium enforcement blocks all Unapproved files from executing but displays a dialog on client computers that lets the user decide whether to run the file. If the user allows the file to run, it is locally approved on that computer and always permitted to run. If an Unapproved file is run remotely from a network share or removable device and allowed by the user, it is temporarily approved to run (the approval remains for 14 days).</p> <p>Platform Note: Some removable or network drives are not recognized by App Control, especially on non-Windows systems. Files run from these drives are treated like local files. Explicitly banned files cannot run under Medium enforcement. Medium enforcement is available to policies in Control mode.</p>
Low (Monitor Unapproved)	<p>When you are not concerned about unknown files and only need to block files that you have specifically banned, use Low enforcement.</p> <p>Low enforcement blocks banned files while allowing users to install software that are Approved or Unapproved (neither banned nor approved). Although Unapproved files are permitted to execute, you can monitor them and respond with emergency lockdown if necessary.</p> <p>Low enforcement is available to policies in Control mode.</p>
None (Visibility)	<p>To track file activity without blocking it, set the Enforcement Level to None (Visibility).</p> <p>Visibility mode tracks executable file activity on your computers through App Control's reporting and asset management features (drift reports, event reports, file inventory, etc.), but enforces no rules. It can be a first step on the way to implementing a more controlled environment.</p> <p>Click Visibility in the Mode line to choose this level.</p>

Enforcement Level	Use when:
None (Disabled)	<p>To stop all enforcement and tracking activities, choose None (Disabled) mode. You might do this if:</p> <ul style="list-style-type: none"> You are instructed to disable an agent by Carbon Black Support staff so that you can debug a system fault. You plan to remove the App Control Agent from a computer; a computer <i>must be</i> in None (Disabled) mode before the agent is deleted and the computer is removed from the App Control Server. <p>If you disable the agent for a computer, that computer's file database is deleted from the agent computer but remains on the server for one day. Computers in Agent Disabled mode re-initialize their files as soon as you move them to a policy at another Enforcement Level.</p> <p>Note: An agent in None (Disabled) mode continues to monitor (but not report to the server) certain operations to avoid gaps in file and process information if the agent is later brought back into an active mode. This normally requires a very minimal amount of resources on the agent computer, although if an extremely large number of writes are performed, the impact may be noticeable.</p> <p>Click Disabled in the Mode line to choose this level.</p>

How Enforcement Levels Affect Policy Setting Enforcement

Enforcement Levels interact with policy settings and other rules to control the conditions under which different types of files actions are allowed. [Table 23](#) shows how file activity is affected for different combinations of Enforcement Level and:

- Advanced Policy Settings and network-wide file bans that are *Active*
- Device Control Settings that are set to *Active*

Table 23: Effects of Active Policy Settings by Enforcement Level

Active Policy Settings	Enforcement Levels				
	None (Disabled)	None (Visibility)	Low (Monitor Approved)	Medium (Prompt Unapproved)	High (Block Unapproved)
Block unanalyzed scripts & executables	off	allow	block	block	block
Block unapproved scripts	off	allow	allow	prompt	block
Block unapproved executables	off	allow	allow	prompt	block
Block banned file names (cannot be disabled)	off	allow & report	block	block	block
Block banned file hashes	off	allow & report	block	block	block
Enforce memory rules (cannot be disabled)**	off	non-blocking action & report	block (if specified)	block (if specified)	block (if specified)
Enforce registry rules (cannot be disabled)**	off	non-blocking action & report	block (if specified)	block (if specified)	block (if specified)
Enforce custom (file and path) rules (cannot be disabled)**	off	non-blocking action & report	block (if specified)	block (if specified)	block (if specified)
Enforce tamper protection (cannot be disabled)	basic	full	full	full	full
Terminate processes with banned images	off	continue & report	terminate	terminate	terminate
Block executables run from a network drive *	off	allow & report	block	block	block
Block writes to unapproved removable devices *	off	allow & report	block	block	block
Block files with banned publishers or certificates	off	allow & report	block	block	block
Block writes to unapproved removable devices *	off	allow & report	block	block	block
Block writes to banned removable devices *	off	allow & report	block	block	block
Report reads from unapproved removable devices*	off	allow & report	allow & report	allow & report	allow & report
Report reads from banned removable devices*	off	allow & report	allow & report	allow & report	allow & report
Block execution from unapproved removable devices *	off	allow & report	block	block	block
Block execution from banned removable devices *	off	allow & report	block	block	block

* Device and Network Drive rules apply to Window computers only.
** The possible actions for memory, registry and custom rules include many non-blocking options.

Notes

- When an attempt to execute an Unapproved file generates a dialog in Medium Enforcement, either choice (block or allow) is recorded as an event. Also, with Enforcement Level set to Low, execution of an Unapproved file generates an event.
- The Related Views menu on the Edit Policy page includes a link called *Unapproved files on computers in this policy*. Since Enforcement Level affects how unapproved files are handled, this link can help you decide how to set Enforcement Level, or whether to leave a given computer in its current policy.

Special Enforcement Level for Local Approval

App Control sets a special Enforcement Level for computers in local approval. This Enforcement Level is reserved for system use, and cannot be chosen directly. It enables local approval of software, especially for computers otherwise under High Enforcement

Changing Policy Enforcement Levels

If you want to change the level of rule enforcement for a group of computers, you might move them to a different policy. Moving computers is described in [“Moving Computers to Another Policy”](#) on page 170.

Another alternative is to raise or lower the Enforcement Level applied to the *current* policy, using one of the following methods:

- If you are already in Control mode and want to stay there, you can switch between control Enforcement Levels by editing a policy's Connected Enforcement Level and Disconnected Enforcement Level menus. For example, to increase protection you can switch policies under Low (Monitor Unapproved) Enforcement Level or Medium (Prompt Unapproved) Enforcement Level to High (Block Unapproved) Enforcement Level.
- If you are already in Control mode and want to eliminate control, you can switch to Visibility mode, which changes the Enforcement Level to None (Visibility).
- If you are in Visibility mode, you can switch to Control mode and choose a new Enforcement Level from the menus.

Important

Disabling and re-enabling a large number of agents in one operation is not recommended. Switching *to* Agent Disabled mode eliminates enforcement, reporting, and tracking provided by the App Control Agent. Switching back *from* Agent Disabled can have significant performance impact, based upon the number of agents in a policy. Each agent switching *out of* Agent Disabled mode reinitializes, going through the same process as a newly installed agent.

To change Enforcement Level for a policy in Control mode:

1. On the console menu, choose **Rules > Policies**. The Policies page appears:

Policy	Connected Enforcement	Disconnected Enforcement	Total
APC-99986-ucdjfp	High (Block Unapproved)	High (Block Unapproved)	1
Default Policy	None (Visibility)	None (Visibility)	0
Help Desk	None (Visibility)	None (Visibility)	0
Local Approval Policy	Local Approval	Local Approval	0
Research Group	None (Visibility)	None (Visibility)	0
Template Policy	None (Visibility)	None (Visibility)	0

2. On the Policies page, click the View Details button next to the policy name you want to edit. The Edit Policy page appears:

Policy Name: Help Desk

Description:

Mode: Visibility Control Disabled

Automatic Policy Assignment For New Computers:

Set Automatic Policy For Existing Computers: There are currently no computers in this policy.

Options: Allow Upgrades Track File Changes
 Load Agent in Safe Mode Suppress Logo In Notifier

Total Computers: 0
Connected Computers: 0

Name	Status	Notifiers
Block writes to unapproved removable devices	Off	<default>: Block writes to unapproved removab
Block writes to banned removable devices	Active	<default>: Block writes to banned removable de
Report reads from unapproved removable devices	Off	<none>
Report reads from banned removable devices	Off	<none>
Block executions from unapproved removable devices	Off	<default>: Block executions from unapproved n
Block executions from banned removable devices	Active	<default>: Block executions from banned remo

3. If you want to switch modes, click the button next to the mode you want.
4. To change Enforcement Level within Control mode, select a Connected Enforcement Level from the dropdown menu:

High (Block Unapproved)
 Medium (Prompt Unapproved)
 Low (Monitor Unapproved)

5. If you chose High or Medium for *Connected Enforcement Level*, you can choose a different *Disconnected Enforcement Level* from its dropdown menu.
6. Make any other needed changes to the policy. See “[Policy Settings](#)” on page 186 for details of policy settings.

7. To save the changes, click the **Save** button at the bottom of the page.

Locking Down all Computers

The App Control Console Home page includes an emergency Lockdown button that changes the Enforcement Level of all agent-managed computers to High. During an emergency lockdown, the following is true for active agents whose policies do not have any enforcement settings disabled:

- Banned files are blocked.
- All Unapproved files that appear *after* the emergency lockdown are blocked.
- All existing Unapproved files that *remain* Unapproved are blocked.
- Certain files become locally approved, as described below, and can be executed.
- Computers that were offline when emergency lockdown was initiated are locked down upon reconnection to the App Control Server if the lockdown remains in effect.
- Lockdown affects all active agents, including those in Visibility Only mode. It does not affect computers whose agents are disabled.

In some cases, locking down a computer causes some Unapproved files to become locally approved. In the Advanced Settings panel of the Edit Policy page, there is a checkbox labeled “[Locally approve unapproved files on transition from Visibility or Low Enforcement Level to Medium or High](#)”. This affects computers whose Enforcement Levels are Low or None when they are moved to Enforcement Levels of High or Medium:

- If the box is checked, existing Unapproved files that first appeared on a computer when it was in Low (or None) Enforcement Level are locally approved upon lockdown.
- If the box is not checked, Unapproved files on computers in that policy remain Unapproved after lockdown and are not allowed to run.

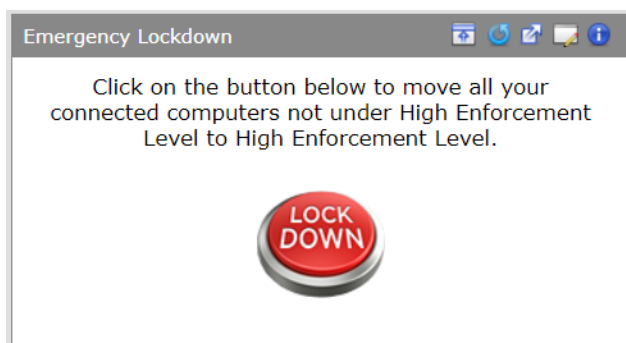
Console users with the default ReadOnly privileges do not have access to Emergency Lockdown. A login account role must have *Manage Computers* privileges for its members to perform an emergency lockdown.

Notes


Emergency Lockdown changes only the *Enforcement Level* of computers. In policies with Advanced Settings of *Off* or *Report Only*, computers might not block certain threats even when in lockdown.

To lock down all computers:

1. From the console menu, choose **Home**. The Home page appears. The default location of the Emergency Lockdown portlet is the bottom right portlet on the page, although you or another administrator may have moved or removed it:




2. In the Emergency Lockdown portlet, click the **Lock Down** button. The Lockdown confirmation page appears:



3. In the confirmation dialog, click **OK** to lock down all computers. All agents except those in Disabled mode are locked down. The Home page appears and the **Lock down computers** button toggles to **Restore computers**:



4. After you resolve the issue that lead to the Lockdown, click the **Restore computers** button to restore all computers to their former Enforcement Level. The Restore confirmation page appears:



5. In the confirmation dialog, click **OK** to restore all computers.

Deleting Policies

You can delete policies when you no longer need them. However, policies cannot be deleted if any computer is assigned to the policy. If a policy you want to delete has associated computers, either uninstall the App Control Agent on those computers or, to keep the computers protected by App Control, move the computers to another policy. See [“Uninstalling App Control Agents”](#) on page 154 and [“Moving Computers to Another Policy”](#) on page 170. When you delete a policy, its associated agent installer is deleted from the App Control Server.

The following built-in policies cannot be deleted:

- Default Policy
- Local Approval Policy
- Template Policy

To delete a policy:

1. On the console menu, choose **Rules > Policies**. The Policies page appears:

<input type="checkbox"/>	Policy	Connected Enforcement	Disconnected Enforcement	Total	Connected
<input type="checkbox"/>	APC-99986-ucdfp	High (Block Unapproved)	High (Block Unapproved)	1	1
<input type="checkbox"/>	Default Policy	None (Visibility)	None (Visibility)	0	0
<input type="checkbox"/>	Help Desk	None (Visibility)	None (Visibility)	0	0
<input type="checkbox"/>	Local Approval Policy	Local Approval	Local Approval	0	0
<input type="checkbox"/>	Research Group	None (Visibility)	None (Visibility)	0	0
<input type="checkbox"/>	Template Policy	None (Visibility)	None (Visibility)	0	0

2. On the Policies page, click the Delete (trashcan) button next to the name of the policy you want to delete. A confirmation dialog appears.

Confirmation

Confirm Policy Deletion

Policy Name:	Help Desk
Description:	
Connected Enforcement:	None (Visibility)
Disconnected Enforcement:	None (Visibility)

Yes No

3. Click **Yes**. You will return to the Policies page.

Note

If a policy contains computers, clicking **Yes** in the confirmation dialog displays a deletion failure message on the Policies page. You must move these computers to another policy or delete them (on the Computers Page) before deleting the policy.

Chapter 6

Managing Virtual Machines

This chapter explains how App Control can efficiently manage virtual machines, called *clones* in the console, and the *template* computers on which they are based. To manage virtual machines, you also will need to be familiar with [Chapter 4, “Managing Computers.”](#)

Sections

Topic	Page
Overview	208
Creating a Template Computer	209
Deploying Clones	214
Making Changes to a Template	216
Configuring Clone Inventory	217
Deleting a Template	217
Deleting Clones	219
Converting a Template to a Regular Computer	221

Overview

When the App Control is installed on a virtual machine, App Control can manage the virtual machine just as it manages physically distinct computers. However, you can *improve* the way virtual machines are managed if some special steps are taken.

When you provision a computer on a virtualized software platform that includes the App Control Agent and then convert that computer to a template using the App Control Console, much of the file inventory processing on clones based on this template can be optimized. The App Control Server can automatically initialize a clone's inventory based on its template, or optionally, you can choose to have the server track only file changes that happen after a clone is created.

These options reduce the network traffic and server load associated with cloned computers, potentially allowing much larger number of virtual machines to be managed by a App Control Server. In addition, the server maintains an association between the template and its clones so that you can easily discover which computers are based on a particular template and manage them accordingly.

Notes

- While this chapter primarily describes how you manage virtual machines as clones, the procedures are applicable to re-imaging of physical computers (such as "ghosting") in which the clones are actually physical machines with a common disk image from a template.
- If you worked with support representatives to implement a custom solution to manage templates and clones in pre-7.0 releases, that solution will still work but is not integrated with the new, standard template management features.

The following key terms are used throughout the chapter and in the App Control Agent to describe the components of virtual and ghosted machine management:

- **Template Computer** - A computer that is pre-installed with required software, including the App Control Agent, and will be used to clone one or more computers through VMware or some other mechanism (e.g., "ghosting" of the hard drives of multiple computers from a common image). Before a computer can be used as a App Control template computer, it must be taken offline.
- **Cloned Computer** - A computer that originated as a clone of a template computer. It will register with the App Control Server as a new computer, but it will also remain identified as a clone of a specific parent template.
- **Parent Template** - Each cloned computer points to its parent Template Computer. This mapping persists until either the clone or the template is deleted.

The login account used to log in to the console must have Manage Computers permission to be able to manage templates and clones.

Creating a Template Computer

App Control does not provide the software (such as VMware View) for creating virtual machines or managing cloned disk images for physical machines, nor does this chapter provide instructions for using those systems. A prerequisite of using the features described here is that you have, and know how to use, a product that creates clones from a primary disk image. The App Control Server can manage the clones produced by those systems, but is not integrated with the systems themselves.

App Control requires the following for a template computer:

- it must have a App Control (or Bit9 or Parity) agent 7.0.0 or greater installed
- it must *not* be the home of any Trusted Directories used by the App Control Server
- it must be fully initialized
- it can be either a physical computer or a virtual machine
- it must be shut down and show as *offline* in the console before becoming a App Control template computer, and should remain offline afterward

To create a template computer:

1. On the computer you plan to use as a template, install the platform, application, and other files you want in the template image.
2. Install (or upgrade to) App Control Agent 7.2.1 or greater on the computer.
3. After App Control Agent installation, make sure the computer is connected to the App Control Server and let it fully initialize. You can monitor initialization progress by choosing **Assets > Computers** on the console menu and clicking on the View Details button next to the name of the computer. Initialization progress is on the **Connection History** tab of the Computer Details page.


Carbon Black App Control Agent	Connection History	Policy Override	System Details	AD Details	Carbon Black EDR
First Registered:	Jun 23 2020 05:33:54 PM				
Last Polled:	Jun 29 2020 01:29:44 PM				
Last Register Date:	Jun 29 2020 01:29:37 PM				
Initialization:	Complete				
Synchronization:	100%				
Server Backlog:	0 files				
Last Logged In User(s):	WORKGROUP\APC-99986\$ APC-99986\bit9qa				

4. When initialization shows as *Complete*, also make sure that Synchronization is at 100%. Files added to the template computer after the App Control Agent is installed will be included in synchronization, not initialization. Wait for both initialization and synchronization are completed before proceeding to the next steps.

Carbon Black App Control Agent	Connection History	Policy Override	System Details	AD Details	Carbon Black EDR
First Registered:	Jun 23 2020 05:33:54 PM				
Last Polled:	Jun 29 2020 01:29:44 PM				
Last Register Date:	Jun 29 2020 01:29:37 PM				
Initialization:	Complete				
Synchronization:	100%				
Server Backlog:	0 files				
Last Logged In User(s):	WORKGROUP\APC-99986\$ APC-99986\bit9qa				

5. Shut down the computer.
Note: Prior to v7.2.3, if you used sysprep to prepare a template, you needed to disable tamper protection on the computer before shutdown. This is no longer necessary.
6. Go to the Computer Details page for the computer, and click **Convert to Template** on the Advanced menu. The Computer Details page changes to a Template Details page.
7. By default, the Template Name is the name of the computer from which the template was created, but you can change it, add a description, and change the cleanup and inventory parameters on the Template Settings tab (see [“Deleting Clones”](#) and [“Configuring Clone Inventory”](#) for details).
8. When you are satisfied with the configuration on the Template Details page, click **Save**. The computer now appears in the Computers table as a template.
Note: Except for specific tasks described later in this chapter, do not bring a computer back online after it is converted to a template. If you bring a template computer back online, it will appear as a clone of itself.
9. Create clones from the computer using your virtualization software. They will appear as new computers in the console.

Viewing Templates in the Computers Table

On the Computers page, you can view a table of computers and information about them, including their policies, Enforcement Levels, and whether they are currently connected to the server. By default, the full Computers table includes a Connected column, which indicates template computers by a white circle with a gray border . You also can add a Template column to the Computers table using the Show Columns button. This column will show *Yes* for templates and *No* for computers that are not templates.

If all you want to see is template computers, you can use the Template Computers Saved View.

To view the template computers in the table of computers:

1. In the console menu, choose **Assets > Computers**. The Computers page appears.
2. Choose **Template Computers** on the Saved Views menu to see computers that are templates for cloned computers.

Computers

Computers connected: 1 Total computers: 1 Current CL version: 1186 CL version for upgrade: 1142

Saved Views: Template Computers Add Group By: (none) Ascending Days Disconnected: (none)

Hide Filters Show Columns Export to CSV Refresh Page

Filters

Add filter

Template Yes No

Apply Cancel Reset

Action Search: Go Clear

Computer Name	Template Date	Upgrade Status	Connected Enforcement	Clone Inventory
MYCORPWIN8-64-IT	Jul 28 2016 12:41:40 PM	Upgrade requested	High (Block Unapproved)	All files
MYCORPWIN7-ENG	Feb 3 2017 09:04:44 PM	Up to date	High (Block Unapproved)	All files
MYCORPWIN7-32-GEN	Jul 19 2016 09:03:33 PM	Not requested	High (Block Unapproved)	All files
MYCORPWIN10-64-ENG	Oct 7 2016 10:57:53 PM	Up to date	High (Block Unapproved)	All files

3. The Saved View uses the filter checkbox *Template/Yes*. Instead of (or in addition to) the Saved View, you can click on **Show Filter** to further customize the view you have of the Computers table.

The default Template Computers view includes a Clone Inventory column that shows whether the file inventory of clones from this template includes all files on the clone computers or just files that were added or modified after creation of the clone. Note that the file inventory might also be affected by exclusion of tracking for Microsoft-signed support files. See [“Excluding Tracking of Microsoft Support Files”](#) on page 230 for details.

Viewing and Editing Template Details

As with non-template computers, there are several ways to locate a template computer and display its details. You can use the Find Computer portlet on the Home Page to locate the template computer and then drill down to its details. The following procedure describes how you can locate and get details for a template computer through the Computers page.

To view the Template Details page for one computer:

1. In the console menu bar, choose **Assets > Computers**. The Computers Page appears.
2. In the Computers table, locate the template computer for which you want complete details (for example, searching by name, using the *Template Computers* Saved View, or using the Computer filters panel).

3. In the table, click either the name of the template computer or the View Details button next to its name. The Template Details page appears.

Template Details

General

Template Name:

Health Check: Passed

Platform: Windows

Description:

Computer Tag:

Policy

Policy: [APC-99986-ucdjfp](#)

Policy Mode: Control

Connected Enforcement: High (Block Unapproved)

Disconnected Enforcement: High (Block Unapproved)

Carbon Black App Control Agent Connection History System Details AD Details Template Settings Carbon Black EDR

CL Version: 1186

Debug Level: None (default)

Carbon Black App Control Agent Version: 8.1.8.306

Enabled Trusted Directories: 0

Tamper Protect: Enabled

Much of the information is the same as for the Computer Details page ([Table 16](#) on page [161](#)), but there are important differences, as shown in [Table 24](#).

Table 24: Differences between Template Details and Computer Details

Field/Menu/Tab	Description in Template Details Page
Template Name	Replaces Computer Name on the details page. By default this is the name of the computer from which the template was made. Must be unique.
IP Address	Not present on the Template Details page (has no meaning for a computer that is required to be offline).
Connection Status	Not present on the Template Details page (has no meaning for a computer that is required to be offline).
Health Check	On the Template Details page, this is the last Health Check done before the computer became a template.
Policy Override tab	Not present on the Template Details page.
Template Settings tab	<p>Details about the template. It includes the following:</p> <ul style="list-style-type: none"> • Date Created – When the template was created in the App Control Agent. • Original Computer Name – The name of the computer when it was converted to a template. • Original IP Address – The IP address of the computer when it was converted to template. • Clone Count – The current number of clones from this template. • Clone Inventory – Whether the file inventory for each clone should include all files, including the files cloned from the template computer, or only new and modified files. See “Configuring Clone Inventory” on page 217. • Clone Cleanup – When clones for this template should be deleted when offline. See “Deleting Clones” on page 219.
Related Views menu	<p>Includes:</p> <ul style="list-style-type: none"> • Show All Cloned Computers, which shows all clones for this template that have been connected to the App Control Server and not yet deleted. • Health Check Events shows the table of Health Check events for this computer before it became a template. • Files on this Computer takes you to a Find Files page with a table of all tracked file instances on the template computer.
Actions menu	<p>The single item on this menu changes depending upon conditions:</p> <p>Delete Offline Clones - Appears if the template has clones listed in the console. Deletes all clones of this template that are currently offline.</p> <p>Convert to Computer - Appears if the template has no clones managed by the App Control Server. In this case, you can convert the template computer back to a regular computer and reconnect it to the server, if needed. This is primarily intended to allow you to undo an unintended template conversion.</p> <p>The menu does not appear if neither condition applies.</p>
Advanced menu	Not present on the Template Details page.

Deploying Clones

Once you have registered a computer as a template, any clones of that template are automatically recognized by the App Control Server. Because they are clones, initialization will occur much faster than it would for non-clone computers.

Any manual or automatic methods of reverting the clones to their snapshot images will result in new clones being added to the console Computers list, still associated with the same template. The “old” clones go offline as far as the App Control Server is concerned, and they can be cleaned up by whatever method you choose (see “Deleting Clones” on page 219).

Viewing Clones in the Computers Table

On the Computers page, you can view a table of computers and information about them, including their policies, Enforcement Levels, and whether they are currently connected to the server. You also can add a Parent Template column to the Computers table using the Show Columns button. Any computer that has a value in this column is a clone. Computers that are not clones show nothing in this column.

If you only want to see clones, you can use the Cloned Computers Saved View on the Computers page to see all cloned computers known to the App Control Server. By default, this view is grouped by Parent Template, so you know what the clones are based upon.

Computers

Computers connected: 1 Total computers: 2 Current CL version: 1186 CL version for upgrade: 1142

Saved Views: (The Current View Has Unsaved Changes - Discard) Cloned Computers Add

Group By: Parent Template Ascending Days Disconnected: (none)

Hide Filters Show Columns Export to CSV Refresh Page

Filters

Add filter

Parent Template is not empty +

Apply Cancel Reset

<input type="checkbox"/>	Computer Name ^	Connected	Parent Template	Policy Status	IP Address	Policy	Inventory
Parent Template: MYCORPWIN10-64-ENG							3 items
<input type="checkbox"/>	MYCORPENG-10-VM1	●	MYCORPWIN10-64-ENG	Up to date	10.29.40.139	Eng-Win-M	All files
<input type="checkbox"/>	MYCORPENG-10-VM2	●	MYCORPWIN10-64-ENG	Up to date	10.29.40.25	Eng-Win-M	All files
<input type="checkbox"/>	MYCORPENG-10-VM3	●	MYCORPWIN10-64-ENG	Up to date	10.29.40.138	Eng-Win-M	All files
Parent Template: MYCORPWIN7-32-GEN							55 items
Parent Template: MYCORPWIN8i64-IT							1 item

The Saved View for Cloned Computers uses the filter *Parent Template is not empty*. Instead of (or in addition to) the Saved View, you can click on Show Filter to further customize the view you have of the cloned computers.

The default Cloned Computers view includes an Inventory column that shows whether the file inventory of this clone includes all files (including those in the template image) or just files that were added or modified after creation of the clone.

Finding the Clones for a Template

There are several ways to identify the clones created from a template:

- On the Computers page, you can choose the Cloned Computers Saved View. This displays clones grouped by their Parent Template.
- On the Template Details page, you can choose Show All Cloned Computers on the Related Views menu.
- On the Computers page, you can use the Parent Template filter to locate all clones from a particular template. This is also useful if you are not sure of the exact template name, since you can enter partial strings to match the name.

Finding the Template for a Clone

You can find the template for a clone computer in the following ways:

- On the Computers page, you can choose the Cloned Computers Saved View. This displays clones grouped by their Parent Template.
- On the Computer Details page, the information listed for a clone is almost the same as the information listed for any other computer, but in addition to the standard information, there is a Template Computer field if the computer is a clone.

Computer Details

General

Computer Name:	WORKGROUP\APC-99986
Template Computer:	Default Helpdesk Template
Inventory:	All files
IP Address:	10.210.65.144
Connection Status:	Connected
Health Check:	Passed
Platform:	Windows
Description:	<input type="text"/>
Computer Tag:	<input type="text"/>

Server Backlog for Clones

The Connection History tab on the Computer Details page includes a field called Server Backlog. This is the number of files that have been received from the computer but not yet fully processed on the server. Files in backlog appear in the File Catalog but not in the Files on Computers tab or Find Files page.

This is particularly significant for clones that are configured to inventory all files, including those in the template image. When a clone is discovered by the App Control Server, if it is configured to inventory all files, the file inventory from its parent template is copied into that computer's backlog. In this case, the Server Backlog field will show a large increase in the number of files. The file inventory of the cloned machine will not be available until this backlog is cleared.

Making Changes to a Template

You might need to modify an existing template for all users, for example, to install new operating system updates. Another possibility is that you might need to keep the original template image but create a new template that is slightly modified to be appropriate for a different purpose or a different group of users.

To modify an existing template, you will have to bring the template computer back online. When it is online, it will be treated as a new clone computer of the original template. You can install updates and make any other needed modifications on the computer while it is considered a clone. When you are finished, you can convert the “clone” into a template. New templates made from an existing template computer automatically inherit the clone cleanup parameters from the original template.

Clones of original template are not automatically deleted – they are still valid as long as they remain online. You can use your virtualization/imaging infrastructure to manage these clones as you see fit.

What you do with the old template depends upon why you updated it and whether there are still online clones associated with it. If the new template was truly an update and the old version is obsolete, you could delete the old template, preferably after any of its clones are offline. See [“Deleting a Template”](#) for more information.

If the new template was a variation, and not necessarily a replacement of the old template, you might want to keep both templates available.

To update a template computer:

1. Bring the template computer back online. It will appear in the console as a clone of its original template.
2. On this “clone” computer, make whatever file additions, deletions and modifications you want for the updated template.
3. Using your virtualization or imaging software, update the image for this computer or create a new one.
4. Wait for the file inventory of the clone to be fully synchronized. You can monitor synchronization progress by choosing **Assets > Computers** on the console menu and clicking on the View Details button next to the name of the computer. Synchronization progress is on the **Connection History** tab of the Computer Details page.

Carbon Black App Control Agent	Connection History	Policy Override	System Details	AD Details	Carbon Black EDR
First Registered:	Jun 23 2020 05:33:54 PM				
Last Polled:	Jun 29 2020 01:29:44 PM				
Last Register Date:	Jun 29 2020 01:29:37 PM				
Initialization:	Complete				
Synchronization:	100%				
Server Backlog:	0 files				
Last Logged In User(s):	WORKGROUP\APC-99986\$ APC-99986\bit9qa				

5. When Synchronization is 100%, shut down the computer or remove it from the network.
6. Go to the Computer Details page for the *clone computer you just updated* (not the original template), and click **Convert to Template** on the Advanced menu. The Computer Details page changes to a Template Details page.

7. The default name of the updated template is the old template name with a number appended to it to indicate how many times it has been updated. For example, if the original template was MYCORP\WIN7-64-IT, the edited template would be MYCORP\WIN7-64-IT (1), the next edited version would be MYCORP\WIN7-64-IT (2), and so on. You can change the name if necessary.
8. Create clones from the new template computer using your virtualization software.

Deleting a Template

You can delete a template at any time. If you delete a template that has clones, those clones become freestanding computers; that is, they lose their association with the template. Even if you restore the template computer at a later time with the same name, the clones do not reconnect with it.

To delete a template computer from the console:

1. On the console menu, choose **Assets > Computers**.
2. Locate the template computer using the Template Computers view or some other method.
3. In the Computers table, check the box next to the template computer, choose **Delete Computers** from the Action menu, and confirm the deletion.

Note

If a template has no clones, you also can convert it to a regular (non-template) computer and manage it with the App Control Server. See [“Converting a Template to a Regular Computer”](#) on page 221.

Configuring Clone Inventory

A primary reason to use App Control’s virtual machine management features is to optimize file inventory processing on future clones created from a template computer. There are two options for clone file inventory management:

- **All files** – The App Control Server can automatically initialize a clone’s file inventory based on the files present on the template. As soon as a clone is detected, the inventory from its template is copied into the clone’s inventory. Any future file additions or changes are also applied to the clone inventory. This is the default setting.
- **New and modified files** – You can choose to have the clone start with an empty file inventory and have the server track only file additions and changes for each clone that happen after it is created.

These options are set on a per-template basis. They affect how the Files on Computers inventory for *clones* are managed. Regardless of your choices here, all files from the template image are included in the App Control File Catalog on your server.

If you choose *New and modified files*, the clone inventory will track the following changes from the baseline template inventory:

- creation of a file
- modification of a file

- deletion of a file
- renaming of a file
- changes in a file's state (i.e., approvals and bans)

Changes in the *path* for a file (other than a change in the file name itself) will not cause a file that was in the template inventory to be tracked as part of the clone inventory.

Choosing an Inventory Option

The best choice for clone inventory will depend upon your environment and your priorities. The most obvious advantages of choosing to inventory only new and modified files are a reduction in network and server traffic and minimization of data that might not be of interest to you. This can be particularly important in an environment with thousands of cloned computers and a large base image of files.

Balance limitation of traffic against the impact of limiting the clone inventory. If new and modified files is chosen:

- Neither the Find Files page nor the Files on Computers page will be able to show all files on a clone computer.
- Drift reports that involve cloned machines will be incomplete. The only type of drift reporting that will work correctly is self-drift (comparison of the files currently on the clone computer with its own previous inventory), and unmodified files from the initial template image will not be included in this report.
- Snapshots created from cloned computers will include only new and modified files.
- File prevalence will not be accurate for unmodified files from the template image because instances on clones will not be counted (and also deletion of files from the original image will not be accounted for).
- Because unmodified files in the template image will not be visible in the clone computer's inventory, direct local approval of such files will require that the specific file instance (on the specific clone computer) appears in an event. Otherwise, global approval might be required.

Note

App Control also provides an option to exclude tracking of certain Microsoft-signed operating system and application files, and this can significantly reduce traffic and database demands. This affects all computers, not just clones. See [“Excluding Tracking of Microsoft Support Files”](#) on page 230 for more details.

To configure the clone inventory setting for a template:

1. On the console menu, choose **Assets > Computers**.
2. Locate the template computer using the Template Computers view or some other method, and click on the View Details button or the computer name.
3. Click on the **Template Settings** tab.

Carbon Black App Control Agent | Connection History | System Details | AD Details | **Template Settings** | Carbon Black EDR

Date Created: Jun 29 2020 04:51:03 PM
Original Computer Name: WORKGROUPVAPC-99986
Original IP Address: 10.210.65.144
Clone Count: 1 online, 0 offline
Clone Inventory: All files New and modified files
Specify method for deleting of offline cloned computers...
Clone Cleanup: Manual

Save Cancel

4. In the Clone Inventory field, choose the radio button for either **All files** or **New and modified files**.
5. If you have no other configuration changes, click **Save**.

Note

Even if you choose *New and modified files* for the inventory option, if a clone goes offline and then a clone with the same name is connected after its files are marked as deleted, the server will do a full inventory, including the files provided as part of the template image.

Deleting Clones

If you create and retire virtual machines on demand in the environment managed by App Control, you will want to make sure that old clones no longer in use do not remain on the Computers page. For example, you might have virtual machines automatically revert to their snapshot on a timed basis or every login, or you might frequently update the template image for your clones. App Control offers several ways of cleaning up old clones.

- **Manual cleanup** – If you choose, you can leave all cleanup to manual methods, periodically deleting offline clones through the Template Details page.
- **Automatic cleanup for all clones** – You can configure a cleanup rule that deletes offline clone computers on a schedule. You can delete *all* offline clone computers or only those matching a particular filter. For example, you could delete all computers that are running in a virtualized environment and are offline for more than 5 days.
- **Automatic cleanup per template** – You can configure different cleanup rules for different templates.

As with regular, non-clone computers, the file inventory for a deleted clone is deleted 24 hours after the clone is deleted.

Manual Cleanup of Clones

There are two primary methods of manual clone cleanup:

- You can locate a particular clone through the Cloned Computers Saved View and delete it as you would any other computer.
- You can go the Template Details page for a template and use the Delete Offline Clones command in the Action menu.

Automatic Cleanup for All Clones

The Advanced tab of the System Configuration page includes settings that remove offline computers from the list of managed computers. You can either choose to remove *any* computer from the console after it is offline for a certain period of time or you can set filters that selectively remove computers.

If you leave the Clone Cleanup configuration for each template on Manual, you can use the filtered global cleanup methods to remove offline clones. If you set an automatic cleanup method for one or more templates and set one of the global removal methods, offline clones will be removed whenever they meet *either* rule.

To create a global cleanup rule for offline clones:

1. On the console menu, click the configuration (gear) icon and choose **System Configuration**. The System Configuration page appears.
2. Click the **Advanced Options** tab. The Advanced Options configuration page appears.
3. Click the **Edit** button.
4. In the Old Computer Cleanup panel, configure Computers Matching Filter to delete clone computers after an amount of time you specify:
 - a. Check the box to the right of Computers Matching Filter.
 - b. Enter the number of days offline after which you want the computers deleted from the console Computers page.
 - c. On the Add Filter menu, choose an appropriate filter. For example, choose **Parent Template** and in the menu that appears next to Parent Template, choose **is not empty**. This assures that any computer with a template will be deleted. You also can choose **Virtualized** and check the Yes box to cleanup all virtual machines (whether or not they are clones) but not any clones created by other means. Or, you can choose **Virtual Platform** and enter **VMware** in the field to cleanup VMware computers.

Old Computer Cleanup

All Computers: delete after 30 Day(s) offline

Computers Matching Filter: delete after 3 Day(s) offline

Add filter

✖ Parent Template is not empty +

5. To save the changes, click the **Update** button and click **Yes** on the confirmation dialog.

Automatic Clone Cleanup for One Template

Each template has its own clone cleanup setting. You can choose manual cleanup or one of two automatic settings. If you also set a global clone cleanup rule on the System Configuration pages, templates are also subject to that rule.

To configure automatic clone cleanup for a specific template:

1. On the console menu, choose **Assets > Computers**.
2. Locate the template computer you want to configure for clone deletion and click its View Details button to open the Template Details page.
3. Click on the **Template Settings** tab. It shows when the template was created, the computer's original name and IP Address, and how many clones from the template have been seen by the App Control Server. It also includes a menu on which you can choose how to cleanup clone computers for this template.
4. On the Clone Cleanup menu, you can choose one of the following:
 - **Manual** - No automatic cleanup. Clones based on this template must be deleted manually, or by the global cleanup rule defined on the System Configuration Advanced tab.
 - **When offline** - Clones based on this template are scheduled for deletion as soon as they are offline. Depending upon other server activities, they will actually be deleted within 10 to 15 minutes.
 - **Based on time** - Clones based on this template are deleted if offline for a period of time you set in a field that appears when you make this choice. If there are two different times defined for the template and for global cleanup, the first deadline to be reached triggers the cleanup.
 - **Based on name** - When a clone based on this template is newly registered with the App Control Server, any *offline* clones with the same name are automatically deleted. Online clones are not affected. This method is safe to use unless you want to retain old reverted computer data for analysis. This will not cleanup offline clones if new clones always get a new name.
5. When you have completed any changes you want to make to the Template Details page, click **Save**.

Note: To prevent accidental deletion of clones that are still in use, clones that are detected as offline will not be deleted unless they have not communicated with the server for at least 10 minutes. This helps mitigate situations in which network interruptions erroneously make it appear that a clone is not in use.

Converting a Template to a Regular Computer

You can convert a template back to a regular computer. This feature is primarily intended as a remedy in case you accidentally convert a computer to a template. However, you can use it for any reason. If you want to convert a computer that was actually used as a template, make sure it does not have any clones listed on the Computers page in the console before the conversion.

To convert a template computer back to an agent-managed computer:

1. On the Template Details page for this template, click on **Show all cloned computers** in the Related Views menu.
2. If there are any clones listed for this computer, delete them from the App Control Server or leave the template in place as a template (see [“Deleting Computers”](#) on page 173). Otherwise the clones become freestanding computers (i.e., with no connection to a template).
3. When you have made certain that the template has no clones, return to the Template Details page and click **Convert to Computer** in the Action menu. The computer returns to App Control Server management and the Template Details page is converted to a Computer Details page.
4. After the conversion is complete, reconnect the computer so that the server can manage it.

Chapter 7

File, Publisher, and Application Information

This chapter describes:

- the location and contents of file information in App Control
- information about the publishers associated with these files
- information App Control collects about applications
- how to exclude tracking of certain files
- how to show all computers with or without certain files

File information from multiple App Control Servers can be viewed through a central management server. See [Chapter 27, “Unified Management of Multiple Servers.”](#)

Sections

Topic	Page
Overview	224
File Catalog	225
Files on Computers	227
Showing Individual Files	227
Finding Computers With or Without Specified Files	229
Excluding Tracking of Microsoft Support Files	230
File Groups	234
File Details Page	236
File Instance Details Page	242
Summary of File Views	247
Global File State	249
Local File State	250
Publisher Information	253
Application Information	256

Overview

App Control collects many different kinds of information about the “interesting” files it discovers on your computers. Interesting files are files that are either determined by App Control to be executable (for example, .EXE or .DLL files) or that match file extensions defined as scripts. You can use this information to be aware of file activity, or to make decisions about controlling execution and writing of particular files or classes of files.

Many files discovered by the App Control Agent have an identified *publisher*. As with other file information, a publisher’s name can be useful simply to know where a file came from, or it can be used to automatically approve or ban files.

Notes

Some file and publisher information is provided by the Carbon Black File Reputation. You must have Carbon Black File Reputation integration activated to receive this information. See [“Activating Carbon Black File Reputation”](#) on page 763 for more information.

For information about using file and publisher information to approve or ban files, see [Chapter 8, “Approving and Banning Software.”](#)

File information is presented in table form in several places within the console, but the primary starting point is the Files page, which you access by choosing **Assets > Files** on the console menu. The Files page has two tabs:

- The **File Catalog** tab shows the unique “interesting” files discovered on your computers. Cataloged files includes those currently present and tracked on the fixed, local drives of agent computers, files considered “interesting” but not tracked in inventory, and files that were once present on an agent system but have been deleted.
- The **Files on Computers** tab shows tracked file instances. This includes every instance of every “interesting” file on the fixed, local drives of every agent-managed computer reporting to your App Control Server (once their files are fully processed), with these possible exceptions:
 - You can exclude common Microsoft operating system and application support files from the file inventory to reduce tracking overhead and database size. See [“Excluding Tracking of Microsoft Support Files”](#) on page 230 for details.
 - You can exclude instances of files that are in the template used by a VDI product to create a clone. See [“Configuring Clone Inventory”](#) on page 217 for details.
 - You can disable file tracking on a policy-by-policy basis.

If any of these conditions affects files on your computers, the Prevalence value for those files will not be accurate – only tracked file instances contribute to Prevalence.

For complete information about one file in a table, you can go to a details page for the file:

- The **File Details** page shows the global information about one unique file and provides a link to a list of all instances of that file.
- The **File Instance Details** page shows information about a specific file instance on a specific computer.

The table of publishers for files discovered on agents is shown on the **Publishers** tab of the **Software Rules** page. If you want complete information about one publisher in the table, you can go to the details page for the publisher.

App Control agents collect information about Windows applications installed on agent systems and displays the information on the **Applications** page. The page includes an **Application Catalog** tab, which shows each unique application on agent systems reporting to this App Control Server, and an **Applications on Computers** tab, which shows each instance of each application on each computer.

Most of the information about an application is from the metadata of the installer file for the application or the application's registry data. It is similar to the information on the Programs and Features page of the Windows Control Panel.

See [“Application Information”](#) on page 256 for more information.

Viewing File Tables

File Catalog

The File Catalog tab on the Files page shows unique files discovered on computers running the App Control Agent and reporting to your App Control Server(s). In addition to displaying tables of files and their details, the File Catalog page has an Action menu for taking file-related actions, including approving, banning and looking up information about files in Carbon Black File Reputation. These actions are described in other chapters.

From the File Catalog, you can open a File Details page by clicking on the View Details button next to a file name. The column headings available in the File Catalog correspond in most cases to fields on the File Details page for a single file. See [Table 26, “File Details and File Catalog Page Fields,”](#) on page 237 for a description of this information.

The screenshot shows the 'Files' page with the 'File Catalog' tab selected. The interface includes a 'Saved Views' dropdown set to '(none)', a 'Group By' dropdown set to '(none)', and a 'Max Age' dropdown set to '(none)'. There are also buttons for 'Add', 'Show Filters', 'Show Columns', 'Show Snapshot', 'Export to CSV', and 'Refresh Table'. Below the controls, it says 'Showing 75 out of 18997 item(s)'. The table below has the following data:

	First Seen Date	First Seen Name	Publisher or Company	Product Name
<input type="checkbox"/>	Jun 23 2020 05:34:05 PM	parity.sys	Carbon Black, Inc.	Cb Protection™
<input type="checkbox"/>	Jun 23 2020 05:34:05 PM	ntosknl.exe	Microsoft Windows	Microsoft® Windows® Operating System
<input type="checkbox"/>	Jun 23 2020 05:34:06 PM	wbemsvc.dll	Microsoft Corporation	Microsoft® Windows® Operating System
<input type="checkbox"/>	Jun 23 2020 05:34:06 PM	ntosknl.exe	Microsoft Windows Publisher	Microsoft® Windows® Operating System

By default, the File Catalog shows all files. You can choose a different Saved View of the catalog or create a view of your own to focus on particular types of files or search for one file. If you have not already become familiar with modifying views in console tables, see [“Console Tables”](#) on page 66. You also can choose to show unique *top-level* files only (files not known to have been installed by or copied from another file) instead of all files. See [“Showing Individual Files”](#) on page 227 before choosing this option.

The File Catalog shows the first seen name of a unique file, and the unique file is identified by its hash. The name of a file instance *on a particular computer* might not appear in the File Catalog even though it appears in the Files on Computers tab. Use Find Files or the Files on Computers tab to locate a particular instance by name.

Table 25 shows the Saved Views provided on the File Catalog tab.

Table 25: Saved Views on the File Catalog tab

Saved View	Description
Applications by Publisher/Company	Files that are identified as Applications or Packages; in this view, they are grouped by Publisher (if available) or Company. Note: You can also view information about Windows applications in the separate Applications Catalog (Assets > Applications).
Approved Files	All executable files approved by a global approval method.
Banned Files	All files explicitly banned by hash. Files banned by name do not appear in the table on the File Catalog tab. Files that are banned for some policies but not others do not appear in the Banned Files table, but can be found in the File Catalog tab by using the File State filter.
Categorized Files	Files that exist on at least one computer and fall into one of the application categories identifiable by Carbon Black File Reputation (such as Hacking Tools and Instant Messaging). In this view, the files are grouped by category.
Existing Files	Files that exist on at least one agent-managed computer reporting to your server.
Installed Programs	Files grouped by the installed program with which they are associated. This view shows the full package or application name for the installed programs. Platform Note: Only Windows files are identified as Installed Programs.
Malicious Files	Files that exist on at least one computer and have been identified by Carbon Black File Reputation as having a Threat level of 1-Potential risk, or 2-Malicious.
New Unapproved Files	Unapproved files that appeared on computers <i>after</i> file initialization, that have not been Acknowledged, and that still exist on at least one computer.
Removed Files	Files that no longer exist on any agent-managed computer reporting to your App Control Server.
Reputation Approvals	Files that have been approved because of the trust rating of the file or its publisher in Carbon Black File Reputation.
Trusted Packages	Top-level files, located in a Trusted Directory, that are the common source or installer files for other files. Click the View Details button to display the File Details page for the package itself. Click on the package name for a table of associated files written by the package. The root file for each package may also appear in other tabs.

Files on Computers

The Files on Computers tab provides a table of files that are on agent computers or, for disconnected computers, were on those computers when their agents last communicated with the App Control Server. Files from deleted computers may continue to appear for one day but will be marked as being from a deleted computer during that time and will no longer appear after the grace period.

By default, the Files on Computers table shows all individual files on all computers. You can choose a different Saved View of the catalog, however, or create a view of your own to focus on particular types of files or search for one file. If you are not already familiar with modifying views in console tables, see [“Console Tables”](#) on page 66. You also can turn off the *Show individual files* setting to instead show only top-level files (files not known to have been installed by or copied from another file) plus groups of initialized files (i.e., files on a computer when the App Control Agent was installed). See [“Showing Individual Files”](#) before choosing this option.

The Files on Computers tab includes the following subset of the Saved Views shown in [Table 25, “Saved Views on the File Catalog tab”](#) on page 226:

- Applications by Publisher/Company
- Banned Files
- Categorized Files
- Installed Programs
- Malicious Files
- Threat Report - Suspicious Files by Name
- Unapproved Files

[Table 26](#) shows the fields that can appear in the File Catalog table, most of which also can appear in the Files on Computer table. [Table 27](#) shows additional fields that are available on the Files on Computers tab. Note that not all fields appear by default.

Showing Individual Files

The checkbox labeled *Show individual files*, in the top right area of both Files page tabs, has a major effect on what files are shown.

When this box *is* checked (the default), the Files page shows all files – both top-level files (files *not* known to have been installed by or copied from another file) *and* files installed by other files. When *not* checked, the File page shows only top-level files (files *not* known to have been installed by or copied from another file). On the Files on Computers page, unchecking this box adds an <Initialized Files> row for each computer, grouping all of a computers files that were present during initialization into one “group” that can be clicked on to expand.

With the default view showing individual files, a complete File Catalog listing of the unique files reported to the App Control Server might number in the tens of millions. Files on Computers, which is an inventory of files actually on your computers, can be significantly larger. In rare cases, especially with a particularly large number of App Control Agents and/or an underpowered database server, attempting to show all individual files can cause the App Control Server to time out. In that case, consider modifying the view. For example, you could turn off *Show individual files*, change the *Group by* choice, or sort by a different column. You also can use a filter to limit the total number of files shown.

A possible side-effect of requesting a table with a very large number of files is that the number of items on all pages of the table, shown in the lower left corner, will show as an approximation, such as *More than 10000 items*. This can also occur if a view you request requires extra processing by the App Control Server, even if the number of results is not especially large. Clicking **Refresh Page** after the results are displayed often shows the exact number.

Keep in mind that you can click on the name of a top-level file in the File Catalog or Files on Computers page to get a list of the individual files associated with it.

Platform Note: For this release of App Control, only Windows files are grouped by installer, so checking *Show individual files* does not change the files shown from non-Windows computers in the File Catalog.

Initialized Files

File *initialization* is the file inventory process that begins immediately after installation of the App Control Agent on a computer. The agent takes an inventory of all executable files on the client computer's fixed, local drives and creates a hash of each file. When a computer first connects to the server, its agent sends each hash to the App Control Server to update the server's file inventory. Files on a computer at initialization receive a *local* state of Approved unless they already have been identified and globally banned or banned by policy on the App Control Server.

In the Files on Computers table, when *Show individual files* is not checked, each agent-managed computer has a row with the file name **<Initialization files>**.

	Date Created	Computer	File Name	Publisher or Company	User Name
<input type="checkbox"/>	Jun 29 2020 07:40:38 PM	WORKGROUP\APC-79895	<Initialization files>		
<input checked="" type="checkbox"/>	Jun 29 2020 08:25:14 PM	WORKGROUP\APC-79895	103e3.msi	VMware, Inc.	
<input checked="" type="checkbox"/>	Jun 29 2020 08:25:14 PM	WORKGROUP\APC-79895	17fb1.msi	Microsoft Corporation	

Clicking on **<Initialization files>** in a row opens a table showing all initialized files for one computer. This is a useful way to determine what was on each system before the agent was installed.

If you disable and then re-enable an agent, a new initialization process begins, and the **<Initialization files>** group will change. Other than that, the files in this group do not change unless there is a problem with the agent. Upgrading the agent does not change the list of initialized files.

If you click on one of the files, it will show a list of File Groups that contain the file but it will not identify the group containing it for the current computer. This is because a file that predates the agent may have been installed or copied from any one of a variety of places.

If you use a filter with *Initialized = Yes* on the Files on Computers page with the *Show individual* box not checked, the table shows rows for **<Initialization files>** and usually several other files. The other files are known installers, but are also included under the **<Initialization files>** group.

Menus on the File Tables Pages

The File Catalog and Files on Computers tables each have an Action menu in the upper left above the table. [Table 28, “Menus on File Tables and Details Pages,”](#) on page 246 shows the available choices on these menus. Note that some choices are available only for certain file states.

Finding Computers With or Without Specified Files

If you add an application to your environment or update an existing program with a new file, you might want to determine whether any computers are missing the file or files involved in this change. On the other hand, you might have found one or more files that you want removed from your environment. In that case, it would be helpful to be able to get a list of all computers that still have these files. The Files pages include menu choices that provide this information, using the file hash as the search parameter.

On the Related Views menu of both the File Details and File Instance Details pages, you have the following computer-search options:

- **Computers with this file** – Shows the Computers table with a list of all computers that have the file described on the details page.
- **Computers without this file** – Shows the Computers table with a list of all computers that *do not* have the file described on the details page.

On the File Catalog, Files on Computers, and Find Files results pages, you can check multiple files in the table and use any of the following Action menu commands:

- **Find computers on <this server> with at least one of the selected files** – Shows the Computers table with a list of all computers that have at least one of the files checked on the Files page.
- **Find computers on <this server> with all of the selected files** – Shows the Computers table with a list of only those computers with *all* of the files checked on the Files page.
- **Find computers on <this server> missing at least one of the selected files** – Shows the Computers table with a list of all computers that are missing any of the files checked on the Files page.
- **Find computers on <this server> missing all of the selected files** – Shows the Computers table with a list of only those computers missing *all* of the files checked on the Files page.

When the Computers page shows the results of any of these commands, a legend appears above the Saved Views menu indicating what kind of command was used and the SHA-256 hash value for the file. You can hover the mouse cursor over the hash value to get more information about the file (including its first seen name).

If you want to eliminate the filtering that produces these results and instead view the standard Computers page view, click the **Remove** link next to this legend.

Computers

Computers connected: 1 Total computers: 1 Current CL version: 1143 CL version for upgrade: 1142

Computers with file with SHA-256: 0f91ba57b0db9eb41c083762fa5bbac71958a42cfd91b3cc8c6ec46dab53aaf [Remove](#)

Saved Views: (none) Add Group By: (none) Ascending Days Disconnected: (none)

Show Filters Show Columns Export to CSV Refresh Page

Action Search: Go Clear

<input type="checkbox"/>	Computer Name	Connected	Policy Status	Upgrade Status	Connected Enforcement
<input type="checkbox"/>	WORKGROUP\APC-79895	●	Up to date	Up to date	Low (Monitor Unapproved)

1 item Page 1/1

Note

Files excluded from the Files on Computers inventory cannot be located using these commands. For example, if tracking of Microsoft support files has been turned off on the Advanced tab of the System Configuration page, you cannot get accurate results for those files from any of the find computers commands.

Excluding Tracking of Microsoft Support Files

By default, App Control inventories and tracks all instances of interesting files on all agents attached to a server. Many of these files are Windows operating system and Microsoft application files and related system updates. As Windows has evolved, the number of operating system files has multiplied to several times what it was in Windows XP, and applications have had similar increases in file number. Windows updates are also significantly increasing in size. Because of these increases, Microsoft files may account for more than half and in some cases three-quarters of all of the files found in your inventory for Windows computers.

If you trust and approve files from Microsoft, you might also prefer not to track them. App Control provides two options that eliminate file tracking for certain files that have been signed by the publishers "Microsoft Windows" or "Microsoft Corporation". By turning off file tracking for a significant percentage of the file instances on your systems, you can reduce the size of the server database needed for a given number of agents as well as reducing the load on the server that would be required to process these files.

The two options allow you to choose where you exclude the information about the support files signed by "Microsoft Windows" or "Microsoft Corporation" publishers:

- **Discard at the server** – If you choose this option, information about locally approved instances of these files will still be sent to the server and included in the File Catalog but the files will subsequently be purged from the Files on Computers inventory. This eliminates your visibility into these files on endpoints. While reporting of events related to these files will be limited, events such as approvals and bans continue to be reported.
- **Discard at the agent** – If you choose this option, information about locally approved instances of these files will be discarded at the agent and not sent to the server. In

addition, events associated with the files will be further suppressed (although not completely eliminated). These files will not appear in Files on Computers, and they generally will not appear in the File Catalog unless an execution or other tracked action related to them occurs.

For both of these options, you will see a warning that exclusion of tracking for these files could hamper further investigations.

Tracking of Microsoft-signed support file instances is controlled on the Advanced Options tab of the System Configuration page.

The screenshot shows the 'System Configuration' page with the 'Advanced Options' tab selected. The 'Database Backup' section includes fields for Backup Type (Network), Backup Path, Username, Password, and Windows Domain, with an 'Enabled' checkbox and a status of 'Idle'. The 'Carbon Black App Control Agent' section shows 'Automatic Agent Upgrades' set to 'Enabled'. The 'Full OS Inventory Tracking' section has three radio button options: 'Track inventory for locally approved support files signed by "Microsoft Windows" or "Microsoft Corporation" publishers' (selected), 'Discard information about locally approved support files signed by "Microsoft Windows" or "Microsoft Corporation" publishers at server', and 'Discard information about locally approved support files signed by "Microsoft Windows" or "Microsoft Corporation" publishers at agent'. A red box highlights these three options. The 'Resource Download Location' is set to 'https://**ServerIP**/hostpkg/pkg.php?pkg='.

To disable or re-enable tracking of Microsoft-signed support file instances:

1. On the console menu, click the configuration (gear) icon and choose **System Configuration**. The System Configuration page appears.
2. Click the **Advanced Options** tab, and at the bottom of the page, click the **Edit** button.
3. To disable tracking of Microsoft-signed support file instances, choose the radio button for one of the following options in the Full OS Inventory Tracking section:
 - *Discard information about locally approved support files signed by "Microsoft Windows" or "Microsoft Corporation" publishers **at server*** – Details of these files are still sent to the server and included in the File Catalog, but the files are purged from the Files on Computers inventory. Events related to these files will be fewer but not completely eliminated.
 - *Discard information about locally approved support files signed by "Microsoft Windows" or "Microsoft Corporation" publishers **at agent*** – Details of these files are not be sent to the server. Instances of these files (including those discovered while full tracking was enabled) will not appear in Files on Computers, and unless they were discovered before this option was configured or are part of a tracked event, they will not appear in the File Catalog. Events associated with them will generally be suppressed.
4. At the bottom of the page, click the **Update** button and, if you are certain you want to make this change, click **Yes** on the Confirmation dialog. Support file tracking is disabled.

5. To re-enable tracking of these files, click the *Track inventory for locally approved support files ...* radio button and click the **Update** button. If you are certain you want to make this change, click **Yes** on the Confirmation dialog.

Note: The Confirmation dialog for re-enabling Full OS Inventory Tracking includes a warning about possible effects on product performance. If you have been operating with this setting off for more than a brief time, consider whether your environment meets the requirements for significant additional file traffic.

Files Instances Affected

When you choose an option to discard files for Full OS Inventory Tracking, instances of files meeting *all* of the following criteria are no longer tracked:

- The publisher must be "Microsoft Windows" or "Microsoft Corporation". This includes directly signed files and those signed with a detached publisher. Files signed by other Microsoft publishers, even if legitimate, continue to be tracked.
- The file must be a support file, such as a .DLL, that would normally be considered "interesting" and therefore be tracked by App Control. Tracking of .EXE files or the events related to them is not affected by this option.
- The file must be locally approved, either directly or because of an approval rule.

Changes that Affect OS Inventory Tracking

As with other rules, there are interactions between the Full OS Inventory Tracking rule and several other rules and conditions in App Control:

- **File State Transitions** – If file exclusion is enabled (that is, Full OS Inventory Tracking is *disabled*), unapproved file instances that otherwise meet the exclusion criteria are inventoried and tracked. If these files are later approved, they are no longer tracked, but the server will not incorporate their state change; they remain in inventory even though their prevalence will show as zero. Conditions that can cause this include:
 - A Microsoft support file was locally *unapproved* and so not excluded from inventory, but it was later locally approved.
 - The criteria for publisher trust was high when Full OS Inventory Tracking is disabled (for example, minimum key size for approval is 2048), and so Microsoft support files were *not excluded*. Then, the publisher trust criteria was lowered (for example to a 1024-bit minimum), approving most of the support files.
- **Disabling Tracking** – If you disable Full OS Inventory Tracking, the following occurs:
 - All affected files are deleted from the file inventory on the Files on Computers page. Deletion will happen in the background, while the server is not busy, and could take several days to complete, depending on inventory size. An event will report how many files were deleted from the inventory.
 - New, approved instances of these files (and changes to them) will not be inventoried or tracked.
- **Re-Enabling Tracking** – If you re-enable Full OS Inventory Tracking after it has been disabled, there will not be an automatic re-inventory of Microsoft-signed files from agent computers. New instances or activity related to relevant files will be tracked. If you want to collect an inventory of all pre-existing Microsoft support files, you can Resynchronize all File Information on a computer-by-computer basis. This option is available on the Computers page Action menu.

- **Agent Version** – You can apply either of the options for turning off tracking of Microsoft support files to agents at version 7.2.1 and greater, and these agents will behave as documented. You can also turn off tracking on (supported) older agents, but the behavior is different. The server cannot always exclude files from older agents immediately because it is missing some of the necessary information. For example, it will not always be able to detect that a file is a supporting file, or that the file is signed by Microsoft. However, if you choose one of the exclusion options, information about these files will be deleted on the server in the background during a regular daily update of file information.

Information about Excluded File Instances

Even if you turn off tracking of approved Microsoft support file instances, information about them is available. Some of this is generic information about the file itself, not its specific instances. These files still appear in the File Catalog if a file with their hash has appeared on any agent monitored computer. Because instance tracking is turned off, the file Prevalence (the number of computers they are found on) number will not be reliable (and might be zero), and a tooltip will indicate that prevalence cannot be calculated.

Publisher or Company	Prevalence	Excluded from Inventory ▾	Product Name
Microsoft Corporation	Prevalence for this file is not accurate because file has been excluded from the inventory		
Microsoft Corporation	1	Yes	Assembly imported from type library *
Microsoft Corporation	1	Yes	Microsoft SQL Server
Microsoft Corporation	1	Yes	Microsoft SQL Server

It is possible that you will want to turn off tracking of these files in general but track specific instances, for example, if a particular version of a Microsoft DLL has a reported vulnerability and you want to replace it. There are several ways to maintain the general setting so that you can reduce the load from these files but also track executions of certain files:

- **Report Bans** – You can create a report-only ban for a file. This will cause all instances of this file on all computers to be added to the inventory.
- **Meters** – If you create a Meter for a file hash, the meter will report all executions of an excluded file as events but not add instances of it to the Files on Computers inventory.
- **Exports of Data to Analytics Tools** – If you have integrated App Control with an External Analytics tool, such as Splunk, data from excluded file instances is included with all the other file and event data. You can use the external tool to find all instances of excluded files as they appeared historically on all computers, and executions of these files are also tracked in the data provided to the external tool.
- **Excluded from Inventory Column in File Catalog** – The optional column *Excluded from Inventory* is available in the File Catalog. If you add this to the table, it identifies files whose instances are not in the file inventory because they are excluded OS support files.
Note: Files locally approved *after* Microsoft support file exclusion was activated will continue to appear as *unapproved* files and so appear in the Files on Computers inventory.

If you have Carbon Black EDR sensors installed on your computers in addition to the App Control Agent, Carbon Black EDR will continue to detect and report executions of these files.

File Groups

Platform Note: For this release of App Control, only files on Windows computers are grouped by installer, so this section does not apply to other platforms.

As files are being installed on a computer, the App Control Agent groups them according to its analysis of what process is installing them. This group name might be unique, or it might be an installer name common to multiple groups – “setup.exe”, for example.

Once installation is complete, the agent scans the Windows program database to see whether these files can be associated with a “Programs and Features” entry. If so, files will be regrouped under the file that is used for modifying or removing corresponding programs. If no Programs and Features entry is found, installed files will retain the initial group name.

The screenshot shows the 'Files on Computers' view in the VMware Carbon Black App Control console. A table lists 12 files. The file '103e3.msi' is highlighted, and its details are shown in a pop-up window. The details window shows files associated with 'VMware Tools' on computer WORKGROUP\APC-79895. The details window also shows a table of 75 files associated with this group.

Date Created	Computer	File Name	Installed Program
Jun 29 2020 07:40:38 PM	WORKGROUP\APC-79895	<Initialization files>	
Jun 29 2020 08:25:14 PM	WORKGROUP\APC-79895	42f49.msi	Microsoft Visual C++ 2017 x86 Minimum Runtime - 14.14.26405
Jun 29 2020 08:25:14 PM	WORKGROUP\APC-79895	42f4d.msi	Microsoft Visual C++ 2017 x86 Additional Runtime - 14.14.26405
Jun 29 2020 08:25:14 PM	WORKGROUP\APC-79895	103e3.msi	VMware Tools
Jun 29 2020 08:25:14 PM	WORKGROUP\APC-79895	42f51.msi	Microsoft Visual C++ 2017 x64 Minimum Runtime - 14.14.26405
Jun 29 2020 08:25:14 PM	WORKGROUP\APC-79895	42f55.msi	Microsoft Visual C++ 2017 x64 Additional Runtime - 14.14.26405

Date Created	Computer	File Name	Publisher or Company
Jun 29 2020 08:13:17 PM	WORKGROUP\APC-79895	vsocklib.dll	VMware, Inc.
Jun 29 2020 08:13:17 PM	WORKGROUP\APC-79895	vmhgfs.dll	VMware, Inc.
Jun 29 2020 08:13:17 PM	WORKGROUP\APC-79895	vmquestlibjava.dll	VMware, Inc.
Jun 29 2020 08:13:17 PM	WORKGROUP\APC-79895	vmquestlib.dll	VMware, Inc.
Jun 29 2020 08:13:17 PM	WORKGROUP\APC-79895	vm3dum_loader.dll	VMware, Inc.
Jun 29 2020 08:13:17 PM	WORKGROUP\APC-79895	vm3dum_10.dll	VMware, Inc.
Jun 29 2020 08:13:17 PM	WORKGROUP\APC-79895	vm3dum_10-stats.dll	VMware, Inc.

Group names are used wherever files are listed in the console. Examples include:

- On the File Catalog and Files on Computers pages, you can choose the Installed Programs Saved View to see a list of applications.
- In Baseline Drift Report Results, if you are looking at a Files view, you can group by Installed Program to see how much drift is attributable to each application.
- If you click on a highlighted file name in the File Catalog, you see a File Group Details page that lists all of the files associated with the file you clicked on, and usually showing the application they are part of. This is the aggregate of all unique files installed by the highlighted file, on all computers running the App Control Agent.
- If you click on a highlighted file name in the Files on Computers page, you see a File Group details page listing all files associated with the file instance you clicked on.

- If you click on a <Initialization Files> in a row on the Files on Computers page, you see a list of all files that were present on the computer named in that row at the time the App Control Agent was last initialized (normally, when the agent was installed).

Viewing Details Pages

The console provides two different details pages for files it manages:

- **File Details** – For *each unique file* discovered on computers running the App Control Agent, you can open a File Details page, which provides global information about the file and allows you to modify various global parameters for the file. The File Details page presents complete information for unique files listed in the File Catalog table.
- **File Instance Details** – For *each instance of a file* discovered on a computer running the App Control Agent, you can open a File Instance Details page, which provides information specific to that instance in addition to some of the global information seen on the File Details page; it also allows you to modify both instance and global attributes of the file. The File Instance Details page presents complete information for instances of files listed in the Files on Computers table.

The following sections provide an overview of file details pages, including tables of menu commands on these pages. More detailed descriptions of activities you can perform on these pages are provided elsewhere in the *VMware Carbon Black App Control User Guide* guide, especially in [Chapter 8, “Approving and Banning Software.”](#)

File Details Page

The File Details page shows details of the global state of a file. In any table showing unique files, such as the File Catalog, you click the View Details button to open the File Details page.

File Details

General

First Seen Name: mpavdlta.vdm
First Seen Date: Jun 29 2020 09:11:23 PM
Last Updated: Jun 29 2020 09:11:23 PM
First Seen Path: c:\windows\temp\ddaa7a80-edfb-4fa7-98e8-aeb1f36b9a021240.1d64e7af91e87a2\
First Seen Computer: [WORKGROUP\APC-79895](#)
First Seen Platform: Windows
Extension: vdm
Global State: Unapproved
Global State Details: File is unapproved, Publisher is unapproved, Certificate is Unapproved
Flags: (none)
Installer / Updater: No
File Prevalence: File exists on 1 computer

File Properties

Publisher: [Microsoft Corporation](#)
Publisher State: Unapproved
Certificate: [Microsoft Corporation Microsoft Corporation Redmond Washington US](#)
Certificate Type: Embedded Signer
Certificate Global State: Unapproved
Company: Microsoft Corporation
Product Name: Microsoft Malware Protection
Product Version: 1.319.470.0
File Size: 4,439,984 bytes
Description: AntiVirus Definition Update
File Type: Application
SHA-256: 812090C15A994B3284D1AEA5AAB35D8A4F74DB04672E3AB22BB65D81EF2AD85E
MD5: 652C13EEB68260571F03946439597C58
SHA-1: 19232FFCD763747C9887901B6529E1792564D546

Groups that contain this file

[Mpsigstub.Exe](#) Find all files contained in this group

History

Jun 29 2020 09:11:23 PM The file appeared on [WORKGROUP\APC-79895](#) post installation

[Table 26](#) shows the information and actions available on the File Details page. Certain global file attributes are captured only for the “first seen” instances of the file seen by a App Control Agent. These are labeled as such on the File Details page.

Table 26: File Details and File Catalog Page Fields

Field	Description
General panel	
First Seen Name	File name of the first file to have this hash discovered by an agent managed by this App Control Server.
First Seen Date	Time the first file with this hash was seen on a network computer, displayed in the format: MM DD YYYY hh:mm:ss(AM/PM).
Last Updated	Last date and time when the file metadata was updated. (Not affected by App Control-provided data, e.g., prevalence or trust).
First Seen Path	Path of the first file to have this hash reported to this server.
First Seen Computer	Name of the computer on which the file was first seen. Click on this name to get the Computer Details page for this computer. If you later delete the first-seen computer from the system, it is no longer associated with the file and this field is blank.
First Seen Platform	Platform (Windows, Mac, Linux) on which this file was first seen by this server.
Extension	File extension of the first seen file to have this hash.
Global State	Global State is a combination of File State and Publisher State, and indicates the overall approval state for all systems or by policy. Files can be globally approved by hash or publisher. The possible values are Approved, Banned, Unapproved, Approved by Policy, Banned by Policy, and Mixed. Global State is Mixed when a file is approved in some policies, but banned in other policies. For example, a file could be banned by hash in some policies, and approved by publisher in the remaining policies.
Global State Details	The File State and Publisher State contributing to Global State.
Flags	File-state metadata for use by Carbon Black Support engineers. Your support representative may ask you to report this information.
Installer/Updater (in File Details) Installer (in File Catalog)	Indicates whether either App Control analysis or a user choice has identified this file as an installer or updater (which means that if the file is approved, so are all files that it creates). Yes – Treat the file as an installer that will expand to create more files. If this file is approved, files it writes will be locally approved. No – Treat the file as non-expandable.
Reputation Enabled	Indicates whether reputation-based approval is enabled for this file (Yes or No). Note: This only displays if Carbon Black File Reputation is properly configured.
File Prevalence (in File Details) Prevalence (in File Catalog)	The number of computers on which this file exists. You can use the Add Alert command on the Actions menu to add an alert that triggers when the prevalence of a file reaches a set level. See “Using App Control Alerts” on page 608 for details.

Field	Description
View Carbon Black File Reputation Data (button)	Click to get a detailed analysis (if available) of this file from Carbon Black File Reputation. Button appears on the File Details page after you activate Carbon Black File Reputation. See “Activating Carbon Black File Reputation” on page 763.
File Properties panel	
Publisher	If the file is digitally signed or was included in a digitally signed package, the console displays the publisher (software manufacturer) of the associated application.
Publisher State	The approval state of the publisher. Values are Approved, Approved by Policy, Banned, Banned by Policy, and Unapproved. Does not appear if the publisher is unknown.
Certificate	The Subject Name for the certificate that signed this file.
Certificate Type	For leaf certificates, certificate type indicates what the leaf certificate is being used for and how it is associated with a file. Type is some combination of the following terms: Embedded, Detached, Signer, Cosigner.
Certificate Global State	The effective state of the certificate. Values are Unapproved, Approved, Banned, Approved By Policy, Banned By Policy, Mixed
Company	The Company name (if provided) in the file metadata.
Product Name	The Product Name (if provided) in the file metadata.
Product Version	The Product Version (if provided) in the file metadata.
File Size	Size of the file (in bytes).
Description	The Description (if provided) in the file metadata.
File Type	<p>One of the following:</p> <p>Application – Any executable (e.g., .exe or .com) except for Packages</p> <p>Supporting File – Any library loaded by an executable (e.g., .dll, .ocx, .sys)</p> <p>Package – Any installer (.exe with contents, such as a self-extracting zip or setup program)</p> <p>Script File – Any script or batch file (e.g., .bat, .vbs, .wsf)</p> <p>Other – Reserved for future types</p> <p>Unrecognized Executed File – A file that was not identified as an executable by App Control during initialization or later analysis, but that some process attempted to execute. The execution attempt adds the file to the lists of files tracked and managed by the App Control Server and Agents.</p> <p>Unknown – Files reported by older App Control Agents that don't provide file type information</p>

Field	Description
SHA-256	<p>Hash (data signature) of the file created using App Control's proprietary SHA-256 algorithm. SHA-256 is used internally as the preferred hash for files tracked by App Control.</p> <p>SHA-256 hashes created by the App Control algorithm may be identical to those created by other means. However, some files change their hash every time they are installed because they include date, location, or other context-specific information not relevant for tracking purposes. For files known to do this, App Control uses a special fuzzy hashing algorithm that eliminates this extraneous variation, and so shows every instance of such files on computers running App Control Agents to be identical. When this algorithm has been used, the hash is identified as "SHA-256 (Normalized)".</p> <p>You can search for files by hash using filters on the Files page or the Find Files page. All File Instances in the Related Views menu provides a way to do this directly from the File Details page.</p>
MD5	MD5 is a widely used hashing algorithm. App Control provides this alternate hash in case you or the system needs to identify the file against a list of published MD5 hashes.
SHA-1	SHA-1 is another widely used hashing algorithm. App Control provides this alternate hash in case you or the system needs to identify the file against a list of published SHA-1 hashes.

Carbon Black File Reputation Information panel

Trust	<p>Indicates the level of trust for the file based on Carbon Black File Reputation information such as file source and certificates. The trust rating is showing on a scale of 0 (none) to 10 (most trusted), along with a graphic meter reflecting this rating. Trust for a file also might be unknown, in which case this field is blank in the column for that file and shows "(unknown)" in its details page.</p> <p>The value of this field is a subjective assessment of the file's integrity. As an indication of whether the file appears to be safe based on information derived from Carbon Black File Reputation analysis, the trust value does not signify actual approval on the App Control Server. However, you can use Reputation Rules to automatically approve files based on their trust rating or the trust rating of their publisher.</p>
Threat level (in File Details) Threat (in File Catalog)	<p>If Carbon Black File Reputation is configured, discovered files are automatically submitted for threat analysis. Carbon Black File Reputation flags known malware with a red x icon. No flag indicates that the file was not recognized as malware, not necessarily that it is safe. Threat levels include:</p> <ul style="list-style-type: none"> 0 - Clean 1 - Potentially malicious 2 - Malicious Unknown - Not identified

Field	Description
Category	If you have configured Carbon Black File Reputation, this shows the category this file is in (e.g., Entertainment, Hacking Tools, Instant Messaging, Media Players). Category may be unknown, and is not displayed on the details page in this case.
Policy Specific States	Indicates ways in which the file is treated differently in particular policies. For example, if the file is under a policy-specific hash ban or approval, the policy name is shown here. Does not appear if there is no policy specific treatment of the file.
Carbon Black EDR panel (all data is from Carbon Black EDR Server)	
First Seen Activity	Date and time when activity involving this file was first reported to the Carbon Black EDR server.
Watchlists	Number of Carbon Black EDR Watchlists that this file is on.
Frequency Data	Number of hosts that have observed the binary with this MD5 hash value.
Unique Paths	Number of unique paths in which this file has been seen
Network Connections	Number of network connections that the execution of this process either attempted or established.
Registry Modifications	Number of registry modifications made because of execution of this file.
File Icon	Desktop icon associated with this file (if any).
More information	Links back to the Carbon Black EDR console to get additional information about the file.
External Analysis Results panel	
<productname>	If you integrated App Control with a supported network security device or service, and you have correlated notifications from that source with the App Control File Catalog, files that match malicious or potential risk notifications from the third party source show those results in this panel. Possible options are: <ul style="list-style-type: none"> • Check Point • Palo Alto Networks WildFire • CB Inspection
Group Information panel	
<group name>	If a file is the root of a group, this indicates the group name (usually the file name) and how many files are in the group. Note that tools such as browsers may appear as the root of a group because they download files. These files may appear as group members even though they are unrelated to the tool in any other way.

Field	Description
Groups that contain this file panel	
<group names>	<p>If a file is associated with a group, this panel indicates the group(s) with which this file is associated and the root file(s), if known, of the group(s). Some files may be installable by multiple root files (or be copies of another file), and so they will show multiple groups here.</p> <p>Each group shown includes a <i>Find all files contained in this group</i> link that opens the File Group Details page to show the results.</p>
History panel	
<dates and times>	<p>Indicates whether the file was identified on the first-seen computer during initialization or detected after initialization.</p> <p>Also indicates any approvals or bans applied to the file.</p> <p>Files detected <i>after</i> initialization are tracked as unapproved files until approved or banned, and may be viewed in the New Unapproved view on the Files page File Catalog tab.</p>
Fields in table only	
Acknowledged	Indicates whether a console user acknowledged this file (Yes or No). You can acknowledge a file using the Action menu on the File Catalog tab. This can help distinguish files you already know about from new arrivals. Acknowledging a file removes it from the New Unapproved Files view but does not change its state.
Approved by Reputation	Indicates whether the file was approved by either its own or its publisher's reputation. (Yes or No).
Assessment	Shows an overall assessment of the risk of this file, combining file threat, file trust, and publisher trust from all available sources.
Certificate Hash	Unique, App Control-generated hash identifier for this file.
Certificate State Reason	For Approved or Banned certificates, how its state was specified. The possible values are: Manual, Trusted Directory, Imported, External (API), Unknown.
Certificate Subject Name	Distinguished name of the subject of the certificate, in this case the signer of the file.
CL Version	For individual files, the configuration list number in which the current global state for this file was defined. Agents at or beyond this CL Version have the correct global state for the file.
Ever Blocked	Indicates whether this file was ever blocked by a App Control rule.
Excluded from Inventory	Indicates whether this file is excluded from the Files on Computers inventory. See "Excluding Tracking of Microsoft Support Files" on page 230.
File Size	Size in bytes of each file.

Field	Description
File State	The approval/ban state of the file hash (Unapproved, Approved, Banned, Approved by Policy or Banned by Policy). The effective “Global State” of a file combines File State and Publisher State. You can change File State using the Action menu on any of the tables on the Files page or any of the details pages for files. On details pages, you can edit an existing approval or ban.
File State Reason	For Approved or Banned file hashes, how its state was specified. The possible values are: Manual, Trusted Directory, Reputation, Imported, External (API), Unknown.
ID	Unique numeric ID for this file.
Initialized	Indicates whether this file was present during agent initialization (Yes or No).
Installed Program	Full package or application name of the installed program (if any) with which this file is associated. Platform Note: Only Windows files are identified as Installed Programs.
Marked as Installer	Indicates whether a file not identified by App Control as an installer has been marked as in installer by a console user. Yes – File was marked as an installer by a user. No – File was not marked as an installer by a user (although it might have been identified by App Control as an installer).
Publisher or Company	Publisher (if available) or company (if available and there is no publisher information) for the file.
Trusted Package	Indicates whether this file is part of a trusted package. (Yes or No). A trusted package is a common source or installer located in a Trusted Directory. Platform Note: Only Windows files can be in a trusted package.
Unified Server Source	If a unified rule has been applied to this file, the name of the unified management server that created or copied the rule.

File Instance Details Page

The File Instance Details page shows information about a file instance on a computer plus some of the global file information you see on the File Details page. In any table showing file instances – for example, the Files on Computers page or Find File Results – you click the View Details button to open the File Instance Details page.

File Instance Details

Details for file on computer: WORKGROUP\APC-79895

File Name:	mpengine.dll
Date Created:	Jun 29 2020 09:08:34 PM
File Path:	c:\programdata\microsoft\windows defender\definition updates\{09fcb399-43c8-4d93-a3d9-7d41adffe5cf}\
Computer:	WORKGROUP\APC-79895
Platform:	Windows
User Name:	(none)
Local State:	Approved
Local State Details:	Locally Approved
Detached Publisher:	(none)
Executed:	Yes
Present At Initialization:	Yes
Top-Level File:	No
Deleted:	No
Root File Name:	(none)

General

First Seen Name:	mpengine.dll
First Seen Date:	Jun 29 2020 07:43:08 PM
Last Updated:	Jun 29 2020 07:43:08 PM
First Seen Path:	c:\programdata\microsoft\windows defender\definition updates\{09b9ff9d-4b4e-4359-890a-6ee11af3c64f}\
First Seen Computer:	WORKGROUP\APC-79895
First Seen Platform:	Windows
Extension:	dll
Global State:	Unapproved
Global State Details:	File is unapproved, Publisher is unapproved, Certificate is Unapproved
Flags:	(none)
Installer / Updater:	No
File Prevalence:	File exists on 1 computer

File Properties

Publisher:	Microsoft Windows
Publisher State:	Unapproved
Certificate:	Microsoft Windows Microsoft Corporation Redmond Washington US
Certificate Type:	Embedded Signer
Certificate Global State:	Unapproved
Company:	Microsoft Corporation
Product Name:	Microsoft Malware Protection
Product Version:	1.1.17200.2
File Size:	15,223,192 bytes
Description:	Microsoft Malware Protection Engine
File Type:	Application
SHA-256:	2090F1A6B417D0F72603B8BF6777A4CB043A2D53D8B1C04928AD52D22B6664A(
MD5:	72B64218DCC56E5EC239F2636F2D848E
SHA-1:	43A4B8FE96C5BF5C861CDECBF19EE2C5DA7BB7E

External Analysis Results

Check Point:	
FireEye:	
Microsoft SCEP:	
Palo Alto Networks WildFire:	

History

Jun 29 2020 07:43:08 PM	The file appeared on WORKGROUP\APC-79895 during initialization
--------------------------------	--

Many File Instance Details fields are identical to those on the File Details page ([Table 26](#)) and you can take many of the same actions from the File Instance Details page. [Table 27](#) shows the additional fields available on the File Instance Details page and Files on

Computers table. On the details page, these appear in the top panel, which is labeled **Details for file on computer:** <computername>.

Table 27: Additional Fields: File Instance Details and Files on Computers

Field	Description
File Instance Details: File on computer panel	
File Name	File name of this instance.
Date Created	Exact time this instance was created in its current location, displayed in the following format: MM DD YYYY hh:mm:ss(AM/PM).
File Path	Path of the this file instance.
Computer	Name of the computer this instance is on.
Platform	Platform (Windows, Mac) of the system the instance is on.
User Name	Name of the user logged in when this file was created.
Local State	The local state of the file instance (Unapproved, Approved, Banned, Deleted). If the local state is Unapproved, you can choose Approve Locally on the Actions menu. If it is Approved, you can Remove Local Approval . If it is Banned, you cannot change it.
Local State Details	File-state metadata for use by Carbon Black Support engineers. If necessary, your support representative may ask you to report this information. See Table 33 for details.
Detached Certificate Subject Name	If this file did not have its own certificate but was indirectly signed via a “detached certificate,” this field appears and shows the subject name from the certificate.
Detached Publisher	If this file did not have its own certificate but was indirectly signed via a “detached certificate,” this field appears and shows the name of the publisher. Some publishers distribute updates as collections of unsigned files with a <i>catalog</i> that contains hashes of all indirectly signed files and is itself signed. App Control can use these catalogs to verify publishers and allow publisher-based approval of files signed in this way.
Detached Publisher State	(If there is a detached publisher) These options are the same as for Publisher State: Approved, Approved by Policy, Banned, Banned by Policy, Unapproved.
Executed	Indicates whether this file instance has been executed or not.
Present at Initialization	Indicates whether this file instance was among the files present on the computer when the App Control Agent was installed, or whether it appeared after installation.

Field	Description
Top-Level File (in Details)	Indicates whether the file is a top-level file; that is, one that was not installed by or copied from another file. Platform Note: On Windows systems, files that were discovered during initialization can be later assigned top-level status if they are discovered to be installers.
Top-Level (in Files on Computers)	
Deleted	Indicates whether this file instance has been deleted from the computer it was on. This is a temporary state immediately after file deletion and before it is removed from the database for this App Control Server.
Root File Name	File that wrote the current file. If this is a top-level file, there is no root file and the name is <i>(none)</i> .
Fields in table only	
Assessment	Shows an overall assessment of the risk of this file, combining file threat, file trust, and publisher trust from all available sources.
Computer Tag	For the computer on which the file appears, displays the optional Computer Tag if provided.
IP Address	The IP address of the computer on which the file appears.
Operating System	The operating system of the computer on which the file appears.
Policy	The App Control security policy of the computer on which the file appears.

Menus on the File Details and File Instance Details Pages

The File Details page includes three menus to the right of the file information: a Related Views menu, an Actions menu, and an Advanced menu. The File Catalog and Files on Computers tabs have an Action menu in the upper left above the table. [Table 28, “Menus on File Tables and Details Pages,”](#) shows the available choices on file page menus. Note that some choices are available only for certain file states.

The File Instance Details page includes three menus to the right of the file information: a Related Views menu, an Actions menu, and an Advanced menu. It is similar to the File Details page menu, except that includes options for local approval. [Table 28, “Menus on File Tables and Details Pages,”](#) shows the available choices on file page menus.

Notes

- Some menu choices are available only for certain file states.
- Many of these commands are also available on the Events page Action menu when the view includes file-related events.
- If you have enabled Unified Management, additional commands appear to allow you to manage files on multiple servers. See [Chapter 27, “Unified Management of Multiple Servers,”](#) for more details.

Table 28: Menus on File Tables and Details Pages

Menu Choice	File Catalog	Files on Computers	File Details	File Instance Details
Related Views menu				
All File Instances			X	X
File Events			X	X
Carbon Black EDR Details			X	X
Computers with this file			X	X
Computers without this file			X	X
Actions menu				
Approve Locally		X	X	X
Remove Local Approval		X		X
Approve Globally	X	X	X	X
Ban Globally	X	X	X	X
Approve by Policy	X	X	X	X
Ban by Policy	X	X	X	X
Edit Global Approval/Ban Edit Approval/Ban by Policy			X	X
Remove Approval or Ban	X	X	X	X
Acknowledge	X			
Find computers with at least one of the selected files			X	X
Find computers with all of the selected files			X	X
Find computers missing at least one of the selected files			X	X
Find computers missing all of the selected files			X	X
Add/Edit Meter			X	X
Add/Edit Alert			X	X
Advanced menu				
View Carbon Black File Reputation Data	X	X	X	X
Enable/Disable Reputation for this File			X	X
Mark as Installer/Not Installer			X	X
External Pages menu				
File Analytics			X	X

Summary of File Views

The previous sections provided details of the main views of file information in the console. [Table 29](#) summarizes how to “drill down” for access to particular views of this information.

Table 29: File Views and File Details in the App Control Console

To view...	...do this
A table of all unique <i>top-level</i> files (files not installed by another file) discovered on computers managed by your App Control Server.	<p>Go to Assets > Files, click on the File Catalog tab, and make sure the <i>Show individual files</i> box is <i>not</i> checked.</p> <hr/> <p>Notes: Top-level files are files that do not have an associated installer, or whose installer is unknown. If a top-level file <i>is</i> an installer, its name shows as a highlighted link to its associated files.</p>
A table of all unique individual files discovered on computers managed by your App Control Server.	<p>Click on the File Catalog tab, and check the <i>Show individual files</i> box.</p> <hr/> <p>Notes: This view shows both files installed by other files and top-level files. Names of known installers are highlighted.</p> <p>Important: There can be millions of unique files discovered by the App Control Server, and this view can cause performance issues on underpowered servers.</p>
The global file details for one unique file.	<p>Click on the File Catalog tab, and click the View Details button next to the file for which you want details.</p>
A table of all files on all computers managed by your App Control Server that are associated with (usually meaning installed by) one top-level file:	<p>Click on the File Catalog tab, make sure the <i>Show individual files</i> box is not checked, then click the name of the file for which you want a list of associated files.</p> <hr/> <p>Notes: This is an aggregate list of associated files, not based on installations seen on one particular computer. For example, if installer X was seen installing files A and B on one computer and installing files B and C on another computer, all installed files (A, B and C) would be listed in the File Group Details page of installer X.</p> <p>For details on any file in the table, click the View Details button next to it.</p> <p>Platform Note: In this release, only files on Windows computers are grouped by installer.</p>

To view...	...do this
<p>A table of all <i>top-level</i> file instances (not installed by another file) on all computers managed by your App Control Server:</p>	<p>Click on the Files on Computers tab, and make sure the <i>Show individual files</i> box is <i>not</i> checked.</p> <p>Notes: Top-level files are files that do not have an associated installer, or whose installer is unknown. If a top-level file <i>is</i> an installer, its name shows as a highlighted link to its associated files.</p> <p>This table view also includes an entry named <Initialization files> for each agent, which is a grouping of the files found on the computer at the time the agent was installed.</p>
<p>A table of all file instances found on one computer at <i>initialization</i>, which occurs either when the agent is initially installed or when a disabled agent is re-enabled:</p>	<p>Click on the Files on Computers tab, and make sure the <i>Show individual files</i> box is <i>not</i> checked. Then click on <Initialized files> in the row containing the name of the computer you are interested in.</p>
<p>A table of <i>all</i> individual file instances on all computers managed by your App Control Server:</p>	<p>Click on the Files on Computers tab, and check the <i>Show individual files</i> box.</p> <p>Notes: This view shows both top-level and “individual” files that were installed by them on an agent-managed computer. Top-level files that have been analyzed by the agent to determine their contents show as highlighted links.</p> <p>Important: Avoid checking this box unnecessarily, especially if you have a large number of agent-managed computers. The number of individual files could number in the tens or hundreds of millions. Attempting to load a list of this many files can cause the App Control Server to time out.</p>
<p>The details for one file instance on one computer.</p>	<p>Click on the Files on Computers tab, and click on the View Details button next to the file instance for which you want details.</p> <p>Notes: Opens the File Instance Details page. Shows both local state and other information about this instance and global details for the file.</p> <p>Top-level files can still appear in Files on Computers tables after they are no longer present. Clicking View Details for a removed file no longer present on a computer will show global details only.</p>
<p>A table of all files on one computer that are associated with one top-level file:</p>	<p>Click on the Files on Computers tab, and click on the name of the highlighted top-level file instance for which you want a list of associated files.</p> <p>Notes: Shows the results of a Find Files search for all files on <i>the named computer</i> in the rows for the files whose name you clicked on.</p> <p>For details on any file in the table, click the View Details button next to it.</p>

Global File State

Files in the File Catalog tab on the Files page have the following high-level states:

- **File State** indicates the approval/ban state of the file itself.
- **Publisher State** is the approval state of the file's publisher (if known). The only choices are Approved, Approved by Policy, and Unapproved.
- **Global State** combines File State and Publisher State to determine how the file is to be treated on agent-managed computers. The File State and Global State are the same except when:
 - Publisher State is *not* Unapproved, *and*
 - File State is *not* approved or banned in the same policies as the publisher.

Global State cannot be modified directly. [Table 30](#) shows the possible Global States.

Table 30: Global State (for files) cataloged by a App Control Server

State	Description
Approved	Allowed to execute on all computers.
Banned	Banned by hash, and not allowed to execute on any computer running in Control mode.
Approved by Policy	Allowed to execute on computers in one or more policies.
Banned by Policy	Banned by hash from execution on computers in one or more policies (in Control mode).
Unapproved	Not Approved or Banned (globally or by policy). App Control blocks or permits execution of an unapproved file based on the Enforcement Level of the Policy of the computer attempting the execution.
Mixed	Effective state varies by policy because File State is Banned for some policies but the Publisher State is Approved for some or all policies.

Flags

Global State is the effective App Control classification of each unique file in the File Catalog. It is a combination of the File State and the Publisher State for the file. *Flags* are primarily for use by Carbon Black Support, but you might find them useful in determining how a file is being labeled or handled in your App Control environment.

Table 31: File Flags

Flag	Description
Report Only Ban	File was identified by a console user so that attempts to execute it are reported as if they would have been banned, but it is not blocked from execution.
Installer	File was identified as an installer by App Control and is allowed to execute. If executable files are written out by it, they are locally approved. Platform Note: For this release, on Mac computers, only files associated with the native Mac updater (i.e., .pkg files) are identified as installers.
Installer (Override)	File was identified as <i>not</i> being an installer by App Control, but a console user changed it to “installer”. If it is allowed to execute, the executable files it writes out are locally approved.
Not installer (Override)	File was identified as an <i>installer</i> by App Control, but a console account user changed its installer status to “Not installer”.

Local File State

Files that are globally Banned or Approved have the same local and global state. Files with a Global State of Unapproved may have different Local States. In particular, you can locally approve a file by a variety of methods, as long as that file was not globally banned. You can view local file state on the Files on Computers tab of the Files page.

Table 32: Local State

State	Description
Approved	This instance of the file is approved for execution. Local approval can be due to approval by name or hash for all computers in a policy or all computers managed by this App Control Server. It also could be due to a global approval method, a change in Enforcement Level, or an explicit Local Approval of this single file instance. Locally approved files can have a <i>global state</i> of Unapproved or Approved, but not Banned.
Banned	This instance of the file is banned from execution. A file that has a local state of banned might be banned on all computers in certain policies or all computers managed by this App Control Server. Note that banning a file by name does not change its local state.
Unapproved	This instance of the file has not been approved or banned. Its execution is blocked or permitted based on the Enforcement Level of the computer it is on.
Deleted	This instance of the file has been deleted, but the record of it still exists in the database for this App Control Server.

Local State Details

Local State is the App Control classification of a particular instance of a file on a particular computer. This information is primarily for use by Carbon Black Support, but you might find it useful in determining why a file was assigned its top-level Local State.

Table 33: Local (File) State Details

State	Description
Approved	Approval state on the local computer for files that are globally approved in the File Catalog.
Approved (Not Persisted)	Approval state on the local computer for files approved by certain pre-version-6.0 methods but are not globally approved in the File Catalog. If you delete a file in this state, new instances would not necessarily be locally approved.
Approved as Installer	Approval state for top-level installers (in Windows) that indicates that the installer and the files it contains have been hashed, analyzed, and globally approved by App Control. When users execute these files, the App Control Agent permits them to run as globally approved files. This state is not common and unnecessary for local approval of files generated by an installer.
Approved as Installer (Top Level)	Approval state for top-level installers. The installer was globally approved and when executed, the files it generates are locally approved. Platform Note: On Mac computers, only files associated with the native Mac updater (i.e., .pkg files) are identified as installers.
Banned	Files with specified hash are not allowed to execute on the computers specified (all computers or by policy).
Banned (Report Only)	Test file state for files that are to be banned by hash. App Control permits files that are banned but in Report-Only to execute but records a “would have blocked” message in the event log to show how the file would have been handled if the ban were active.
Locally Approved	File is approved to run on the local computer but unapproved (globally or for the current policy) in the File Catalog. Files can be locally approved so that they can be installed on one computer without approving them for any other computer running the agent.
Locally Approved (Auto)	File is approved to run on the local computer because it was written by a trusted installer or updater. Other than the source of its approval, this is the same as Locally Approved.

State	Description
Unapproved	File appeared after agent initialization and has not been approved. Depending on Enforcement Level on each computer, the agent either blocks the file or permits its execution. These files might become locally approved if a computer transitions from Low (or no) Enforcement to Medium or High, depending upon policy settings. Files are assigned Unapproved local state details if the first local instance was found when the Enforcement Level was Low (Monitor Unapproved) or None (Visibility Only). See “Automatic Local Approval on Enforcement Level Change” on page 290 for details.
Unapproved (Persisted)	File appeared after agent initialization and has not been approved. Unapproved (Persisted) files do not become locally approved when a computer changes from Low or None (Visibility) Enforcement to High or Medium Enforcement. Files are assigned Unapproved (Persisted) local state details if the first local instance was found when the machine was in High or Medium Enforcement Level.

Publisher Information

The Publishers tab on the Software Rules page shows file publishers discovered on computers running the App Control Agent in your organization. It also shows any publishers that have been added manually to the File Catalog for your App Control Server. This page includes an Action menu that allows you to approve or ban a publisher, remove approvals or bans, and acknowledge a publisher to indicate that you have reviewed it already. These actions are described in “[Approving or Banning by Publisher](#)” on page 281.

To view the list of discovered or added publishers:

1. On the console menu, choose **Rules > Software Rules**. The Software Rules page appears.
2. Click the **Publishers** tab. All publishers of signed software installed on agent-managed computers reporting to your server, plus any publishers you manually added using certificates, appear in the Publishers table:

Software Rules

Updaters Rapid Configs **Publishers** Users Directories Files Custom Memory Registry Scripts

Saved Views: (none) Add Group By: State

Show Filters Show Columns Export to CSV Refresh Table

Action Add Publisher Search: Automatically apply

Select 23	Name	State	Date Approved or Banned
State: Unapproved		20 item(s)	
<input type="checkbox"/>	Microsoft Windows	Unapproved	
<input type="checkbox"/>	Microsoft Windows Publisher	Unapproved	
<input type="checkbox"/>	VMware, Inc.	Unapproved	
<input type="checkbox"/>	Microsoft Windows Hardware Abstraction Layer Publisher	Unapproved	
<input type="checkbox"/>	Microsoft Windows 3rd party Component	Unapproved	

You can view a Publisher Details page for any publisher shown in the Publishers table by clicking on the View Details button next to the publisher name. In addition to details (see [Table 34](#)), the Publisher Details page has shortcuts with which you can Approve or Remove Approval for the publisher. The Related Views menu also includes a command that shows all files from the publisher as well as commands that show computers where the approval state for this publisher is up-to-date.

To view complete details for one publisher:

1. On the console menu, choose **Rules > Software Rules**. The Software Rules page appears.
2. Click the **Publishers** tab. All publishers of signed software installed on agent-managed computers on your network appear in the Publishers table.
3. From the table of publishers, locate the publisher you want to authorize and click on the View Details button. The Publisher Details page opens.

Publisher Details
?

General

Publisher Name: Google LLC

State: Approved

Acknowledged: No

Trust: Not Trusted (Unknown)

Description:

Rule Applies To: All Current and Future Policies
 Selected policies

Platforms: All platforms
 Selected platforms

▶ All Certificates For This Publisher (click to expand)

History

Date First Seen: Jun 29 2020 07:51:57 PM

Platform First Seen: Windows

Computer First Seen: WORKGROUP\APC-79895

Date Approved: Jun 30 2020 01:18:43 PM

Approved By: admin

CL Version: 1151

Save & Exit
 Save
 Cancel

Related Views

All files signed by this publisher

All Computers that have received this rule

All Computers that have not yet received this rule

Table 34: Publisher Details

Field	Description
General panel	
Publisher Name (in Details) Name (in Table)	The name of this publisher as it appears in its certificate.
State	Approved, Unapproved, or Banned.
Enable reputation approvals... (in Details) Reputation Enabled (in Table)	This checkbox appears if you have reputation approvals enabled. <i>Enable reputation approvals for this publisher</i> is checked by default and allows this publisher to be approved by reputation. Removing the check disables reputation approvals for this publisher, but if reputation approvals were already globally enabled, removal only affects files first seen after the change.
Acknowledged	You can Acknowledge a publisher, which indicates that you have reviewed it. This can help distinguish new publishers from those you already know about.
Trust	This field appears if you have Carbon Black File Reputation enabled. Shows the trust rating for this publisher, which can be High, Medium, Low, or Not Trusted.
Description	Optional additional description of this publisher and its state.
Rule Applies To (in Details) Policy (in Table)	For publishers that do not have reputation approval enabled, you can apply the publisher state to computers in all policies or only to those in some policies.

Field	Description
Platforms (in Details) Approved Platforms (in Table)	You can apply the publisher state to computers in all platforms or choose a specific platform (Windows, Mac). Platform Note: Publisher approvals work on Windows only.
All Certificates For This Publisher panel	
	<p>The Publisher Details page includes a panel entitled All Certificates for This Publisher. Because this panel has the potential to be long, it can be collapsed and expanded on the page by clicking the panel name.</p> <p>The panel shows all leaf certificates for this publisher, and all root and intermediate certificates associated with these leaf certificates. See Chapter 11, “Managing File-Signing Certificates,” for more on certificate details.</p>
History panel	
Date First Seen (in Details) First Seen Date (in Table)	When this publisher was first seen on a agent-managed computer reporting to your server.
Platform First Seen (in Details) First Seen Platform (in Table)	The platform (Mac or Windows) of the computer on which this publisher was first reported to your server.
Computer First Seen (in Details) First Seen Computer (in Table)	The computer on which this publisher was first reported to your server.
Date Approved	If the publisher is approved, when that approval was made.
Approved By/ Banned By	The console user who approved or banned the publisher. Publishers approved by reputation may show “System” in this field.
Date Acknowledged	If the publisher has been acknowledged, when it was acknowledged.
Acknowledged by	If the publisher has been acknowledged, the console user that acknowledged this publisher.
CL Version	The version of App Control rules containing current publisher state. This helps determine whether an agent has the rule.

Application Information

App Control agents collect information about applications installed on agent systems and report it to the server, which displays the information on the Applications page. This page is similar to the Files page. It includes an Application Catalog tab, which shows each unique, identifiable application on agent systems reporting to your server, and an Applications on Computers tab, which shows each identified instance of these applications on each computer.

Notes

- The Applications page shows Windows applications only.
- Applications are shown only for agents reporting to the current App Control server; there is no Unified Management for applications.

Most of the information about the application itself is from the metadata of the installer file for the application or the application's registry data. It is similar to the information on the Programs and Features page of the Windows Control Panel. If the installer was an MSI file and that file still exists on a computer, application information is extracted from the installer metadata. For other installer types, or if the MSI installer has been deleted after installation, application information is extracted from the Windows registry.

Other information describes the application in your environment, such as where it is installed and its prevalence. It also includes information about App Control features that could affect the application, such as the policy of the agent on a computer with the application, and whether there is a Rapid Config that (if enabled) would control the behavior of the application.

Application data is not about a specific file or hash but about the application as a whole. Some views on the Files pages can show files associated with an application. See [“File Groups”](#) on page 234 and [Table 25, “Saved Views on the File Catalog tab”](#) on page 226.

Important

Keep in mind that most *file* data in App Control comes from file metadata while much of the *application* data comes from the registry. It may be easier for a malicious actor to spoof this data, making an application appear to be something it is not.

Viewing Application Data

Application data is shown in the tables only. Unlike files, applications do not have separate details pages in App Control. However, as with other tables, you can add columns with additional information using the Show Columns link on the Applications page. For example, you might want to include the Upgrade URL for each application in the table. [Table 35](#) shows the information available on both tabs of the Applications page. Not all fields will contain information for all applications.

To view the Applications page:

1. On the console menu, choose **Assets > Applications**.
2. To view the table of unique applications, click on the **Application Catalog** tab.

3. To view all application instances, click on **Applications on Computers**.

Applications

Application Catalog
Applications on Computers

Saved Views:

(none) ▾
Delete
Save

Add

Group By:

(none) ▾
Ascending ▾

[Show Filters ▾](#) | [Show Columns ▾](#) | [Export to CSV](#) | [Refresh Table](#)

Showing 26 out of 26 item(s)

	Date Created	Name	Version
Q	Jun 29 2020 07:44:58 PM	Google Chrome	83.0.4103.116
Q	Jun 29 2020 07:44:58 PM	Python 3.8.3 Test Suite (32-bit)	3.8.3150.0
Q	Jun 29 2020 07:44:58 PM	Python 3.8.3 Utility Scripts (32-bit)	3.8.3150.0
Q	Jun 29 2020 07:44:58 PM	Python 3.8.3 Standard Library (32-bit)	3.8.3150.0
Q	Jun 29 2020 07:44:58 PM	Microsoft Visual C++ 2017 Redistributable (x64) - 14.14.26405	14.14.26405.0
Q	Jun 29 2020 07:44:58 PM	Microsoft Visual C++ 2017 x86 Additional Runtime - 14.14.26405	14.14.26405

In addition to viewing application details, you can list all computers with an application.

To list all computers that have an application:


- In the Application Catalog, click on the Find button  on the left side of the row for the application.

Table 35: Applications Page Fields

Field	Description
Fields on Both Tabs	
About URL	The URL for the “about box” of the application, which can show information such as the general purpose, copyright and version of the application; from the registry or installer metadata.
Comments	Comments about the application; from the registry or installer metadata.
License Owner	The company name entered by the user who installed the application; from the registry or installer metadata. Usually empty.
Contact	Contact information for the application; from registry data or installer metadata.
Date Created	On the Applications on Computers tab, the date an instance of the application was first discovered by the agent on its computer. In the Application Catalog, the discovery date of the first seen instance of the application.
Date Modified	On the Applications on Computers tab, the date of the most recent modification of one instance of the application. In the Application Catalog, the date of the most recent modification of this application reported by any agent.

Field	Description
Help Telephone Number	The phone number for support/help for this application; from registry data or installer metadata.
Help URL	The URL for application help; from registry data or installer metadata.
Name	The application name; from registry data or installer metadata.
Package Code	The application package code; from registry data or installer metadata. Note: See https://blogs.msdn.microsoft.com/pusu/2009/06/10/what-are-upgrade-product-and-package-codes-used-for/ for more information about Package, Product, and Upgrade codes.
Product Code	The application product code; from registry data or installer metadata.
Manufacturer	The application publisher; from registry data or installer metadata. Note: The publisher identified in the Manufacturer field is not necessarily the same as the publisher identified on the console Publishers and Certificates pages, which is drawn from certificate metadata, not installer metadata.
Readme	The readme information about the application; from registry data or installer metadata.
Upgrade Code	The upgrade code for this application; from registry data or installer metadata.
Upgrade URL	The URL for application upgrade downloads; from registry data or installer metadata.
Version	The four-part, decimal version number of this product; from registry data or installer metadata.
Version Number	The hexadecimal version number of this application; from registry data or installer metadata.
Windows Installer	Whether this is a Windows Installer (Yes/No).
Fields on the Application Catalog Tab Only	
Computer Count	The number of computers reporting to this server that have this application.
First Seen Computer	The name of the computer on which the application was first seen and reported to the current App Control Server. Click on this name to get the Computer Details page for this computer. If you later delete the first-seen computer from the system, it appears with "(Deleted)" appended to its name and no link.
Rapid Config	The name of any App Control Rapid Configs that applies to this application, whether or not it is enabled.

Field	Description
Fields on the Applications on Computers Tab Only	
Catalog ID	A locally unique, numeric identifier for each application in the catalog on your App Control Server. Begins with '1' and increments by 1.
Computer	The computer reporting to this App Control Server on which an instance of the application was installed.
Estimated Disk Usage	The estimated amount of disk space used by an instance of the application, based on information from the registry. Applications that add content after the installation has completed will use more space than this indicates.
Estimated Install Date	The estimated date the application was installed on this computer in the format provided by the software vendor. This is an estimate because the time zone and format of this date differ from vendor to vendor.
Install Directory	The directory in which the application was installed on this computer. This is either a path or "Default". Note: "Default" appears if an application did not record an "InstallLocation" in the registry when it was installed. If an application did specify an installation directory in the registry, that directory appears in this field, even if it happens to be the same as the default.
Installer	The name and path of the file used to install the application. This field has data only for MSI installers.
Platform	The operating system platform this application can run on.
Policy	The App Control policy of the computer on which this application is installed.
Repair Command Line	The command line for repairing the application; from registry data or installer metadata.
Source	The directory in which the installer that was first used to install the application was run from on each agent computer. This is often a temporary folder. Depending on the application, the Windows installer may need to reference the source file for repair installations, installation or removal of application components, upgrades, or for uninstalling the application.
Uninstall Command Line	The command line for uninstalling the application; from registry data or installer metadata.
Uninstall Key	The registry key for uninstalling the application; from registry data or installer metadata.
Users	The users for whom this application was installed. This will either by "All-Users" or one or more specific users; from the registry.

Chapter 8

Approving and Banning Software

This chapter describes how to approve or ban software using App Control. It includes information about both global and local file approval. Many of the methods for approving and banning software are found on one of the tabs of the Software Rules pages.

In addition to explicit approvals and bans, App Control allows you to define Custom Rules for allowing or blocking file execution or writing at specified locations, and if you choose, by specified users and/or processes. See [Chapter 14, “Custom Software Rules.”](#)

Sections

Topic	Page
What is App Control Software Approval?	261
What are App Control Software Bans?	263
Approving by Updater	266
Approving by Trusted Directory	271
Approving by Trusted User or Group	279
Approving or Banning by Publisher	281
Locally Approving Files	289
File-Specific Rules: Approvals and Bans	301
Approving or Banning Lists of Files	310
Enabling Bans to Stop Running Processes	312

What is App Control Software Approval?

Software approval ensures that users of computers running the App Control Agent can freely install and run *known-good* applications regardless of the App Control security settings and Enforcement Level in effect. App Control supports several complementary methods for approving software on computers. Based on the method(s) you select, installation of approved software may be permitted on all computers, on computers in selected policies, or on individually selected computers.

You can choose the combination of methods that best conforms to your established settings and procedures, especially the software distribution process in place at your site:

- When you need to pre-approve applications to run on all computers (or all computers in selected policies), designate trusted directories, approve specified publishers to allow installation of their applications, or enable certain updaters to update applications automatically.
- When you would like to pre-approve low-threat applications to run on all computers (or all computers in selected policies), enable reputation rules based on the trust level reported by Carbon Black File Reputation for specific files and publishers.
- When you discover an individual file or installer that you want to allow to run on all computers or all computers in selected policies, create a File Approval rule.
- When you have a list of hashes for files you want to approve, you can create approvals for the entire list in a single operation.
- When you need to approve software for installation on selected individual computers, either designate trusted users (or groups) to perform installations, or choose a local approval method.
- When you have a special need for a rule to allow installation or execution of files in particular locations, or by particular users or processes, create a Custom Rule.

Tip

At all Enforcement Levels except for High, users can install unapproved software. Although not required, Carbon Black recommends approving (or at least Acknowledging) widely used software even if you plan to run at Low Enforcement Level. Approval reduces the number of files with the unapproved status, which can enable you to focus on files that are of potential concern. For example, approving known-good files generally reduces the size and increases the readability of Baseline Drift reports.

Similarly, computers operating in Visibility mode can run *any* software, regardless of its approval state. Even if you are running all your computers in Visibility mode, you might want to approve known-good files to reduce the amount of event data collected about those files. This also helps prepare you for possible transition of some or all computers into High or Medium Enforcement Level in the future.

Based on your internal standards and procedures, and on the required scope of the approval (network-wide or computer-specific), you can choose to approve files in any of the ways shown in the table that follows.

Table 36: File Approval Methods

Approval Method	Software Is Approved for	When to Use
Approving by Trusted Directory	All computers (global)	When you have a trusted, secure server (e.g., for software deployment) on which to create an authorized approval directory.
Approving by Trusted User or Group	Installation computer only (local)	When you want to give unlimited installation privileges to a Windows user account or all users in a Windows or AD group. Trusted users are allowed to install on any computer on which they log in with their credentials.
Approving or Banning by Publisher	Installation computer only (local), but can be installed on demand by any computer	When you want to approve all software from a vendor for which App Control can confirm a valid digital certificate. You also can approve or ban certificates that identify a publisher, and this affects file state. See “Using Certificates for Enforcement” on page 348.
Approving by Publisher Reputation (see Chapter 10, “Reputation Approval Rules”)	Installation computer only (local), but can be installed on demand on any computer	When you want to automatically approve all software from all publishers considered trustworthy by Carbon Black File Reputation.
Approving by Updater	Installation computer only (local), but can be installed on demand on any computer	When you want to permit installation of application updates as they become available for download via specified application update programs.
Automatic Local Approval on Enforcement Level Change	Installation computer only (local)	When you want to locally approve <i>unapproved files found while in Low enforcement or higher</i> when you move the computer from a less secure Enforcement Level to either Medium or High.
Moving Computers to Local Approval Mode	Installation computer only (local)	When you want to permit users on computers in High Enforcement policies to install software. Local approval occurs when a user installs an unapproved file while in this mode.
Locally Approving All Unapproved Files on a Computer	Installation computer only (local)	When you want to locally approve all existing unapproved files on a specific computer.
Locally Approving Individual Files	Installation computer only (local)	When you want to select specific files on a computer for local approval. You can locally approve files, or remove local approval.
File Approval Rule	Approved for all computers or those in selected policies	When you want to ensure that a known-good application can run on any computer, approve it by hash.

Approval Method	Software Is Approved for	When to Use
Approving by File Reputation (see Chapter 10 , “Reputation Approval Rules”)	Approved for all computers or those in selected policies	When you want to automatically approve (by hash) all software that Carbon Black File Reputation considers trustworthy.
Approving by Event Rule (see Chapter 19 , “Event Rules”)	Varies by rule	When you want to automatically approve a file, either locally or globally, when it is included in a reported event.

Platform Considerations for Rule Specifications

Many App Control rules involve specification of a file name and/or path, or other manually entered information such as a user, group, or computer name. On both Mac and Windows computers, file names, paths, and user names in rules normally are *not* case sensitive. On Linux computers, file and user names in rules normally *are* case sensitive; for example, if you create a rule to ban `/temp/myfile.exe`, it will not block the files `MyFile.exe` or `/TEMP/myfile.exe`. There are two additional factors to consider in determining how case sensitivity works for rule parameters:

- Regardless of the general case-sensitivity rule for an operating system, it is actually the file system that determines case sensitivity. If a case-sensitive file system is attached to a computer whose standard file system is not case-sensitive, App Control rules will be case sensitive, and vice versa. Keep this in mind when you connect an external drive or mount a network file system to a App Control-managed computer.
- The case of text entered in rule fields is preserved even if it is not relevant in its current use. This might be significant if you copy information to a place in which it applies to a different platform.

When you enter a path, be sure to use the correct directory delimiters for the platform it applies to, and to use only characters and formats legal for paths in the chosen platform. The App Control Server does not convert paths between platforms (e.g., ‘\’ to ‘/’), although it will display a warning in some cases if the delimiter is known to be a mismatch for the platform.

What are App Control Software Bans?

File bans are rules that block specific files from executing on computers running the App Control Agent, based on the agent Enforcement Level (see [Table 37](#)). You can ban files reported by your App Control Agents in the course of day-to-day operations, and you also can preemptively ban files not yet seen on your computers but for which you have obtained information from third-party sources. App Control supports bans by file name or hash. Bans can affect all agents running in Control mode or be targeted to computers in selected policies only. You also can configure App Control to terminate processes already running when you ban their file image.

As the table shows, file bans do not prevent software from running on computers operating in Visibility mode. However, even in Visibility, a ban will produce an event that you can use to monitor how often the banned file is run. Banning undesirable files while in Visibility mode also helps you prepare for a transition into full Control mode in the future.

Table 37: How File Bans affect File Execution, by Enforcement Level

Policy Settings	Enforcement Levels				
Active Bans	None (Agent Disabled)	None (Visibility Only)	Low	Medium	High
Banned files (by hash or name)	Off/Permit	Permit & Report	Block	Block	Block

When you ban specific files by name or hash, the bans appear as rules on the Software Rules page Files tab. One fundamental decision about how you ban a file is whether you ban it by name or by hash. [Table 38](#) describes the differences between the two.

Table 38: Name vs. Hash Bans

Ban Type	Description
File Name Ban	<p>Block execution of the named file everywhere (if you enter only the file name) or at specified locations (if you enter a path), and on all computers or computers in selected policies. File name bans do not change the Global State of a file, but assure that all instances of files by the specified name are locally banned wherever they appear.</p> <p>Be careful not to ban a file required for system or application operation, especially when you specify paths using the (*) wildcard character.</p> <p>As a precaution, you can execute file-name bans in Report-Only state to test the effects of the ban. Ban (Report Only) bans remain unenforced until you change them to a blocking Ban.</p> <p>When you search by state for a file that is banned by both name and by hash, the file appears in the list of files in the Banned state but not in files with Local State Details of Banned by Name.</p> <p>Platform Note: Each file name ban is specific to one platform only. If you enter a path, be sure to use the correct directory delimiters, and to use only characters and formats legal for paths in the chosen platform.</p>
Hash Ban	<p>Block execution of the specified hash in any location on all computers or on computers in selected policies. Hash bans are not platform-specific.</p> <p>Although you can copy and paste hashes from external sources, it is easier to ban hashes discovered by an agent directly from console pages that list files. You can create a Ban directly from most console pages that show a hash. Bans initiated from these pages automatically direct you to the Add File Rule page, fill in the hash for you, set the Type as <i>Ban</i>, and allow you to modify other ban properties before creating the ban.</p>

File Ban Options

In addition to bans applied directly to specific files to prevent future execution, App Control provides many other methods and options. The following list summarizes options for banning software:

- When you want to prevent certain software from running on all computers or all computers in selected policies, create a File Ban rule for each file, which blocks it on all computers running in Control mode (or if you are running in High Enforcement, simply do not approve it). See [“File-Specific Rules: Approvals and Bans”](#) on page 301 for details on how to create these bans.
- When you have a list of hashes for unwanted files you want to ban, you can create bans for the entire list in a single operation. See [“Approving or Banning Lists of Files”](#) on page 310 for details on how to create these bans.
- When you want to ban all files from a particular publisher, ban the publisher. See [“Approving or Banning by Publisher”](#) on page 281 for more details. You can further fine-tune publisher bans by banning a specific certificate from a publisher. See [“Using Certificates for Enforcement”](#) on page 348 for more details.
- When you have a special need for a rule to block or allow installation or execution of files in particular locations, or by particular users or processes, create a Custom Rule that blocks execution – this is not a ban but can act like a ban when conditions match its criteria. See [Chapter 14, “Custom Software Rules,”](#) for more details.
- When you want to ban currently running processes with banned images in addition to future attempts to execute a file, configure Policies to do so. See [“Enabling Bans to Stop Running Processes”](#) on page 312 for more details.
- When you want to ban files when they are referenced in certain events, including malware reports from external notifications, create an Event Rule. See [Chapter 19, “Event Rules,”](#) for more details.

Note

Beginning with version 8.1.0, you can use CB Console commands to delete files on endpoints. See [Chapter 9, “Deleting Files,”](#) for details.

Approving by Updater

Updater Approval Rules permit users of computers under High Enforcement protection to install application updates from approved sources as they become available for download. You can approve updater programs for commonly used enterprise applications, including anti-virus, anti-spyware, personal firewall, and desktop productivity programs. All computers can run approved updaters, but applications installed by these updaters via the Web are locally approved by the App Control Agent for use on the installation computer only.

Platform Note: Updaters are platform-specific. Most of the updaters are disabled by default but can be enabled. Mac and Linux are listed as separate updaters and disabled by default to allow greater control of your environment. The Mac App Store Downloads updater is one of the few enabled by default, but you may disable it.

Keep in mind that enabling a product-specific updater approves only the *upgrade procedure* for that product, not the application's full installation package.

As new applications or new application versions are introduced, and old products or versions become obsolete, the list of updaters you need may change. The list of available updaters is refreshed in the following ways:

- When you install a new version of App Control, the updaters list is refreshed to add any new updaters, delete any obsolete updaters, and make any necessary modifications to existing updaters.
- To keep your updaters current, you can allow automatic updating of your updaters by the Carbon Black File Reputation cloud; this feature is enabled by default when Carbon Black File Reputation integration is enabled.
- For update programs currently not supported, you can make a request on the Carbon Black User Exchange. If approved and made available, the new updater can be downloaded automatically through the Carbon Black File Reputation.

Note

To avoid unwanted file blocking, before you install any App Control Agents, it is best to enable any supported updaters for any applications your organization runs. However, if a product whose updater is not enabled attempts to modify files, and this results in the application being blocked, you can use global or local approval methods to manually approve the blocked files.

You can view the complete list of updaters available on your server by opening the Updaters tab of the Software Rules page on the console. In addition to supported updaters for this release, this page might show a manually added updater or, if you have upgraded from a previous version of Bit9 Platform or Parity, older updaters you have enabled in the past.

[Table 39](#) provides information about updaters whose names might not make their purpose obvious or that require special implementation notes. If you do not have access to the console and need a complete list of supported updaters, contact Carbon Black Support.

Table 39: Updater Notes

Updater	Platform	Description
Note: This table describes only updaters requiring addition explanation. For a complete list of updaters, see the Software Rules/Updaters page in your console.		
Adobe Application Manager	Windows	Allows updates of products <i>managed by</i> the Adobe Application Manager.
Adobe Products Not Listed	Windows	Allows automatic approval of updates to certain Adobe products for which a specific App Control updater is not shown.
Allow Printer Installations	Windows	Allows a print server to automatically install a printer driver not currently on an agent computer (Windows 2003 and later). This updater should not be enabled as a means to allow installation of drivers for locally attached printers.
Carbon Black EDR	Mac	Allows updates to the Carbon Black EDR sensor on endpoints running OS X.
CSC.exe Temporary Files - Do Not Report	Windows	This updater significantly reduces the number of new file reports on the server when the Microsoft Visual C# Compiler (CSC.exe) creates or modifies DLLs in locations dedicated to temporary files. You may still approve or ban files at these locations when this “updater” is enabled, and you can disable it if you prefer to see all temporary file traffic from this process.
Java	Windows	Allows updates to the Java Virtual Machine and updates that install or update add-ons (search bars or third-party applications, etc.) included in some versions of Java. This is equivalent to the Java and Bundled Software updater from previous releases.
Mac System Updates	Mac	Allows updates to the OS X operating system. Note: In pre-7.2.1 releases, Mac System Updates were automatically allowed and there was no updater listed. You can now control whether these updates are allowed.
Microsoft .NET Framework	Windows	Allows the .NET just-in-time compiler to run. It must be enabled if you run any applications that require.NET. Although Windows Update provides updates for both Windows Defender and Microsoft .NET , successful installation of updates for either of these products requires that you trust their specific updater in addition to Windows Update.
Microsoft Office 2013	Windows	Allows updates based on Microsoft's Click-to-Run streaming technology. If you used the MSI installer for Office and did not enable Click-to-Run, Office updates will be provided by Windows Update and this updater does not need to be enabled.

Updater	Platform	Description
Red Hat Prelinking	Linux	Carbon Black recommends disabling Prelinking on RedHat and CentOS computers before installing agents. Prelinking has negative impacts on performance and App Control features (see the Release Notes). However, if you must enable Prelinking on your RedHat and CentOS systems, enable the RedHat Prelinking updater before installing agents.
Red Hat Software Update	Linux	Allows automatic updates to supported RedHat and CentOS operating systems.
Symantec Endpoint Protection for Mac	Mac	Enable the Symantec Endpoint Protection for Mac updater if SEP is run in your environment. It allows SEP updates and improves performance on file operations. Use the SEP Auto Protect Preferences Pane to configure SEP to include the following endpoint SafeZone: /Library/Application Support/com.bit9.Agent
Windows 8, 10 and Server 2012 Updates	Windows	Allows updates for these platforms on pre-7.0.1-Patch 11 agents. These updates are enabled automatically beginning with version 7.2.0 agents, and for 7.0.1, on Patch 11 agents and later.
Windows Defender	Windows	Although Windows Update provides updates for both Windows Defender and Microsoft .NET , successful installation of updates for either of these products requires that you trust their specific updater in addition to Windows Update. Note: Defender is activated by default in Windows 10 unless there is another AV product installed on the system.
Windows Update (for pre-6.0.2 agents)	Windows	This updater allows Windows Updates to run on pre-6.0.2 agents. Windows Updates are enabled by default for v6.0.2 and later agents.
Windows Update Temporary Files - Do Not Report	Windows	This updater significantly reduces the number of new file reports on the server when Windows updates are applied. Because the files not reported are in temporary locations and supplied by Microsoft, they should not be of interest for tracking or investigation. You may still approve or ban files at these locations when this “updater” is enabled, and you can disable it if you prefer to see all updater file traffic.

To automatically approve files installed by application updaters:

1. In the console menu, choose **Rules > Software Rules**.
2. On the Software Rules page, click the **Updaters** tab. A table of updater programs for various applications appears, grouped by default by whether they are enabled:

Software Rules

Updaters Rapid Configs Publishers Users Directories Files Custom Memory Registry Scripts

Saved Views: (none) Add Group By: Enabled Descending

Show Filters Show Columns Export to CSV Refresh Table

Action Showing 35 out of 64 item(s) Showing 2 out of 2 group(s)

<input type="checkbox"/> Select 35	Name ▲	Platforms	Enabled	Date Created
Enabled: Yes		10 item(s)		
<input type="checkbox"/>	Allow Printer Installations	Windows	Yes	Jun 29 2020 07:00:41 PM
<input type="checkbox"/>	Apple System Performance	Mac	Yes	Jun 29 2020 07:01:24 PM
<input type="checkbox"/>	Detection of Linux Shutdown sequence	Linux	Yes	Jun 29 2020 07:00:57 PM
<input type="checkbox"/>	Linux Self Protection	Linux	Yes	Jun 29 2020 07:01:26 PM
<input type="checkbox"/>	Linux System Performance	Linux	Yes	Jun 29 2020 07:00:56 PM
<input type="checkbox"/>	Mac App Store Downloads	Mac	Yes	Jun 29 2020 07:00:52 PM
<input type="checkbox"/>	Mac System Updates	Mac	Yes	Jun 29 2020 07:00:55 PM
<input type="checkbox"/>	Microsoft .NET Framework	Windows	Yes	Jun 29 2020 07:00:42 PM
<input type="checkbox"/>	Windows 8, 10 and Server 2012 Updates	Windows	Yes	Jun 29 2020 07:00:50 PM
<input type="checkbox"/>	Windows Update Temporary Files - Do Not Report	Windows	Yes	Jun 29 2020 07:00:50 PM
Enabled: No		54 item(s)		
<input type="checkbox"/>	Adobe Acrobat Reader 10.0	Windows	No	Jun 29 2020 07:00:39 PM
<input type="checkbox"/>	Adobe Acrobat Reader 9.0	Windows	No	Jun 29 2020 07:00:39 PM

3. Check the box on the far left of the row for any currently disabled updaters you want to enable, and then choose **Enable Updaters** on the Action menu. The updaters are enabled, and with the default grouping in effect, they are moved into the *Enabled: Yes* group. Automatic updaters for these applications can now install software on computers running the App Control Agent.

Note

Some software manufacturers include multiple products in the same product family. Verify that the updater you select corresponds to the correct product and version for your application.

4. If you would like Carbon Black File Reputation to keep your updater list current with updater changes, additions, and deletions, leave the “updater updates” option enabled. See [“Automatic Cloud Management of Updaters”](#) on page 270.
5. If an updater you want to include does not appear in the table, you can submit a new updater request to the Carbon Black [User Exchange](#). See [“Automatic Cloud Management of Updaters”](#) on page 270 for more information.
6. To disable updaters, check the box next to the Name of each updater you want to disable and then choose **Disable Updaters** on the Action menu.

Updater History

Viewing the history of an updater can show whether it is current and when any modifications were made to it. For example, the *Date Created* field in the history might suggest that Carbon Black File Reputation added a new updater.

To view an updater's history:

- On the Updaters tab, click the View History button next to the name of the updater. Click the **Return** button to go back to the full list of updaters.

The history page includes the following information about the updater:

- Updater Name
- Platform
- Enabled (Yes/No)
- Updater Version number
- Date Created (on this App Control Server)
- Created by (on this App Control Server)
- A history of any modifications to the updater

Using the Related Views menu of the Updater History, you can see which agent-managed computers have the latest rule for this updater.

Automatic Cloud Management of Updaters


Changes in the products or product versions from software providers might change the list of updaters you need. App Control tracks changes to the updaters for supported products as well as the arrival of new products with their own updaters. When you install a new version of App Control, the updater lists are modified to reflect these changes (if any). However, you might need to have the updater list refreshed between releases.

By allowing Carbon Black File Reputation to maintain the updaters, you can get new and modified versions as soon as they become available. Enabling Carbon Black File Reputation updates also means that obsolete updaters are deleted from the updater list. Note that this feature is enabled by default if you have Carbon Black File Reputation enabled.

To enable or disable cloud updates of updaters:

1. On the console menu, choose **System Configuration** on the Administration (Gear) menu.
2. On the System Configuration page, click **Advanced Options** on the menu. The Advanced Options Configuration page appears, with the Software Rules Options panel at the bottom.
3. At the bottom of the page, click the **Edit** button.

4. In the Software Rule Options panel, the Carbon Black File Reputation *Rapid Configs and Updaters* option is enabled by default:
 - a. If you *do not want* Carbon Black File Reputation to keep your updaters current, *un-check* the *Rapid Configs and Updaters* box and then click the **Update** button at the bottom of the page.
 - b. If you want to *re-enable* automatic updates from Carbon Black File Reputation after they have been disabled, check the box and click the **Update** button.



5. In the Confirm Server Setting Change dialog, click **Yes** to save your changes.

Alerts for Updater Changes from the Cloud

You can enable an alert that will notify you each time an Updater is created, modified, or deleted. This is recommended if you enable automatic cloud updates.

To enable alerts for Updater changes delivered from the cloud:

1. On the console menu, choose **Tools > Alerts**.
2. Check the box next to the Updater Modified Alert.
3. On the Action menu, choose **Enable Alerts**.

Approving by Trusted Directory

If your organization uses software deployment tools, or if you want to dedicate a computer for software approval, you can use a trusted directory to automatically approve software during regular roll-outs. Trusted directory approval easily integrates with existing software deployment processes. All software in the specified trusted directory of your deployment server is automatically approved. The level of approval provided by a trusted directory depends upon the platform on which it is located and applicable policies, if any.

App Control has been tested with and fully supports trusted directory approval with common deployment technologies. Please contact Carbon Black Support to determine whether your deployment method is supported and for guidance on any special considerations for integrating it with App Control.

Trusted Directory approvals are not sent to agents immediately upon activation of the directory or addition of files to it. There are three conditions that cause a trusted directory file approval to be sent to endpoints:

- **Blocked files** – If the App Control Server has a record of a file being blocked *on any endpoint* and that file is later approved by trusted directory, the server begins sending the approvals of the file to agents immediately.
- **Execution attempts** – If a user attempts to execute an instance of a file approved via trusted directory on a computer connected to the App Control Server, the server will allow the agent to run the file immediately, and also will begin sending the approval to other agents.

- **Installers** – If a file approved via trusted directory is identified as an installer, the App Control Server begins sending the approval of the file to agents immediately.

Even if a file is approved by trusted directory and not blocked by another rule, until its approval is sent to agents because of one of the cases above, instances of the file may be locally unapproved and may block if the agent computer is disconnected from the server before the approval is distributed.

Important

Avoid using removable media for trusted directories. If a removable device is disconnected and then reconnected, it is not rescanned, and so any new content is unprocessed and untrusted. You would have to disable and re-enable the trusted directory to trust the new content. Configure trusted directories on permanently attached fixed media so that the agent can monitor modifications and additions, and can process any new content.

Windows Trusted Directories

On Windows computers, files found in a trusted directory (and any of its sub-folders) are themselves approved.

Installers and Archives in Trusted Directories

Archives and installers are file types that can generate other files. It can be convenient to put both types of files in a trusted directory to make file approvals more efficient, but note that they are treated differently:

- **Installers** – App Control recognizes these common Windows formats as *installers*: Nullsoft, Wise, Install Shield, and MSI. You also can manually mark files as installers.

In a trusted directory, an installer file is globally approved and added to the File Catalog. If the system hosting the trusted directory is running an agent, the installer is also added to the Files on Computers list. Installer files are *not* analyzed to determine the files they *will write* when run, nor are the files an installer will write added to the File Catalog or Files on Computers list until the installer is actually run. Files instances written by an installer are locally approved but not globally approved.

- **Archives** – App Control recognizes the following Windows formats as *archives*: 7Zip, BZip2, CAB, GZip, ISCab, ISO, MSCompress, RAR, ZIP and TAR.

In a trusted directory, archive files are analyzed by App Control to determine what files they will write when expanded. The files that will be written by the archive file are globally approved and added to the File Catalog, even if there are no instances of them yet. They are not, however added to the Files on Computers inventory until the archive is expanded on some computer. The top-level archive file (e.g., myfiles.ZIP) is not added to the File Catalog.

Windows Image (WIM) files are commonly used to package operating system files. By default, they are *not* recognized by App Control as archives, but you can enable analysis and approval of the content of WIM files completing the following procedure for each system hosting a trusted directory.

To enable trusted directory approvals of WIM file contents:

1. Choose or create the trusted directory in which you want to be able to approve the content of WIM files. On the system where the trusted directory is located, download the Microsoft Windows Automated Installation Kit (AIK). This includes **imagex.exe**, which is required for WIM approval.

<https://www.microsoft.com/en-us/download/details.aspx?id=10333>

Although they are not included in the list of supported operating systems, Windows 7, 8, and 10 (and their server equivalents) should be able to use this kit.

2. Disable tamper protection on the agent so that ImageX.exe can be added to the agent folder.
3. From the Windows AIK download location, copy ImageX.exe into the agent installation directory (typically C:\Program Files (x86)\Bit9\Parity Agent).
4. In the console, approve the ImageX.exe file(s) on the agent with the trusted directory.
5. Re-enable tamper protection on the agent.
6. In the console, enter the URL for the Support page:
`https://<serveraddress>/support.php`
7. Click on the **Advanced Configuration** tab, and in the Agent Configuration panel, check the box for **Enable Deep Crawl**.
8. In the **Deep Crawl Files** line, add "*.wim" to the end of the list of file extensions if it is not already there. Use a comma to separate the new extension from the previous one in the list. Click **Update** when you are finished.

Mac and Linux Trusted Directories

On Mac and Linux computers, top-level files found in a trusted directory (and any of its sub-folders) are approved, but their contents are not analyzed or approved. For example, files that an installer *would* install or files that *could be* extracted from an archive file are neither analyzed nor approved when their top-level file is placed in a trusted directory.

However, on Mac computers, if a PKG file is placed in a trusted directory, it becomes an approved *installer*. This means that even though the PKG file was not analyzed, anything written from the PKG by the installer process will be approved.

For both Mac and Linux trusted directories, you can accomplish global approval of the files for an application or archive by expanding or extracting the package so that the files it would install or extract are actually in the trusted directory.

Creating a Trusted Directory

Trusted directories must be on a computer with the App Control Agent installed. You specify the deployment server name, the directory to trust on that server, and if that trusted directory applies to all or specific policies.

To use a trusted directory for automatic software approval:

1. If you haven't already done so, install the App Control Agent on the deployment server. Wait for the server's files to complete initialization. You can monitor initialization status of the deployment server on the Computers or Computer Details page (see "[Viewing Complete Details for One Computer](#)" on page 159).

2. On the console menu, choose **Rules > Software Rules**. The Software Rules page appears. The default tab for this page is Updaters.
3. Click the **Directories** tab. The table of Trusted Directories appears.
4. Click the **Add Trusted Directory** button. The Add Trusted Directory page appears:
5. Enter information about the deployment server, the policies the trusted directory applies to, and the status of the trusted directory.

The table that follows shows the trusted directory fields and their possible values.

Table 40: Trusted Directory Parameters

Field	Description
Name	Name used to identify the automatic approval instance in the Trusted Directories table. This can be any text.
Computer	<p>Agent-managed computer that is or will be your software deployment server. Use the computer name as it appears on the Computers page. For computers in domains, this includes both the domain and the computer name, in one of the following formats:</p> <ul style="list-style-type: none"> • DOMAIN_NAME\computer_name (Windows only) • computer_name.domain.extension (all platforms) <p>Note: If you edit the computer name for an existing Trusted Directory and the App Control Server has seen multiple computers by the new name, trusted directories are created for each one.</p>
Directory	<p>Deployment directory for the deployment server. Depending on the deployment technology, you may need to separately specify more than one deployment directory. For example, Microsoft WSUS requires the following directories (substitute actual drive letters):</p> <p>C:\WSUS\WsusContent\ C:\Program Files\Update Services\Selfupdate\ Note: Use of removable drives for trusted directories is not recommended. Removable drives are not re-scanned when removed and reattached, so new software might not be trusted.</p> <p>Platform Note: When you enter a path, be sure to use the correct directory delimiters, and to use only characters and formats legal for paths in the chosen platform. The App Control Server does not convert paths between platforms (e.g., '\ to '/'). Also, keep in mind that Linux files and paths normally are case sensitive.</p>
Description	Optional additional description of this trusted directory.

Field	Description
Policies	<p>Select the policies applied to the trusted directory:</p> <p>All Current and Future policies – Select this option to apply this trusted directory to all current policies and any policies created in the future.</p> <p>Selected policies – Select this option to apply the trusted directory only to the selected policies.</p> <p>NOTES:</p> <ul style="list-style-type: none"> - The policy list populates upon selection. - You can only select policies that you have permission to manage.
Status	<p>Select one of the following:</p> <p>Enabled – Software present in the trusted directory on the deployment server will be approved for installation on all computers.</p> <p>Disabled – Software present in the trusted directory on the deployment server will not be approved for other computers. Software installed from this directory will be treated according to the settings of the policy to which the deployment server belongs.</p>

6. Click the **Save** button. The approval computer and specified configuration information appear in the Trusted Directories table.
7. When you are ready to use the trusted directory, make sure it is enabled if you did not do so when it was created.
8. Deploy software according to your established procedures. If you want to use the trusted directory to approve Mac or Linux applications, see [“Mac and Linux Trusted Directories”](#) on page 273.

When you enable a trusted directory:

- All files (including files in sub-folders) actually present when the trusted directory was enabled are globally approved, as are any files you add after you enable the trusted directory.
- Files identified as installers in Windows trusted directories are globally approved, and files they write are locally approved when and where they are written. Similarly, archive files are globally approved and the files they will write when expanded are globally approved.
- The computer on which the directory resides is configured for permanent prioritization of updates so that any rule changes are applied to it as soon as possible. This status can be changed on the Computer Details page.

Note

If you make an existing Windows deployment folder a trusted directory, the App Control scanning process that analyzes and approves the directory's contents can take several hours to complete if the folder contains a large amount of software.

Verifying Trusted Directories

There are several ways you can confirm that a trusted directory is working, and that files in it are being approved.

To check the status of a trusted directory:

1. On the console menu, choose **Rules > Software Rules**, and on the Software Rules page, click the **Directories** tab. The table of Trusted Directories appears, and shows the status of each trusted directory and the progress of analysis of its contents so far.
Note: See [“Tracking Analysis Progress in Trusted Directories”](#) below for a description of what the progress indicator means.
2. If you choose, you can click the View Details button next to a trusted directory to view just the details for that directory. This details page may include additional status information.

You also can check the Events page for trusted directory-related events. There are event subtypes that show directory creation and modification activity as well as the results of any file analysis that occurs in the trusted directory.

To verify that the files on the deployment server are being approved, you can choose **Approved Files** from the Saved Views menu on the File Catalog tab and search for one of the files you expect to see approved. How quickly newly approved files from a trusted directory appear in the Approved Files table depends upon the number of files in the directory and the amount of other activity on the App Control Server. To update the Approved Files table, use the Refresh Page button on the File Catalog page.

You also can add a filter to the Approved Files view to see all files approved because of trusted directories. On the Add filter menu, choose **File State Reason**, and then complete the filter by choosing **is** and **Trusted Directory** from the File State Reason menus.

Tracking Analysis Progress in Trusted Directories

There is a progress indicator in both the rows for each directory in the Trusted Directories table and the Edit Trusted Directory page for a single directory. The Progress field shows the number of “crawl jobs” that have been processed in the directory versus the total number queued there.

?
Edit Trusted Directory

Trusted directory settings

Name:

Computer:

Directory:

Description:

Policies: All Current and Future policies
 Selected policies

Note: You can only change policies that you have permission to manage

<input type="checkbox"/> Policy
<input type="checkbox"/> APC-79895-uorrrt
<input type="checkbox"/> Default Policy
<input checked="" type="checkbox"/> Help Desk

Status: Enabled Disabled

Notice: This directory has not yet been verified by the agent.

Progress: 672/821

History

Created By: admin
Date Created: Jun 30 2020 01:43:31 PM
Last Modified By: admin
Date Modified: Jun 30 2020 01:44:28 PM
CL Version: 1172

Crawl jobs are investigations to discover and analyze files. There are two crawl job types:

- discovery and enumeration of the contents of a directory
- a “deep crawl” to discover and enumerate the contents of an archive file – see [“Installers and Archives in Trusted Directories”](#) on page 272 for a list of file types that are recognized as archives by App Control

The first crawl job in a Trusted Directory is that directory itself. When the top-level directory is crawled:

- Any individual, non-archive, executables and scripts at the any level are reported to the server and approved (unless banned by another rule) without requiring a crawl job of their own.
- Any archive files are scheduled as crawl jobs.
- Any directories are scheduled as crawl jobs.
- This process is recursive, so, for example, an archive inside another archive is new crawl job.

As you monitor progress for a Trusted Directory, keep in mind that because of the processes described above, the changes in numeric values in the Progress field do not necessarily reflect a linear time progression. Also, as different sub-folders are crawled, the total number of crawl jobs queued might actually increase even if you have not added any files to the directory. The Progress field is cumulative – the numbers do not reset once the queued and completed crawl job numbers match.

Verifying Approval of Windows Packages

For Windows installers, you can verify that App Control recognized and approved the installer in a trusted directory (and so will locally approve files it installs). On the File Catalog tab, the Saved View called **Trusted Packages** lists installers that are globally approved because they are in a Trusted Directory. This list also includes the App Control Agent installers. Files that are not recognized as installers will not appear in this table.

In the Trusted Packages view, click the View Details button next to a package name to display its File Details page. Click the package name for a table of associated files written by the package.

Custom Rules for Installer Access

App Control supports a Custom Rule that creates a “trusted path.” A trusted path can be useful as a network location in which you place installers so that computers in some or all policies can execute them.

The local state of any files written by a file in a trusted path depends upon the *Execute Action* command used. If the Execute Action is *Allow*, an installer is allowed to write files but those files are not locally approved by the action. If the Execute Action is *Allow and Promote*, the installer can write files and those files will be locally approved (unless already banned). In either case, the global state of any files written is unaffected by the trusted path. See “[Trusted Paths](#)” on page 441 for more details.

Removing or Disabling Directory Trust

If you decide to remove trust from a trusted directory, you can do one of two things:

- You can *disable* the trusted directory so that files added *after* you disable it are no longer trusted. You do this by clicking the View Details button next to its name, clicking the **Disabled** status radio button, and then clicking **Save**. Consider this if you want to temporarily suspend installations from your deployment server. Disabling (rather than deleting) gives you the option of re-enabling the directory at a later time without having to reenter all of its properties.
- You can *delete* the directory from the Trusted Directories list by clicking the **X** button next to its name. This deletes its *trusted status* in App Control, not the actual folder. Delete the folder itself if you do not want its contents on your deployment server.

Notes

- Disabling or deleting trusted directory status does not remove approval from files that were already in the directory.
- A Trusted Directory folder that is either deleted from the computer or inaccessible to App Control Agents due to network issues is listed as *Enabled, Inaccessible* in the Trusted Directories table.

Approving by Trusted User or Group

App Control supports installation privileges for users who need to install software on their own or others' computers when the computers are under High Enforcement protection. You can trust individual users or specify trusted groups whose members become trusted users.

Trusted users and users in trusted groups have full permission to install software (unless banned) on any accessible computer that allows them to log in with their credentials. Applications installed by a trusted user are locally approved where they are installed.

Trusted users can also execute unapproved files, however, the file state remains unapproved.

Caution

When you designate a trusted user or group, you grant a broad privilege to install and approve software on all of your endpoints. This privilege should be granted only if absolutely necessary, and should be disabled when not needed.

If the installations you need to perform are limited in scope, consider creating Custom rules that match those limits instead of granting global trusted user status. For example, you can create an "Allow and Promote" Custom rule that promotes processes initiated by a specified user or group, and allows execution of unapproved files by that user, but only when executed from a specified location.

How Groups are Specified

For Mac and Linux, you specify a group by entering its name.

For Windows, you have the following choices for specifying a group:

- If AD is implemented, you can specify an AD group. You enter it by typing in the group and domain name, or an SID.
- You can pick a built-in Windows group from a menu.

If you choose AD users or groups:

- You can specify trusted AD users or groups as long as the App Control Server has access to AD information about that user or group.
- AD-based privileges are determined when a user logs in. If you change an AD group in a way that affects App Control Console privileges, any logged-in users in that group are not affected until the *next* time they log in.

If you choose a built-in Windows group, certain operating system versions may not provide the access you expect. On Windows Vista and later, recognition of membership in pre-defined security groups like Administrators requires that the application run as an administrator. If a group definition is necessary for a rule, consider using security groups *you* have defined rather than the pre-defined groups.

Creating a Trusted User or Group

To trust users to install software on High Enforcement computers:

1. On the console menu, choose **Rules > Software Rules** and click the **Users** tab on the Software Rules page. The Trusted Users or Groups view appears.
2. Click **Add Trusted User or Group**. The Add Trusted User or Group page appears.
3. Choose the Platform from which you will choose a user or group. Some of the fields change if you choose Mac or Linux instead of Windows.
4. If you chose Windows as the platform, enter the name of the user or group to be given trusted privileges in **one** of the following ways:
 - Leave **User or group** checked and enter a valid domain and user name in either of these formats: `DOMAIN_NAME\user_name` or `user_name@DOMAIN_NAME`
 - Leave **User or group** checked and enter a valid AD group name in one of these formats: `DOMAIN_NAME\group_name` or `group_name@DOMAIN_NAME`
 - Leave **User or group** checked and enter a valid User or Group SID.
 - Click the **Pre-defined group** button and choose a Windows group from the menu.
5. If you chose Mac or Linux as the platform, enter the name of the user or group to be given trusted privileges in **one** of the following ways:
 - Leave **User** selected and enter a valid user name for the platform you chose.
 - Click **Group** and enter a valid group name for the platform you chose.
6. Click the **Save** button. The user or group appears in the Trusted Users table.

Removing Trust from a User or Group

If you no longer want a user or group to have installation privileges on locked-down computers, you can remove that user or group from the Trusted Users or Groups table. You do this by clicking the Delete (X) button next to the entry for that user or group.

Important

- If you eliminate App Control trust from a user or group, that user or group loses its trusted status almost immediately, as soon as agents receive the change. This means the user is not trusted to perform new installations. However, a process that was created when the user was trusted remains trusted until the process exits.
- If you remove a user from an AD group that is trusted by App Control, the user continues to be trusted until he or she logs out.
- When a trusted user logs off of a computer, sessions the user initiated may be cached by the operating system. When that user logs in again, trust privileges could continue, even after being disabled through the console. It might be necessary to reboot a system to remove user trust from all processes.

Approving or Banning by Publisher

Platform Note

Publisher approvals and bans currently work only on Windows computers. They have no effect on files on other platforms.

Many files are signed with a digital certificate that verifies the integrity and identity of the file, including the name of its publisher. The Publishers tab of the Software Rules page lists each unique publisher identified in a valid certificate for a file discovered by a App Control Agent. If Windows can find the digital signature on a file, the publisher is discovered and listed in the console.

Once a publisher is listed in the console, it may be approved, banned, or left unapproved. Publisher approvals and bans can be applied to all computers or to computers in specific policies. You may Acknowledge a publisher to indicate that you have seen it and do not need to track it as closely. Acknowledging a publisher does not change its state.

Publisher state affects files differently depending upon whether you have banned or approved the publisher:

- **Bans** – When you ban a publisher, any file signed by a certificate identifying that publisher is banned.
- **Approvals** – When you approve a publisher, a file signed by a certificate identifying that publisher is approved *if its certificate meets additional App Control validation requirements*. These requirements are described in more detail in the section [“Determining Which Certificates Can Approve Files”](#) on page 286.

Note

App Control also allows approval or banning of certificates themselves. This is a more secure but more complex way to identify and control files by identifying their source. See [Chapter 11, “Managing File-Signing Certificates,”](#) for more details.

Publisher Approvals

You might approve files by publisher when it is not practical to approve applications using a trusted directory and you want to permit all users to install all software from a particular source. Applications from approved publishers are permitted to be installed and run on computers in the policies to which the approval applies. The Global State of publisher-approved files is changed (if necessary), but the File State is not changed (see [“Global File State”](#) on page 249). Each instance of such files is locally approved, and therefore allowed to run on the computer on which it is present.

Approving by publisher allows you to assure that new files from a trusted source are pre-approved when they arrive on an agent-managed computer. It also can reduce the amount of rule traffic sent to agents since it is not necessary to send an individual rule for each file.

There are two ways to approve a publisher:

- **Manual Approval** – You can choose to approve publishers that you select from the list on the Publishers tab. Manual approval is described in this section.

- **Reputation Approval** – You can enable automatic approval of all publishers that meet a particular trust threshold as reported by Carbon Black File Reputation. Approving a publisher by reputation has the same effect on *existing* files as approving it manually. In addition, as soon as a file with a new publisher is discovered on one of your computers, the publisher is approved if it is known to Carbon Black File Reputation and meets the trust level you chose. Details and considerations for reputation approval of publishers are described in [Chapter 10, “Reputation Approval Rules.”](#)

Important

Before approving a publisher, consider all possible files that could come from that publisher. Once the approval is added, *all* executables and script files from the publisher will be locally approved. You can remove the publisher from the Approved list, but this only affects files not yet encountered on your network at the time of the change – no single operation can remove *file* approval from all files already locally approved by a publisher approval.

Publisher Bans

When you ban a publisher, computers in policies affected by that ban cannot run software from that publisher. You might ban files by publisher when you know that the publisher is a source of malicious files or applications that you simply don't want running in your environment. When you create a publisher ban, the local state of files from that publisher is changed to Banned. Publisher bans are created manually through the console.

You can ban files by publisher even if they are invalidly signed or do not meet other requirements for approval by publisher.

Important

As with approvals, consider all of the files that might be affected by a publisher ban and be sure that a publisher ban does not inadvertently ban a file required in your environment.

Managing Bans and Approvals from the Publishers Tab

On the Publishers tab, you can approve, ban, or remove bans and approvals from multiple publishers at one time. Publisher state changes done this way apply to all policies.

When you check more than one publisher in the table, you must perform the same state change on them; that is, you must ban them all, approve them all, or remove the ban or approval on all. You cannot ban some publishers and approve others in a single operation.

To approve or ban software from one or more publishers for all policies:

1. On the console menu, choose **Rules > Software Rules**. The Software Rules page appears.
2. Click the **Publishers** tab. All publishers of validly signed software discovered on agent-managed computers reporting to your server, plus any publishers whose certificates you added manually, appear in the Publishers table:

The screenshot shows the 'Software Rules' console interface. At the top, there are tabs for 'Publishers', 'Users', 'Directories', 'Files', 'Custom', 'Memory', 'Registry', and 'Scripts'. Below the tabs, there are 'Saved Views' and 'Group By' dropdowns. A search bar and an 'Add Publisher' button are visible. The main table has columns for 'Name', 'State', and 'Date Approved or Banned'. The table shows 20 items, all in an 'Unapproved' state. The visible rows are:

	Name	State	Date Approved or Banned
<input type="checkbox"/>	Microsoft Windows	Unapproved	
<input type="checkbox"/>	Microsoft Windows Publisher	Unapproved	
<input type="checkbox"/>	VMware, Inc.	Unapproved	
<input type="checkbox"/>	Microsoft Windows Hardware Abstraction Layer Publisher	Unapproved	
<input type="checkbox"/>	Microsoft Windows 3rd party Component	Unapproved	

3. In the table of publishers, locate the publishers you want to approve, or the publishers you want to ban. Keep in mind that the table may be several pages long.

Note

Files from the same company can be identified as being from different publishers, often based on minor changes in punctuation. These appear as separate lines in the Publishers table. For example, you might see both “Adobe Inc.” and “Adobe, Inc.” in the table. You can approve (or leave unapproved) each instance separately. If files signed by a publisher appear as unapproved on the Files page and you want these files approved, be sure to approve the correct version of the publisher certificate.

4. Review the publisher(s) you are interested in approving or banning. If necessary, open the Publisher Details page for specific publishers for more information.
5. Check the checkbox next to the name of each publisher whose state you want to change. You can check as many names as you want on one page. Note that approval and ban actions are applied to the currently visible page only.
6. When you have checked all the publishers (on the current page) whose state you want to change, on the Action menu:
 - a. Choose **Approve Publishers** to approve all of the selected items.
 - b. Choose **Ban Publishers** to ban all of the selected items.
 - c. Choose **Remove Approval or Ban** to return all selected publishers to the Unapproved state.

Managing Bans and Approvals from the Publishers Details Page

For a single publisher, you can use the Publisher Details page to approve or ban the publisher, or to remove an approval or ban. You also can change the policies to which an approval or ban applies.

Publisher Details

General

Publisher Name: Google LLC

State:

Acknowledged:

Trust: Not Trusted (Unknown)

Description:

Rule Applies To: All Current and Future Policies
 Selected policies

Platforms: All platforms
 Selected platforms

▶ All Certificates For This Publisher (click to expand)

History

Date First Seen: Jun 29 2020 07:51:57 PM
Platform First Seen: Windows
Computer First Seen: WORKGROUP\APC-79895
Date Approved: Jun 30 2020 01:18:43 PM
Approved By: admin
CL Version: 1151

Related Views

- All files signed by this publisher
- All Computers that have received this rule
- All Computers that have not yet received this rule

To approve or ban one publisher by policy (Publisher Details page):

1. On the console menu, choose **Rules > Software Rules**.
2. Click the **Publishers** tab. All publishers of validly signed software discovered on agent-managed computers reporting to your server, plus any publishers whose certificates you manually added, appear in the Publishers table.
3. From the table of publishers, locate the publisher whose state you want to modify and click on the View Details button. The Publisher Details page opens.
4. In the State field, choose **Approved** or **Banned**.
5. If you choose, change the Acknowledged state to **Yes**. This indicates that you have reviewed the publisher so that you can concentrate on publishers you haven't yet reviewed. To do this, you can filter the Publishers table using the Acknowledged field. Acknowledging a publisher has no impact on its approval state.
6. In the Rule Applies To field, click the radio button for **All policies** or **Selected policies**.
7. If you chose *Selected policies*, check the box next to each policy for which you want the publisher approval or ban to be enabled.
8. In the Platforms field, click the radio button for **All platforms** or **Selected platforms**.
Platform Note: Publisher approvals and bans currently affect only Windows agents.
9. When you are finished configuring the approval or ban, click the **Save** button.

Adding Publishers

Any publisher already identified through a file on a computer running the App Control Agent appears in the Publishers table, but you might want to approve a publisher before its files arrive on your computers. This could be the case, for example, if you distribute software using a computer that does not run the App Control Agent. To address this, you can *manually* add publishers to the table.

To add a publisher:

1. Open a browser and log in to the console on a computer with access to the file whose publisher you want to add. It might be most convenient to do this on the computer that has the file.
2. On the Publishers tab, click the **Add Publisher** button to view the Add Publisher dialog:
3. Click the **Browse** button and locate an application file validly signed by the publisher. You can browse to any validly signed, executable file and add its publisher:
4. In Windows, confirm that the file is signed by right-clicking on the file and choosing Properties from the menu. If there is a Digital Signatures tab on the Properties window, the file is signed and you can examine its credentials.
5. Double-click the filename to enter it into the File Name field.
6. Click the **Save** button. Publisher information is extracted and the publisher is added to the table, initially in the Unapproved state.
7. If you want to approve or ban this new publisher for all policies, check the box next to its new entry in the Publisher table and choose **Approve Publishers or Ban Publishers** from the Action menu. The publisher is approved, and if you have the table grouped by State, the publisher moves into the appropriate *State* section. Now, as soon as a file from this publisher appears on one of your agent-managed computers, it will be handled as you instructed.

You also can approve or ban the publisher by policy from the Publisher Details page.

Note

When you add a publisher manually, the App Control Server creates a temporary copy of the file you identified and then deletes it after the publisher has been added. If an agent is running on the server computer, the file will appear in the File Catalog, but will have a prevalence of zero.

Removing Publisher Approvals

To change an *approved* publisher to *unapproved*, go to the Publisher tab on the Software Rules page, check the box next to its name and choose **Remove Publisher Approval** on the Action menu. This simply removes approval; it does not ban the publisher. You also can remove approval using the Publisher Details page.

Any computers that have installed or run software from this publisher while it was approved continue to be able to run the software. All existing instances of software from an approved publisher are locally approved, and the local approval is not removed by the change in publisher status on the App Control Server.

Removing Publisher Bans

To change a *banned* publisher to *unapproved*, go to the Publisher tab on the Software Rules page, check the box next to its name and choose **Remove Publisher Approval or Ban** on the Action menu. This simply removes the ban; it does not approve the publisher. You also can remove the ban by choosing **Unapproved** in the State menu on the Publisher Details page.

When a publisher ban is removed, the files from that publisher revert to whatever their state would have been without the publisher ban.

Finding All Files from a Publisher

On the Publishers tab of Software Rules, you can find all instances of files on your computers that are identified as being from a specified publisher. You do this by clicking the Find Files button next to the publisher name. You also can get this list using the Related Views menu on the Publisher Details page.

Determining Which Certificates Can Approve Files

Publisher identification and approval of files by publisher approval are both based on digital certificates. If you are unfamiliar with certificates, the following web sites may provide useful background:

[http://msdn.microsoft.com/en-us/library/ms537361\(v=vs.85\).aspx](http://msdn.microsoft.com/en-us/library/ms537361(v=vs.85).aspx)

<https://sites.google.com/site/ddmwsst/digital-certificates>

It is important to distinguish between approval of a publisher and approval of a file identified as being from that publisher. You can approve any publisher that appears on the Publishers tab of the Software Rules page. A publisher appears in this list if a file had a certificate identifying the publisher and the signature was considered valid by Windows.

However, a *file* identified as being from this publisher can be approved by publisher only if all certificates in the certificate chain for that file are considered valid by Windows. For example, current root certificates must be installed for a certificate to be accepted.

Note

Microsoft security bulletin MS13-098 describes a flaw in the Authenticode signature verification that could allow remote code execution. In response, Microsoft announced availability of an update for all supported releases of Windows to change how signatures are verified for binaries signed with the Windows Authenticode signature format. If this change is enabled, Windows Authenticode signature verification no longer allows extraneous information in the WIN_CERTIFICATE structure, and Windows no longer recognizes non-compliant binaries as signed. Activation of this new behavior could cause files previously approved by publisher to be blocked by App Control.

The change is included with Security Bulletin MS13-098, but (as of July 2014) will only be enabled on an opt-in basis. However, Microsoft states that it may make this a default behavior in a future release of Microsoft Windows. See <https://technet.microsoft.com/library/security/2915720> for more information on this change.

All certificates in the chain for a file must also meet additional App Control requirements. These settings are configurable on the Advanced Options tab of the System Configuration page. Keep the following in mind about these certificate settings:

- It is best to set certificate configuration options *before* generating the agent installation packages (i.e., as soon as possible after installing App Control Server). This assures that all agents, including those disconnected from the server, will handle certificates as you want them to. In addition, changing certificate settings after the agent is installed requires re-evaluation of certificates to occur on each agent. Having these settings correct before deploying the agent avoids a significant amount of processing.
- Changing any of the configurable certificate settings does not remove local approval of files whose certificates met the previous settings and were approved by publisher.
- Changing certificate settings may affect the tracking and inventory of Microsoft Support Files. See [“Changes that Affect OS Inventory Tracking”](#) on page 232.

To view and change configurable certificate approval options:

1. On the console menu, click the configuration (gear) icon and choose **System Configuration**.
2. On the System Configuration page, click the **Advanced Options** tab. The Advanced Options Configuration page appears, with the Certificate Options panel at the bottom.
3. At the bottom of the page, click the **Edit** button.
4. **Expired Certificates:** In the Certificate Options panel, use of expired certificates is enabled by default. See [“Approval with Expired Certificates”](#) on page 288 for information that may assist you in configuring this option:
 - a. To *disable* use of expired certificates, *un-check* the *Expired Certificates* checkbox.
 - b. To *re-enable* use of expired certificates after it has been disabled, check the box.
5. **Exclude Publisher Approvals With These Certificate Algorithms:** Review the currently checked boxes in this field. See [“Excluding Certificate Algorithms”](#) on page 288 for information that may assist you in configuring this option.
 - a. To *prevent* publisher approvals of files signed by certificates with a certain algorithm, check the box next to the algorithm name.
 - b. To *allow* publisher approvals of files signed by a certificates with a certain algorithm, *un-check* the box next to the algorithm name.
6. **Minimum Certificate Key Size for Approval:** To change the minimum certificate key length required for a file to be approved by publisher, choose a new value from the menu. See [“Minimum Key Size”](#) on page 288 for information that may assist you in configuring this option.
7. **Digital Countersignatures:** To require a countersignature for the digital signature of each certificate, check the *Require countersignature* box. If you do not want to require a countersignature, *un-check* the box. See [“Countersignature Options”](#) on page 288 for information that may assist you in configuring this option.
8. **Initial/Background Revocation Check:** Two separate settings control checks for certificate revocation: *initial*, which controls the revocation check when a file is first discovered, and *background*, which controls ongoing checks that occur (if enabled) every 24 hours. See [“Revocation Checks”](#) on page 289 for more on these settings.
9. If you changed any settings, click the **Update** button at the bottom of the page and in the Confirm Server Setting Change dialog, click **Yes** to save your changes.

Approval with Expired Certificates

By default, App Control allows the use of expired certificates whose (verifiable) timestamp is within the certificate validity period to approve files by publisher. If the timestamp is missing, invalid, or is before or after the certificate validity period, then the software cannot be approved by publisher.

You can disable approval by expired certificates that would otherwise be trusted by App Control. This provides extra security, but can prevent approval of legitimate files whose valid certificate is now out of date.

When you disable *Allow approval of software with expired certificates*, all publishers are re-evaluated. However, if a file was locally approved by a publisher with an expired certificate when this was allowed, it remains locally approved when the setting is disabled.

The Expired Certificates setting has no effect on *bans* of publishers, so you can ban files by publisher even if they have an invalid signature or an expired certificate.

Important

It is especially important to set the expired certificate option before generating installation packages for agents that will be primarily or permanently disconnected from the server. This assures that disconnected agents will handle expired certificates as you want them to.

Excluding Certificate Algorithms

With the *Exclude Publisher Approvals With These Certificate Algorithms* option, you can disallow publisher-based approval of files whose certificates use certain algorithms. If an algorithm box is checked, files whose certificates use that algorithm *cannot* be approved by publisher. If not checked, a certificate using that algorithm may be used to approve files by publisher. The choices are: MD2RSA, MD5RSA, SHA1RSA, and SHA256RSA. The default for new Parity installations beginning with 7.0.1 Patch 11 is to allow certificates with any of the listed algorithms to be used for approvals. Upgrades and patches from previous releases also allow certificates with any of the listed algorithms to be used for approvals unless the setting was modified through the console before the upgrade.

Minimum Key Size

The *Minimum Certificate Key Size for Approval* option allows you to specify a minimum key length for a certificate to be used for file approval. Choices range from 512 to 4096. Certificates whose key size is greater than or equal to the chosen value may be used to approve files. Certificates whose key size is smaller than the chosen value may not be used for file approval. The default value for new Parity installations beginning with 7.0.1 Patch 11 is 512. Upgrades and patches from previous releases also use this value unless the setting was modified through the console before the upgrade.

Countersignature Options

You can require that the digital signature for a certificate is countersigned in order for App Control to approve a signed file by publisher. This can provide greater security against manipulation of time stamps on a signature. By default, the box is not checked (i.e., no countersignature is required). If the box is checked, certificates that are not countersigned are not considered valid for use in approval by publisher.

Note the following additional details of countersignature handling:

- If the box is unchecked, signatures lacking a countersigner are only valid for the life of the signing certificate.
- Regardless of this setting, if a countersignature is present, it must be valid for the digital signature to be considered valid.

Revocation Checks

There are two settings that control if and how the agent checks to see whether a file's certificate has been revoked:

- **Initial Revocation Check** – This determines whether, and if so, how a certificate revocation check is done when a file is initially discovered on an agent.
- **Background Revocation Check** – This determines whether, and if so, how a certificate revocation check is done in the background every 24 hours.

For each of the revocation settings, there are three possible values:

- **Network** – If revocation information is not locally available then use the network to retrieve the revocation status of a certificate.
- **Cache** – Use locally available revocation status information when performing certificate revocation (the network will not be used).
- **None** – Do not perform certificate revocation checking.

Consider your agent deployment scenario when setting these values since they can impact agent performance. For example, if you have offline agents, you might want to avoid using the Network option, especially for the Initial Revocation Check. Also keep in mind that the daily revocation check is performed in the background, and is less likely to have a negative impact on agent performance, whereas the initial revocation check setting may have a noticeable effect on agent performance.

Note

Regardless of whether agent-based certificate revocation checks are enabled, the App Control Server validates certificates in its inventory on a recurring basis to make sure that they have not been revoked. This validation generally occurs on a weekly basis and involves downloading certificate revocation lists (CRLs) from registration authorities or making Online Certificate Status Protocol (OCSP) calls to OCSP responders. If you are monitoring network traffic, keep in mind that these downloads might involve a variety of sites in a variety of countries.

Server-based validation checks are provided to inform administrators when the status of a certificate changes, but they do not affect enforcement of rules. Enable agent-based revocation checks if you want revocations to affect rule behavior.

Locally Approving Files

When the App Control Agent is installed on a computer for the first time, the computer goes through an *initialization* process during which all files present on that computer are *locally approved* unless they are already globally approved or banned. This means that they are allowed to run on that computer, regardless of its Enforcement Level. Local approval has no effect on the *global state* of the files, however. Because files present

during agent initialization are locally approved, you can set up a computer with the files it needs to run, saving global decisions about these files for later, after you have used App Control to collect more information about the files and computers on your network.

Files that appear on a computer *after* App Control Agent initialization, if not explicitly banned or approved, are assigned *Unapproved* state. Unapproved files are allowed to run on computers running in Low Enforcement and (with user intervention) Medium Enforcement, but they are not allowed to run on computers in High Enforcement.

You might want a particular computer to be able to run a new application without approving it for any other computers on your network. You also might want to change the state of a file from Unapproved to Locally Approved on one or more computers before putting those computers into High Enforcement. To accomplish tasks like these, App Control offers the following options:

- A per-policy ability to make certain unapproved files Locally Approved when a computer makes a transition to a more secure Enforcement Level
- Local approval of individual files on a specific computer
- Local approval of all unapproved files on a specific computer
- Temporary reassignment of a computer in High or Medium enforcement to the Local Approval policy, during which any files that are installed are locally approved
- Designation of files as installers even when App Control analysis did not identify them as such, and vice versa; local approval of an installer also locally approves all of the files it installs

Note

- You cannot use any of these methods to locally approve a file that has been globally banned or that is banned by policy on the computer with the file. You also cannot remove local approval for a file that has been globally approved or that is approved by policy on the computer with the file.
- Certain approval methods, such as approving a publisher, make all instances of a file locally approved. These are not discussed in this section. See [“Approving or Banning by Publisher”](#) on page 281 for details of how publisher approvals affect file state.
- You must have full Suite licenses (Visibility and Control) to be able to reassign a computer to Local Approval policy; sites with only Visibility licenses cannot perform the reassignment.

Automatic Local Approval on Enforcement Level Change

App Control security policies have an Advanced Setting, enabled by default, that causes unapproved files discovered while App Control Agent is in a policy whose Enforcement Level is Low or None (Visibility) to be locally approved when the policy makes a transition to Medium or High Enforcement.

Automatic local approval of unapproved files allows you to install new files while in Low Enforcement and then change to a more restrictive Enforcement Level without restricting the execution of the files that existed at the time of transition. Files that you explicitly ban remain banned, and unapproved files discovered while in Medium or High Enforcement remain unapproved during transitions to and from any Enforcement Levels.

You can disable this feature if you choose, on a policy-by-policy basis. This will increase security against unwanted execution of unapproved files already on an agent before the transition, but it might also cause more blocks of non-risky software after the transition. If you do not plan to enable automatic local approval, consider other bulk approval methods that might reduce the number of individual files you must approve.

Note

Enforcement level changes can happen because a computer changes policy or because the enforcement level of the policy itself changes. If a computer changes policy, it is the setting in the *policy it begins in*, not the policy it changes to, that determines whether the approval-on-transition takes place.

To disable automatic local approval of unapproved files on Enforcement Level change:

1. On the console menu, choose **Rules > Policies**. The Policies page appears.
2. Click the View Details button next to the name of the policy you want to change. The Edit Policy page for that policy appears.
3. Click the **Show Advanced Settings** button. The Advanced Settings panel appears.
4. At the bottom of the Advanced Settings panel, un-check the *Locally approve unapproved files on transition from Visibility or Low Enforcement Level to Medium or High* checkbox.
5. Click the **Save** button.
6. Repeat steps 2-5 for any other policies you want to change.

You can re-enable automatic local approval by checking the checkbox.

Which Files Are Locally Approved On Transition

There are two types of locally “unapproved” files, and these have different Local State Details:

- Files with Local State Details of *Unapproved* were discovered on a system in None (Visibility) or Low enforcement. They will be locally approved by a change to Medium or High Enforcement Level.
- Files with Local State Details of *Unapproved (Persisted)* were discovered on a system in Medium or High enforcement. They remain Unapproved on transition.

You can view Local State Details on the Files page or Find File results (for multiple files) or the File Instance Details page (for one file). In any of the tables, add the *Local State Details* column if it is not shown.

For one policy, the Related Views menu on the Edit Policy page includes an **Unapproved files from computers in this policy** link that opens the Find Files page with the results of a file search for these files. Viewing this list may be useful before taking actions affecting local approval of unapproved files.

Locally Approving Individual Files

You might discover that one or more files you thought were present during App Control Agent initialization were missing, and as a result, those files are not locally approved. A missing file could be a standalone executable or a file whose absence prevents an application from running. If you can identify the missing files and put them on the computer, you can locally approve them on an instance-by-instance basis.

You can do local approvals from any console table that shows file instances, including:

- the Files on Computers tab on the Files page, which shows instances of tracked files on every agent-managed computer on your network
- any file view of a Baseline Drift Report Results page
- the Find Files page when you have search results displayed

Note

If you are looking for a particular file on one computer, you can add a Computer filter to your Find Files query and enter the computer's name. The resulting search will find the file you are looking for only on the computer you entered.

You can use filters on any of these pages to get exactly the list of files you want, or one particular file.

To locally approve individual file instances from a table of files:

1. Locate the file instance(s) you want to locally approve in the file table.
2. In the table, check the box to the left of each file instance you want to locally approve. Confirm that the computer name next to each file is a computer you want to affect.
3. On the Action menu, choose **Approve Locally**. The Local State of each checked file becomes *Locally Approved* for the computer on which it appeared.

Note

To get more information about a file before you locally approve it, click the View Details button in the file table to bring up the File Instance Details page. That page also includes an **Approve Locally** choice on the Actions menu if the file is not already globally or locally approved.

Removing Local Approval

Just as you can locally approve an individual file, you can *remove* local approval on a file that has been locally approved. You might choose to do this if a file you really didn't want approved happened to be on a computer at App Control Agent initialization, or if you mistakenly approved the file by one of the post-initialization methods. You locate the file or files the same way you would if you wanted to approve them, and then do one of the following:


- In a file table (Files page, Find Files page, Baseline Drift Report Results), check the box next to each file whose local approval you want to remove and choose **Remove Local Approval** on the Action menu.
- On a File Instance Details page, click the **Remove Local Approval** link.

Locally Approving Files Not Yet in File Catalog Inventory

As new files are discovered by agents, the addition of file instances to the server is processed in the background to allow efficient operation of the server and console. Because of this, the Events page might report that a new file has been discovered on a computer before that file appears as a file instance in the Files on Computers page.

You can locally approve a file from the Events page by choosing Approve Locally from the Action menu on the page. You also can click on the highlighted file path in the Event Description to go to the File Details page. If you do this for a file that is not fully processed, you see a note at the top of the File Details page.

Note: Specific file instance cannot be found - file might have been deleted or have not been processed yet. Note that computer MYCORP\LAPTOP-7 still has 35 files to process. Displaying global file details.



You can use the Approve Locally command from the Actions menu on the File Details page even though file was not found.

Locally Approving Transient or Deleted Files

There may be cases in which a file appears briefly on a computer to accomplish a particular task. One example of this is a printer driver installation, during which a temporary file could appear long enough to install the driver and then disappear. Although this file does not appear in the Files on Computers page, you might want to locally approve it by hash so that installation of this driver is not blocked by App Control on a particular computer.

As with files that are present on an agent computer but not fully inventoried, you can locally approve transient or deleted files through the Action menu on the Events page or the Actions menu on the File Details page for the file. For files that are selected and approved by a local approval command, the local approval persists indefinitely. This means that even after all copies are deleted, new instances of the same file that appear on the computer in the future will still be locally approved.

Other local approvals expire several days after all instances of a file are deleted from a computer. If a new instance of the same file appears on the same computer in the future, its state is recalculated. This "linger" time depends upon how the file was approved. Specifically:

- Approvals that are created by an automated method that can always be recalculated (such as a Publisher Rule or Reputation score) linger for 3 days by default.

- Approvals that are created by an automated method that cannot necessarily be recalculated (such as a Custom Rule or enforcement level change) linger for 14 days by default

Note

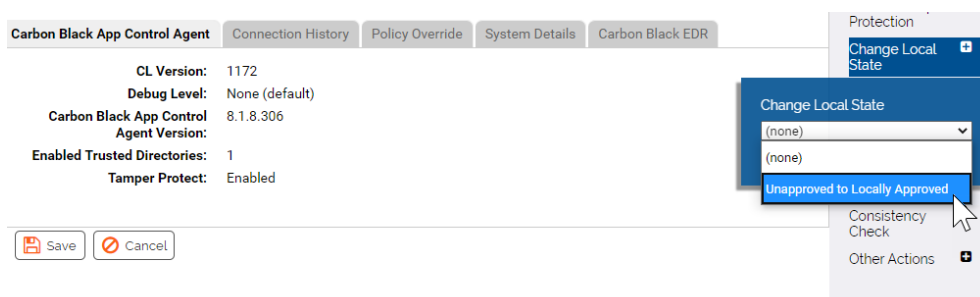
You cannot *remove* local approval of files that do not currently exist on a computer.

Locally Approving All Unapproved Files on a Computer

App Control provides a mechanism for locally approving all unapproved files on a selected computer. You might choose to do this if you have added a large number of known-good files to a computer after initialization, at which point they are in the unapproved state (if not explicitly banned or globally approved).

To change all unapproved files on a computer to Locally Approved:

1. On the console menu, choose **Assets > Computers**.
2. Click the name of the computer whose unapproved files you want to convert. The Computer Details page for that computer appears.
3. In the Advanced menu on the lower right of the page, click on **Change Local State**, choose **Unapproved to Locally Approved** in the *Change Local States* menu, and then click the **Go** button. All files whose local state on the computer was *Unapproved* are now *Locally Approved*.



Moving Computers to Local Approval Mode

Note

You must have full Suite licenses (Visibility and Control) to be able to reassign a computer to Local Approval mode; sites with only Visibility licenses cannot perform the reassignment.

To permit installation of new applications on a selected computer under High Enforcement Level, you may temporarily relax protection and give the computer permission to execute any files that are not banned. Your choice of how to do this depends upon whether the computer is connected to or disconnected from the App Control Server:

- **For an online computer**, you can use the console to move the computer into another Enforcement Level for as long as it takes to complete software installation and then

move it back when you are finished. This option is described in the section [“Moving Online Computers into Local Approval Mode”](#) on page 295.

- **For an offline computer**, you can use the console to generate a system-specific password for use on the computer to move it into another Enforcement Level for a specified time period. This option is described in the section [“Using Timed Policy Overrides”](#) on page 297.

In either case, Local Approval mode should be temporary – it has a specified time limit for the Timed Enforcement Level override, but must be returned manually for online computers, as described in [“Restoring Online Computers from Local Approval Mode”](#) on page 297.

Once you return the computer to its original Enforcement Level, all files that were in the Unapproved state before the computer was placed in local-approval mode *and were not executed while in local-approval mode* remain unapproved. Formerly Unapproved files that were run or installed while the computer was in local approval mode are locally approved on the computer but continue to have a *global* state of Unapproved.

You can move into Local Approval from both High and Medium Enforcement Level. Although you can execute unapproved files in Medium Enforcement, by using Local Approval you eliminate the need to respond to notifiers when you attempt to run unapproved files.

Moving Online Computers into Local Approval Mode

Local Approval mode allows you to install new files that will become locally approved without affecting the local state of any files already on the computer before the mode change or installed after the computer is returned to its normal policy. It is most useful if you have not yet introduced the new files you want to install on a computer.

You can use the console to move an *online* computer into the predefined Local Approval policy for as long as it takes to complete software installation. While in the local approval policy, computer users are permitted to install and run unapproved applications that were previously blocked because of High or Medium Enforcement Level, although banned files remain banned and blocked from running.

After the installation is complete, you can (and should) restore the computer to its original policy, at which point it continues to be able to run all files that were installed and locally approved while it was at the relaxed Enforcement Level.

Notes

- Unapproved software can be installed on computers in a Low Enforcement Level policy. However, you still might want to move the computer into Local Approval to approve known-good files, especially if you might move the computer to a higher Enforcement Level at a later time.
- In Local Approval, the only active Device Control settings are *Block writes to banned removable devices* and *Block executes from banned removable devices*. All others are set to *Off*.

You can move computers into Local Approval mode in several different ways, each of which also allows you to restore the computer to its previous policy:

- You can move one or more computers at a time to Local Approval mode via the Computers page.
- You can move a single computer from High or Medium Enforcement into Local Approval using the Action menu on its Computer Details page.
- You can move a single computer into Local Approval mode using the Change Policy portlet on the console Home Page (or any other dashboard it is on).

Local Approval mode has a number of special features for monitoring and control:

- You can track which machines are in Local Approval mode by choosing the Saved View *Computers in Local Approval* on the Computers page.
- You can set an alert to trigger if a computer is in Local Approval longer than a time interval you specify. See [“Using App Control Alerts”](#) on page 608 for more details.
- Computers manually moved to Local Approval mode can be easily returned to their normal Enforcement Level using the *Restore to Normal Enforcement Level* command on the Computers page Action menu.

To place one or more online computers in Local Approval mode:

1. In the console menu, choose **Assets > Computers**. The Computers Page appears.
2. In the Computers table, locate the computer to be placed in local approval mode. To reduce the number of computers displayed, you can use the Show Filters button and filter on policy or some other relevant field. You also can enter all or part of the computer name in the Search box.
3. Check the names of any computers you want to move to Local Approval mode.

Computers

Computers connected: 1 Total computers: 1 Current CL version: 1186 CL version for upgrade: 1142

Saved Views: (none) Add Group By: (none) Ascending Days Disconnected: (none)

Show Filters Show Columns Export to CSV Refresh Page

Action Search: Go Clear

Computer Name	Connected	Policy Status	Upgrade Status	Connected Enforcement	Disconnected Enforcement
<input checked="" type="checkbox"/> WORKGROUP\APC-99986	●	Up to date	Up to date	High (Block Unapproved)	High (Block Unapproved)

Page 1/1

4. On the Action menu, choose **Move to Local Approval**. The computer(s) moves into the Local Approval policy. Unapproved files may be executed and device control is disabled except for writing to banned devices, which is blocked. Note that if computers in Low Enforcement are included in your selection, the operation will fail and show an error message.
5. On the Computers Page, choose **Computers in Local Approval** on the Saved Views menu. Verify that the computer appears in the table as part of the Local Approval policy. If so, the computer user may now install software on that system and have it locally approved (if not globally banned or approved). The only active Device Control setting is *Block writes to banned removable devices*.

To move one online computer to Local Approval (Computer Details page):

1. On any page displaying a Computer Name field, click on the name. The Computer Details page for that computer appears.
2. In the Actions menu, click on **Change Policy**. The Change Policy dialog appears.
3. On the Change Policy menu, select **Local Approval** and then click the **Go** button. The computer moves into the Local Approval policy. Unapproved files may be executed and the only active Device Control settings will block writes to and execute attempts on removable devices. (Local Approval appears on the menu only for computers in High and Medium Enforcement).
4. On the Computer Details page, confirm that the Policy has changed to Local Approval. If so, the computer user may now install software on that system and have it locally approved (if not globally banned or approved).

Restoring Online Computers from Local Approval Mode

When you have put computers into Local Approval mode, you normally should restore them to their previous policy as soon as possible, after you have finished installing new application(s) on them. As with the transition to Local Approval, restoration to the previous policy can be accomplished from the Change Policy portlet, the Computer Details page, or the Computers page. The last of these is described here.

Note

The method described below works only for online computers. If you used a timed Enforcement Level override to move an offline computer into Local Approval mode, the computer will move back to its normal Enforcement Level automatically when the time period is over. See [“Using Timed Policy Overrides”](#) on page 297 for more information on that case.

To restore Local Approval mode computers to their previous policy:

1. On the console menu, choose **Assets > Computers**. The Computers page appears.
2. On the Computers page, choose **Computers in Local Approval** on the Saved Views menu and verify that the computer appears in the Local Approval policy.
3. In the table, check the box next to the computer you want to restore. If you have multiple computers to restore, select each one.
4. On the Action menu, choose **Restore to Normal Enforcement Level**. The computer moves back to its previous policy and is no longer displayed in the Computers in Local Approval view.

Using Timed Policy Overrides

You might need to install new applications on a selected computer under High Enforcement Level protection. You can do this by temporarily relaxing protection and giving the computer permission to execute any files that are not banned; that is, you move the computer into the predefined Local Approval policy for as long as it takes to complete software installation.

Because disconnected computers cannot be controlled directly from the App Control Server, you need a different way to instruct the agent to make the transition to another Enforcement Level. You can generate a special code that can be entered on a agent-

managed computer to switch its Enforcement Level for a specified amount of time. The code is specific to one agent, and it can be used only once. You can generate codes to switch a computer into any Enforcement Level except *None (Disabled)*, although this feature is primarily intended for temporary transitions to Local Approval mode.

Once the specified time for the override has elapsed, the computer is automatically restored to its original policy. If you had moved it temporarily into Local Approval, it continues to be able to run all files that were installed while it was in Local Approval. Files run or installed while the computer was in the Local Approval policy are locally approved on the computer (unless globally banned or banned for that computer's policy) but continue to have a *global* state of unapproved.

While especially convenient for disconnected computers, a timed policy override may be used for a connected computer. However, the override procedure disconnects the agent during the override. The override is maintained until the designated time period expires, even if the agent or computer is restarted during this period.

You can specify a duration of up to 500 minutes for the Enforcement Level change. If you specify 0 (zero) minutes, the override never expires (and the computer remains disconnected) until you reset it with another override.

If you want to change the duration or Enforcement Level of an override, you can create and apply a new override key. For example, if you want to end an override sooner than the original time period, you can specify a new override that is one minute long.

Caution

If you use a Temporary Policy Override Code to switch a computer's Enforcement Level to *Low* or *None (Visibility Only)*, when the agent transitions back to its original Enforcement Level, it might locally approve certain unapproved files discovered on that computer while in the more relaxed Enforcement Level – this affects files with Local State Details of *Unapproved*, and depends on whether *Locally approve unapproved files on transition from Visibility or Low Enforcement Level to Medium or High* is checked in the Advanced Settings for the policy that computer is assigned to. App Control recommends that unless you are certain that this automatic local approval setting is **off**, you only use the Enforcement Level override feature for temporary transitions to *Local Approval*, *Medium*, or *High Enforcement*.

To generate a code to place a computer in temporary local approval mode:

1. On the console menu, choose **Assets > Computers**. The Computers page appears:
2. In the table, locate the computer for which you want to generate a code and click on its name. The Computer Details page for that system appears.
3. Click the **Policy Override** tab in the panel at the bottom of the page.
4. In the Temporary Policy Override Code panel, unless you want to transition to a different Enforcement Level, leave the default choice for *Temporary Enforcement*, which is **Local Approval**.
5. In the *Enforcement Level Active For* box, enter the number of minutes (up to 500) you want the Enforcement Level change to last.

6. In the *Key Valid For* box, enter the length of time you want the override code to be valid. Your choice for this field should take into account how long it will take to get the key to the computer user who needs it and how quickly they will be able to enter it.
7. When you have entered all parameters, click the **Generate Code** button. A code with nine sets of letters separated by dashes appears in the box next to the button.
8. Copy and save the code from the box (and note the computer name) so that you can deliver it to the person who will be installing new software on the offline computer. The code is *not* saved on the Computer Details page, so you must record it.

Important

Because they provide privileged access to App Control Agents, do not email or otherwise reveal override keys in clear text. They should be encrypted if emailed, and only the administrator who needs them should have access. Once used, any file or other record containing the key should be deleted.

Computers need not be disconnected from App Control Server before an override is initiated. If the agent is connected to the App Control Server, the override procedure automatically disconnects it and then reconnects it after the override period is over. Machine reboots or agent restarts do not cancel the timed override.

The location and syntax of the override command is slightly different on different platforms.

To use a Timed Policy Override code on a Windows computer:

1. On the computer you want to apply the override to, open a command window and change to the App Control Agent installation, which by default is:
 - `c:\Program Files (x86)\Bit9\Parity Agent`
2. Enter the following command, either with the override code as the first argument or by itself:
 - `TimedOverride.exe <code>`
 - or-
 - `TimedOverride.exe`
3. If you don't provide the code in the command line, an authorization dialog box appears. Enter the override code for this agent into the dialog box and click **OK**. For either method:
 - If the code entered is invalid or expired, or if `TimedOverride.exe` is unable to communicate with the App Control Agent for any reason, an error message will be displayed. After three invalid attempts, the program automatically closes.
 - If a valid code is entered and the Enforcement Level transition is successful, no message is displayed but the dialog box closes. If the computer was connected to the App Control Server, it is disconnected.

To use a Timed Policy Override code on a Mac and Linux computer:

1. On the computer you want to apply the override to, open a terminal window and change to the following directory:
 - On Linux, `cd /opt/Bit9/bin`
 - On Mac, `cd /Applications/Bit9/Tools`

2. Enter the following command with the override code you generated as an argument:

```
./b9cli -timedoverride <code>
```

- If the code entered is invalid or expired, an error message will be displayed. After three invalid attempts, the program locks out further attempts for an hour or until the agent is restarted.
- If a valid code is entered and the Enforcement Level transition is successful, the message *Timed override set* is displayed.

When the override is set, the agent is disconnected from the server (if connected) and has the new Enforcement Level specified by the key. If the override code specified Local Approval, you can begin installing new software needed on this system and it will be locally approved unless already banned or approved.

When the configured override period expires, the following actions happen:

- The Enforcement Level returns to its previous setting.
- If the computer was connected when the override code was applied, it is reconnected to its App Control Server.
- When reconnected (whether immediately or at a later time), the agent reports events associated with the Enforcement Level change to the server.

If the computer is off or rebooting when the override expires, these actions occur when it is running again.

Marking a File as an Installer/Not an Installer

When it analyzes a file, App Control determines whether the file is likely to be an *installer* – that is, whether it will generate additional files when executed. By locally approving a file identified as an installer, you make any files it installs locally approved as well. Files not identified as installers do not transfer their approval status to files they generate, if any.

It is possible that a file is mis-categorized, or that you prefer not to have the local approval of a top-level file cause local approval of the files it installs. You can override installer status in both directions using menus on the file details pages. For each file, you see only the menu choice that reverses the current status.

Note

For this release, no Linux files are recognized as installers. The only Mac files recognized as installers are packages – files with .PKG extensions and properly defined *archive* headers. Because of this, using the *Mark as installer* feature might be particularly useful for these platforms.

To mark a file as an installer:

- On the File Details or File Instance Details page, click **Mark as Installer** in the Actions menu.

To mark a file as not an installer:

- On the File Details or File Instance Details page, click **Mark as Not Installer** in the Actions menu.

Notes

- When you override the installer status of a file, that override is shown in the Local State Details for the file.
- In file tables, if you check the box next to a file *not* identified as an installer, and you choose Approve by Policy on the Action menu, you can mark the file as an installer as part of your approval rule. This ensures that new files it writes will be locally approved. Files it has already written will remain in their current state.
- You can create a Custom Rule that *Promotes* files meeting the rule specifications. This treats these files as installers under the conditions of the rule but does not change their global status as an installer or not an installer. See [Chapter 14, “Custom Software Rules.”](#)

File-Specific Rules: Approvals and Bans

The Files tab of the Software Rules page shows all of the approvals and bans created at your site for specific individual files. These rules identify specific files by hash or optionally by file name (for bans only).

Approvals and bans can be global, applying to all computers, or they can be applied to computers in selected policies. Active Bans block file executions for affected computers in Control mode, report an event for computers in Visibility mode, and do nothing for computers in Agent Disabled mode. You also can create a Ban that only reports what it would have done if active.

Unified Management

File Approvals and Bans can be centrally managed for multiple servers through the Unified Management feature. This is described in [“Unified Management of Rules”](#) on page 791.

Because the Files tab shows both Approvals and Bans, you can manage all file rules in one place. You can check to see whether a particular file has any approval or ban affecting it, and you can remove rules from one or more checked files.

By default, file rules are grouped by their *type*, so you see all of the Approvals together, Bans together, and Report Only bans together. As with most console tables, you can change (or eliminate) the grouping by making another choice on the *Group by* menu.

You can create approvals and bans directly on the Software Rules page Files tab if you want to enter the file hash or name manually in a property page. The easier way to create bans, however, is from a table or File Details page that already has the file hash in it. In either case, when you create the approval or ban, it appears on this page.

When you create a new ban or approval, it might affect a file that already has an approval or ban. If you attempt to do this, a warning appears, informing you that if you save the new rule it will delete the old rule. This can be especially helpful if you select a group of files and are accidentally replacing a ban with an approval on some files, or vice versa.

In some cases, creating a ban not only prevents future executions of a file but stops any currently running processes matching that file. See [“Enabling Bans to Stop Running Processes”](#) on page 312 for more details.

Beginning with version 8.1.0, you can delete files on endpoints using CB Console commands. See [Chapter 9, “Deleting Files,”](#) for details.

Note

Approvals and bans on the Files tab are rules created specifically for a given file (by name or by hash). This page does not show *all* approvals or bans that take effect because of other rules, including Reputation and Custom Rules, and it is not a comprehensive list of global *file state*. If you want to see all files whose *global state* is approved, use the File Catalog.

Approvals and bans that appear on the File Rules page are created in the following ways:

- From the Software Rules Files tab, open the Add File Rule page and enter the hash for a single file; for bans, you also can use the file name or a specific path.
- From a File Details or File Instance Details page, choose one of the approval or ban commands on the Actions menu to create a rule for a single file.
- In a table of files (e.g., the File Catalog), check one or more files and choose one of the approval or ban commands on the Action menu to create one or more rules.
- In the Events table, check one or more events that have a file reference in the description and choose one of the approval or ban commands on the Action menu to create one or more rules.
- From the Software Rules Files tab, import a list of file hashes to create multiple rules.
- From the Software Rules Directories tab, create a Trusted Directory. Each file located in a trusted directory has an approval rule created for it.
- An approval or ban might be created through an external API. Rule origin also might be unknown, for example if the rule was created in an older version of App Control (including Bit9 Platform or Parity). The *Source* field on the Files tab or Edit File Rule page shows how a rule was created.

Once you create a rule, you can manage it from the File Rules page, and in most cases you can delete it using commands on the page you used to create it.

Caution

Banning the wrong file can have unintended and possibly harmful consequences. For example, inadvertently banning a legitimate system file could cause computers to immediately crash. Before you ban a file, ensure that you enter the correct name or hash. As a precaution, first search the file name or hash with the Find Files feature to verify that it is the file you want to ban, and review the File Details page. For further assurance, consider using Carbon Black File Reputation to learn more about the file before banning it. For more information, see [Activating Carbon Black File Reputation in Chapter 26, "System Configuration."](#)

One way to test the impact of a ban without actually blocking files is to create a Report Only ban.

Testing a ban through Report only is especially advisable if you have enabled termination of running processes when bans are created. See ["Enabling Bans to Stop Running Processes"](#) on page 312.

Report Only Bans

Creating a *Ban (Report Only)* rule enables you to observe how a ban might affect your users. With a report-only ban, the file is not blocked but *would-have-blocked* and *would-have-terminated* warnings are written to the Events log. If you are certain this is a file you want to block from executing, you can change the rule to a full Ban. See ["Event Reports"](#) on page 590 for more information about App Control event reporting.

Approvals and Bans of MSI Files by Hash

If you plan to approve or ban an MSI file, extra care must be taken to ensure that the right file is identified. MSI files include timestamps that typically are unique on each endpoint, and this causes the hash of the same MSI file to be different on every machine. This affects hash approvals and bans of MSI files in two ways:

- To provide identical hashes for functionally identical MSI files, the App Control Agent creates a "fuzzy" SHA-256 hash of each MSI file, without the unique timestamps. This means that you can apply rules to all copies of an MSI if you use a SHA-256 hash created in App Control. However, if you import or copy a SHA-256 hash for an MSI from another source, rules that use the imported MSI are unlikely to be effective in App Control.
- For MD5 and SHA-1 hashes of MSI files, the App Control Agent hashes the whole file. This means that because of different timestamps, each copy of a functionally identical MSI file will have a different hash, and file rules that use an MD5 or SHA-1 hash of an MSI file are unlikely to work as expected.

Because of these issues, the best practice for approving or banning an MSI file is to use the SHA-256 hash created in App Control. Other hash types, and hashes imported from elsewhere, should be avoided in this case.

Creating an Approval or Ban from the Software Rules Page

If you want to specify all of the parameters for an approval or ban, you can create it on the Add File Rule page.

To create and configure an approval or ban for a single file:

1. On the console menu, choose **Rules > Software Rules**. The Software Rules page appears.
2. Click the **Files** tab. The File Approvals and Bans table appears:
3. Click the **Add File Rule** button. The Add File Rule page appears, with Approval as the default Rule Type.
4. Specify the rule details and the file to be approved or banned ([Table 41](#) shows the full list of possible parameters as well as rule information available after creation):
 - a. Provide a Rule Name so that you can identify the rule in the table.
 - b. Choose the Rule Type (*Approval, Ban, Ban (Report Only)*). Note that if you choose Ban, a warning appears stating that the Ban could stop matching files currently running. See [“Enabling Bans to Stop Running Processes”](#) on page 312 for details.
 - c. If the rule is a Ban, choose the Type (*Hash or File Name*).
 - d. For Hash rules, specify the type of hash you will provide (*MD5, SHA-1 or SHA-256*).
 - e. For File Name Bans, choose the platform to which the rule will apply (*Windows, Mac, or Linux*).
 - f. Enter the Hash Value or File Name that will identify the file.
 - g. Optionally, provide a Description.
 - h. In the Rule Applies To field, choose *All policies* or specify the *Selected policies* to which the rule will apply.
5. To create the approval or ban, click **Save**. The rule appears on the File Rules table. Group the table by Type (the default) if you want to see Bans together, Report Only bans together, and Approvals together.

When you save a rule, the parameters that define the rule and additional information about it are available on its details page. [Table 41](#) shows the details that can appear on the Edit File Rule page.

Unified Management Note

Additional options appear on the Add File Rule page if you have enabled Unified Management of multiple servers. See [“Unified Management of Rules”](#) on page 791.

Table 41: File Rule Parameters

Field	Description
Rule Name	Text description of the files to be approved or banned. This could be a file name or other identifying information to help you manage the rule (the rule is created even if you do not enter a name). Note: This is name for the rule only. Entering a file name here does not create a filename-based rule.
Rule Type	The choices are Approval, Ban, and Ban (Report Only), which reports events for situations in which the file would have been blocked if the rule had been a full Ban.
Source (Read Only)	The source type of the most recent modification of the rule. This could indicate how it was created or a later change. The possible values are: Manual (all manual or from Action menu commands), Trusted Directory, Reputation, Imported (from an uploaded list of files), External (API), Event Rule, Rapid Configs, Unified Management, and Unknown. Appears after the rule is created.
Source Name (Read Only)	The name or additional description for the source of the most recent modification of the rule. This will either be Carbon Black Installation, Trusted Directory (Deleted), or a rule name. Manual, Imported, and Unified Management rules do not have a Source Name.
Type (Bans Only)	To ban a file you must know the Name of the file or its Hash (data signature). Choose one, as appropriate. If you choose Name, you can enter a path so that the rule only applies to a file in a particular location. Approvals are always by hash, so the Type field does not appear for them. Name bans must be platform-specific.
File Name (Bans Only)	(Appears only for bans, and only if you chose File Name as Type) Name of the file and its extension. For example, msblast.exe. Specify a directory path if you want to ban only matching files in a particular location. If you use a path, files with the same name that appear in any other directory are not subject to the name ban. Platform Note: If you enter a path, be sure to use the correct directory delimiters, and to use only characters and formats legal for paths in the chosen platform. The App Control Server does not convert paths between platforms (e.g., '\ to '/'). Also, Linux file names normally are case sensitive.
Platform (Ban by Name Only)	(Appears only for bans, and only if you chose File Name as Type) Platform for which this rule is effective (Mac, Linux, Windows). Name bans must be platform-specific.
Hash Type	Cipher algorithm used to create the hash you want to approve or ban. If you paste in a value, the choices are MD5, SHA-1, and SHA-256. Rules created from a file table or details page use SHA-256, if available.
Hash Value	Hash (data signature) for the file. Hashes not yet seen by this App Control Server can be used in rules. To locate hashes for files already found on your computers, you can use the File Catalog or Find Files pages.

Field	Description
Description	Optional text to further describe the file approval or ban. This information is displayed in File Rules table under the Description column (if visible).
Rule Applies To	Policies for which the approval will be enforced: Select All policies to approve or ban the file for all computers. Select Specified policies to choose which policies to apply the rule. When you click this button, a list of policies appears, each with a checkbox. You also can use the checkbox at the top of the list to check all boxes or clear all checks, but keep in mind that you cannot create a rule that applies to no policies.
History (Read Only)	Shows when and by whom the rule was created and last changed. Also shows the CL version (i.e., the version of App Control rules) in which the current version of the rule is present, which can be used to determine whether the rule is present on an agent.

Editing and Deleting File Rules

You can modify or delete an existing File rule. In [Table 41, “File Rule Parameters”](#) on page 305, some of the parameters can be changed and some are read-only.

To edit an approval or ban rule:

1. On the Files tab of the Software Rules page, click the View Details button next to the rule. The Edit File Rule page appears.
2. Edit the details you want to change. You can change all rule parameters *except for* Type (hash or file), Hash Type, and Hash Value. Also Source and History are read-only fields added to the page to reflect activities related to the rule.
3. When you have finished making changes, click **Save**. The rule is updated.

Note

You cannot disable an existing approval or ban. You can, however, change the Rule Type. For example, you can change a ban from an active ban to Report Only, which will prevent it from blocking but still report file executions it would have blocked.

You also can change a Ban to an Approval or vice versa, but be certain you understand the effects before doing this. If you don't want a rule enabled in any way, you must delete it.

To delete a File rule, you can use the **Remove Approval or Ban** commands on the Action menu of any file table page, or the appropriate *Remove* command on a details page. If you are on the Software Rules page Files tab, you delete rules using the following procedure.

To delete one or more approval or ban rules:

1. On the Files tab of the Software Rules page, check the box next to the approvals and bans you want to delete.

2. Click the **Delete File Rule** button.
3. In the confirmation dialog box, click **OK**. The rules are removed.

You also can delete a single approval or ban by clicking the **Remove Rule** button on its Edit Rule page.

Creating File Approvals and Bans from Table Pages

The following procedure describes creating an approval or ban rule from the Files page (File Catalog or Files on Computers), but it applies to any other console page that lists files as well as pages in which the file is not the primary information but might be included as a link in details of another object. Generally, a row with a checkbox next to a filename allows creation of bans and approvals from the Action menu. This includes:

- Files page (both File Catalog and Files on Computers)
- Baseline Drift Report Results pages that list files
- Snapshot Content page
- Events page (only events that include file hashes)
- Find Files page (when showing results)

The Action menus on table pages provides the following choices for managing approvals and bans:

- **Approve Globally** – Immediately creates a hash-based rule globally approving a file for all computers – no configuration is necessary.
- **Ban Globally** – Immediately creates an active hash ban applying to all computers – no configuration is necessary.
- **Approve by Policy** – Opens the Add Rule page with the file name as Rule Name, Approval as the Rule Type, and the file Hash already in place. You can choose to apply the rule to selected policies or all computers and, you can edit the rule name and add a description.
- **Ban by Policy** – Opens the Add Rule page with the file name as Rule Name, Ban as the Rule Type, and the file Hash already in place. You can choose to apply the rule to selected policies or all computers, you can edit the rule name and add a description, and you can make the rule an active ban or just a report-only ban.
- **Remove Approval or Ban** – Immediately removes the rules for all checked boxes, including mixed selections of approvals and bans.

Unified Management Note

Additional options appear on this menu if you have enabled Unified Management. See [“Unified Management of Rules”](#) on page 791.

The advantage of creating an approval or ban from a console files table is that you can approve or ban multiple files at once. For example, you might use the filtering tools on a files page to get a list of files meeting certain criteria, check the box next to each file’s name, and globally ban them in one operation.

When you create a rule from a table, the rule definition you provide applies to each selected file. When you save the definition, a separate rule is created and named for each selected file. Rules created from checked rows of a table are always hash bans, and use SHA-256 hashes if available.

Notes

- Initially, files that originate from a common source or installer are grouped under the source/installer file name. If you are looking for a file to approve or ban and want to include all *individual* files grouped under an installer in the table so that you can view and search them, check the **Show Individual Files** box in the lower right corner of the Files page, which automatically refreshes the table.
- You can filter the lists of files on the Files page, rearrange display columns, and download results in comma-separated-value format. For more information, see [Console Tables](#) in [Chapter 2, “Using the Console.”](#)

Creating Global Approvals and Bans

The Action menu on files pages has two shortcut commands, one of which creates a global ban and the other a global approval for the files you check on the page. These commands give you a quick way to approve or ban one or more files as long as you do not want to create any special configuration for the rules you create.

When created this way, rules apply to all policies. If you choose *Globally Approve*, checked files are globally approved for all computers and each file has a separate approval rule on the Software Rules page. Likewise, if you choose *Globally Ban*, the files are banned on all computers in Control policies and each file has a separate ban rule on the Software Rules page.

For both approvals and bans, if you checked one file, the file name is used as the rule name. If you checked more than one file, the name is left blank.

Notes

If you select files that already have a rule and apply a different type of rule to them, it is possible that the name of the old rule will be maintained and the rule type will be changed. This could be confusing if you named a rule something like “Approve Files for My Project” and then changed the Rule Type to *Ban*.

To create a global approval or ban for one or more files on a Files page:

1. On the console menu, choose **Assets > Files**. The Files page appears.
2. Locate the files you want to approve or ban and check the boxes next to their names.
3. On the Action menu, choose **Globally Approve** or **Globally Ban**.
4. In the confirmation dialog box, click **OK**.

Custom Approvals and Bans

When you choose Approve by Policy or Ban by Policy on the Action menu of a file table, an Add File Rule dialog appears with the hash(es) for the files you selected already entered. Unlike choosing one of the global options, this choice allows you to customize other parameters before you create the rule.

To create a custom approval or ban for one or more files on a Files page:

1. On the console menu, choose **Assets > Files**. The Files page appears.
2. Locate the files you want to approve or ban and check the boxes next to their names.
3. On the Action menu, choose **Approve by Policy** or **Ban by Policy**. The Add File Rule page opens.
4. You can change the Rule Type, including changing from **Ban**, which actively blocks executions, to **Ban (Report Only)**, which just reports that the file would have been blocked if the ban was fully activated.
5. You can add an optional description of the rule (for example, something the approved files have in common or why you banned the files on them).
6. In the *Rule applies to* field:
 - a. To apply the rule to all computers, leave the *All policies* button selected.
 - b. To apply the rule to selected policies only, click the *Selected policies* button.
7. If the Rule Type is Approval, an Installer Information panel is included at the bottom of the page. If any of the files selected for approval is not currently recognized as installers, a *Mark all files as installers* checkbox appears in the panel. Check the box if you want the files to be approved and marked as installers.

Important

Especially when you have multiple files selected for the rule, be certain you want *all* of the files to become installers before you check the *Mark all files as installers* box. Files created by installers are locally approved, and there is no automatic way to remove this approval. The message in the Installer Information panel will tell you how many files in your selection would be affected by this choice, and whether any files in the selection have created or modified other files.

8. When you have configured the rule as you want it, click the **Save** button. Each file you checked when you started the process appears on the Software Rules page Files tab as a separate approval. The File Approvals and Bans table indicates whether an approval or ban is global or not.

Warnings when Creating or Editing Bans

When you create or edit a ban, the File Rule details dialog will show a warning in red, indicating that the rule could stop currently running files. This appears as a reminder even if you have not enabled process termination in any policy.

In addition, when you add or edit a file ban and click Save on the File Rule details dialog, a confirmation dialog may provide a further warning. The warning appears if a name ban contains wildcards in the name. It also appears for both name and hash bans if the file specified in the rule has a Carbon Black File Reputation threat level of either "0 – Clean" or "Unknown" *and* one of the following conditions is true:

- A ban specifies a file signed by Microsoft (including key system files)
- A ban specifies a file signed by another trusted publisher
- A ban specifies a file with Carbon Black File Reputation trust levels above 7.
- A ban specifies a file that appears on more than 10% of reporting agent computers.

In each of these conditions, terminating the file or files indicated in the ban could have undesirable effects, including shutting down the computer. The default on this dialog is to allow the ban, so be sure to click on **Cancel** if you have any concern about the ban.

Approving and Banning Files from the File Details Page

Although you can approve or ban files from tables, you might want more information about the file before you decide to ban it. For this, you can go to the File Details page.

Note

You can follow this same procedure to approve or ban a file globally or by policy from the File *Instance* Details page, which also includes options for applying or removing local approval of an individual file.

To approve or ban a single file using the File Details page:

1. When you find a file to approve or ban, click the View Details button next to it in a table or click its hash or name if it is in the Events table. The File Details page appears.
2. Examine the information on the File Details page to be certain you want to approve or ban the file. For example, you can see in the File Prevalence line whether any computers currently have the file. To determine which computers have the file before you approve or ban it, click the **All File Instances** link on the Related Views menu.
3. If you have Carbon Black File Reputation enabled, the Carbon Black File Reputation Information panel shows Trust, Threat, and other information about the file, if available. You can click the **Analyze** button to search Carbon Black File Reputation for information if none is shown or to check for updated information.

Note

If you want to analyze the file but the Analyze button is not visible, see [“Activating Carbon Black File Reputation”](#) on page 763.

4. In the Action menu, choose the rule you want to create for this file – note that if the file is already approved or banned, you must remove the current rule (using **Remove Approval** or **Remove Ban**) before you create an opposite rule.

Note

For more information about approving or banning hashes from the Files tab of the Software Rules page, see [“Creating an Approval or Ban from the Software Rules Page”](#) on page 304.

Approving or Banning Lists of Files

If you have a list of hashes for files, you can import the list in a text file as input to the console and change the state of these files in one operation. You can change the file state to Approved, Banned, or Ban (Report Only), and you can do this for some or all policies.

The requirements and recommendations for approving or banning lists of hashes are:

- The file containing the hash list must be accessible to the App Control Server.
- The file must contain a list of MD5, SHA-1, or SHA-256 hashes, with only one hash per line.
- Use only one hash *type* per file; mixing types in one file may cause unpredictable results.
- You must take the same action on all files on the list; that is, you must approve the whole list, ban the whole list, or create a report-only ban for the whole list.
- On some older versions of IE with Advanced Security Settings, you must make the root URL of your App Control server, <https://<your server name>/>, a trusted site in **Internet Options – Security – Trusted Sites – Sites**. Otherwise, bulk hash files cannot be processed. See the *Operating Environment Requirements* for this release for a list of supported browsers.
- Do not navigate away from the page until the Upload Hashes page shows that the process is complete. If you do navigate away, processing of the hashes is interrupted. In this case, you can upload the file again, and any hashes not yet approved or banned will be processed.

When you use this method to approve or ban a list of files by their hashes, each file appears as a separate rule, but the rule name is the same for each.

To create approvals or bans for a list of hashes:

1. Copy or move the file containing the hashes to a location accessible to the App Control Server.
2. In the console menu, choose **Rules > Software Rules**.
3. Click the **Files** tab. The File Rules page appears with a list of Approved and Banned files.
4. Click the **Import** button. The Upload Hashes for Banning or Approving page appears.
5. Enter the rule parameters, as follows:
 - a. Enter the Rule Name as you want it to appear on the File Rules page.
 - b. Use the **Browse** button to locate the file containing the list of hashes and click **Open** in the *Choose file* dialog when you locate the file. The pathname to the file containing the hashes appears in the File name box.
 - c. (Optional) Enter a description for the rule.
 - d. Choose **Approve**, **Ban**, or **Ban (Report Only)** on the Rule Type menu.
 - e. Make the rule effective for **All policies** or **Selected policies**.
6. When you are satisfied with all of the rule parameters, click **Upload**. A two-column progress table appears as the hashes are processed, reporting the success or failure of the rule for each file and also informing you when hashes on the list are already in the state you chose.
7. On the console menu, choose **Rules > Software Rules**. On the Files tab of the Software Rules page, the hashes you created approvals or bans for appear in separate rows in the table, but with the same Rule Name. Once rules have been created for all files on the list, each rule can be modified individually.

Enabling Bans to Stop Running Processes

By default, file bans stop future attempts to execute a file but do not terminate processes that are already running on an agent-managed system. This means that files that are allowed to run but are later determined to be malicious will continue to run unless they are terminated for some reason other than a App Control rule, or if the system restarts. This is especially likely in Low and Medium Enforcement policies, where files not explicitly banned are allowed to run.

You can configure policies so that computers in those policies stop currently running software when they receive a rule that bans it. This capability provides better control over software in your environment. It must be used carefully, however, to avoid interrupting important processes or even preventing a computer from running at all. Also, keep in mind that when enabled for a policy, process termination applies to *all* banned files. So that you can see what the effect of this setting might be, newly created policies are configured to report processes that would have been terminated by a ban, but not to actually terminate them.

Notes

- Termination of processes with banned images is supported on Windows agents only.
- You can also delete files on endpoints using CB Console commands. See [Chapter 9, "Deleting Files."](#)

Any ban, whether on a system that terminates banned processes or one that doesn't, may disrupt a user's system or cause other dependent applications to fail, possibly causing loss of work in progress. On the other hand, allowing bans to terminate running processes provides immediate feedback on the results of the ban. They also make it possible to terminate legitimate processes infected with malware and allow them to restart without the infection. The following are some examples of the potential impact of enabling process termination:

- **Discrete Single Application** – Ban `skype.exe`. On systems affected by the ban, all running instances of Skype are abruptly terminated and any attempt by users to restart Skype will be blocked.
- **Windows Explorer Extension** – Ban a file called `malware.dll`, which is registered as a Windows Explorer extension and is present in all running instances of Explorer. On systems affected by the ban, all instances of Explorer are terminated and then the Explorer is automatically restarted by Windows. On restart, the banned file `malware.dll` is blocked while Explorer continues to load and run, so the ban prevented the unwanted process from running without blocking the critical Explorer process. Without the terminate process setting, the unwanted process would continue to be running in every active Explorer, even after it was banned.
- **Dynamically Loaded DLL** – Ban `wsock32.dll`. Also assume that `wsock32.dll` is dynamically loaded by the application `xyz.exe` when it needs to perform certain network operations and then unloaded when the operation is complete. On systems affected by the ban, if the file `wsock32.dll` is banned while unloaded, it will be blocked the next time it is loaded by `xyz.exe`, likely causing the operation to fail. If the ban takes effect when the file is loaded, the process `xyz.exe` will be terminated.

- **Shared Service** – Ban malware.dll, which is installed as a network service and shares an instance of svchost with other running services. When the file is banned, the instance of svchost is terminated along with all services in the same process.
- **Injection in Critical Process** – Ban malware.dll, which is injected into csrss.exe, a critical system process. On systems affected by the ban, csrss.exe is terminated. Windows detects the termination of critical system processes and immediately shuts down. If the csrss.exe is reloaded again on startup, App Control prevents the image from being injected and allows the system to boot normally without malware being installed.
- **Boot-time Driver** – Ban malware.sys, which is installed as a boot-time driver. If the driver loads before the App Control Agent does, it can continue to be executed and may not be stoppable without crashing the machine. Going to safe mode to remove the infection or restore to an earlier time may be the only remediation.

Keep these and other possible effects in mind when considering whether to enable process termination in a policy.

To enable or disable immediate termination of banned processes:

1. On the console menu, choose **Rules > Policies**.
2. Double-click the View Details button next to a policy for which you want to configure process termination.
3. On the Edit Policy page, click the **Show Advanced Settings** button.
4. In the Advanced Settings panel for the policy, go to the last setting, *Terminate processes with banned images*. Choose one of the following items on the Status menu:
 - **Off** – Creation of a ban does not terminate or report that it would have terminated a running process.
 - **Report Only** – Creation of a ban does not terminate a running process but reports that it would have if this setting was Active.
 - **Active** – Creation of a ban terminates a running process matching the banned image.
5. Click the **Save** button above the Advanced Settings panel.
6. Repeat these steps for any other policies in which you want to change this setting.

Chapter 9

Deleting Files

This chapter describes how to use the App Control Console to delete files on Windows endpoints.

Sections

Topic	Page
Overview	315
Requesting File Deletion from Table Pages	317
Requesting File Deletion from Details Pages	320
Automating File Deletion Requests	321
Monitoring File Deletions	324

Overview

App Control provides many ways to block and ban unwanted activity on endpoints. In some cases, you might want delete a file entirely. (You can delete files on Windows endpoints through the Console.)

The commands for deleting files are available in the following locations:

- the File Catalog page
- the Files on Computers page
- the Find Files page
- the File Details page
- the File Instance Details page

In addition, you can create Event Rules that will automatically delete files when certain events occur, such as a report of a malicious file.

The user on an endpoint is not notified when a file is deleted by App Control.

Scope and Limits

Depending upon the page you are on, you can delete instances of a file on one computer or all current instances of the file on all Windows computers managed by the current App Control Server. On pages showing tables of files, you can also select multiple files for deletion at the same time. Files are identified by hash, not name, so that (for example) malicious files that rename themselves do not escape deletion.

Deletion of files from the App Control console has the following limitations in scope:

- Only files that are in the inventory at the time a delete request is made are deleted. Instances that appear later are unaffected.
- Deletions are applied only to the agents reporting to the server you are logged in to. They cannot be applied to remote servers in a Unified Management environment.
- The file deletion feature applies only to Windows endpoints running the 8.1.0 agent or later. Files cannot be deleted from Windows endpoints running earlier agent versions. When VMware Carbon Black App Control v8.5.0 was released, file deletion was not supported on Linux or Mac agents; monitor the User Exchange for any updates to support status.

Permission to Delete Files

To access file deletion features, a user must have a login account with a User Role that includes the *Delete files* permission. Permission for these features is not granted by default to the *admin* account or members of the Administrator account group. You must explicitly add it, either by enabling *Delete files* permission for an existing role or by

creating a new role for file deletion and adding it to accounts as you choose. See [Chapter 3, “Managing Console Login Accounts,”](#) for more on enabling permissions in an account.

Important

This feature should be used with extreme care, and in full compliance with your organization's policy on accessing other users' files. Be sure that only those App Control Console users that absolutely need access to the file deletion feature are given permission to use it.

Timing of File Deletion

When you issue a command to delete files, the request is sent to the agents the next time the server and agent communicate, normally at 30 second intervals. Depending on the number of files and agents involved, actual deletion of files and reporting those deletions back to the server can take from approximately 30 seconds to several minutes.

Requests to delete files on offline computers remain active until the files are actually deleted. These files will be deleted when their endpoints reconnect to the server. However, new file instances discovered *after* a delete request is processed will not be deleted.

During the time between a deletion request and actual deletion of a file on an endpoint, other actions related to the file might take place. If a user moves a file during this period, it will still be deleted from the server unless it is moved to a location in which file tracking is disabled, for example, because of a Performance Optimization rule.

A server-initiated file delete command can only delete files that are currently tracked in agent's inventory. Deletion of a file by the endpoint user (i.e., moving it to the Recycle Bin) might prevent successful completion of a server-initiated file deletion. By default, the agent only tracks a file in the Recycle Bin if the file has been *executed* from the Recycle Bin. For files in the Recycle Bin that have *not* been executed there, a server-initiated delete action will fail and report that failure in an event.

You can delete a file that has been scheduled for upload to the App Control Server or upload and analysis by a third-party network security service, and it is possible that the deletion will occur before the upload or analysis. Keep this in mind if you want to analyze a file before choosing to delete it.

Files Protected from Deletion

Deletion of certain files could prevent proper operation of your endpoints. To minimize this possibility, some files are protected from deletion via the console, including files identified as operating system files and files necessary for operation of the agent. You can request

deletion of these files, but they will not be deleted on endpoints, and the Events page will show a *File deletion failed* error event.

Caution

When you issue a delete command, you are asked to confirm the decision, but once you do, there is no undo or restore available. The files VMware Carbon Black App Control excludes from deletion are probably not the only files critical for proper operation of your endpoints. Keep this in mind when choosing to delete a file.

Also, the delete file commands do not completely “wipe” a file from a system. This could lead to system instability, including crashes or failure to start an application, if there are still references to the file. For example, if a process that does implicit linking depends on a DLL that has been deleted, the application will not start. It is even possible that malware that you want to delete has registered other services as being dependent upon it, making these services unable to start when the malware is deleted.

Requesting File Deletion from Table Pages

The Action menu on the File Catalog, Files on Computers, and Find Files page provides the following choices for deleting files:

- **Delete file instances** – On the Files on Computers and Find Files pages, this command operates only on the computers associated with the files you checked. It submits a request to remove all instances of files matching the hash of the selected files *only from the computers associated with the file instances you checked*. Instances of the same hash on computers that were not selected will not be deleted.
- **Delete files from all computers** – On the File Catalog page, this submits a request to delete all instances of files matching the hash of the checked files in the table *on all computers reporting to the current server*.

The advantage of executing delete file commands from a console table is that you can delete multiple files at once. For example, you might use the table filtering tools to get a list of files meeting certain criteria, check the box next to each file’s name, and request that they be deleted in one operation. However, this also requires carefully checking the list of files before you execute the command so that you avoid deleting a file by accident.

Note

For more information on filtering the results shown on table pages, see [Console Tables](#) in [Chapter 2, “Using the Console.”](#)

To request deletion of selected file instances via the Files on Computers or Find Files page:


1. On the console menu, choose either **Assets > Files** or **Tools > Find Files**.
2. If you are on the Files page click on the **Files on Computers** tab.

3. Search or filter the list to locate the files you want to delete and check the boxes next to their names.
4. On the Action menu, choose **Delete file instances**.

The screenshot shows the 'Files on Computers' view in the File Catalog. The 'Action' menu is open, and 'Delete file instances' is selected. The table below shows a list of files with their hashes, names, and publishers.

File Hash	File Name	Publisher or Company	User Name
79895	cygevent_extra-2-0-5.dll		
79895	cygevent_openssl-2-0-5.dll		
79895	cygevent-2-0-5.dll		
79895	cygevent_core-2-0-5.dll		
79895	cygedit-0.dll		
79895	c++filt.exe		
79895	cygdb_sql-5.3.dll		
79895	cygdb_cxx-5.3.dll		
79895	cygdb-5.3.dll		
79895	cygcurl-4.dll		
79895	cygcheck.exe		

5. Examine the information in the confirmation box, and if you want to proceed with the deletion, click **OK**.



A banner indicating that the delete request was successfully created should appear. Each instance checked is scheduled for deletion from the computer on which it was found, and all files with the same hash on the same computers are also scheduled for deletion.

You can check the Events page to confirm that the files are actually deleted or see any error messages related to the action, such as failure to delete a file that cannot be deleted or failure to delete a file that no longer exists. When one deletion request affects multiple instances of a file, one event appears to report the outcome of the deletion attempt (success or failure) for each instance.

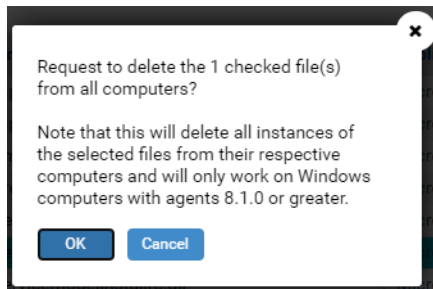
To request deletion of all instances of selected files via the File Catalog page:

1. On the console menu, choose **Assets > Files**. The Files page appears.
2. If File Catalog is not the current view, click on the **File Catalog** tab.

3. Search or filter the list to locate the files you want to delete and check the boxes next to their names.
4. On the Action menu, choose **Delete files from all computers**.

The screenshot shows the VMware Carbon Black App Control interface. At the top, there's a 'File Catalog' header with a sub-tab 'Files on Computers'. Below this, there are controls for 'Saved Views' (set to '(none)') and 'Group By' (set to '(none)'). There are also links for 'Show Filters', 'Show Columns', 'Show Snapshot', 'Export to CSV', and 'Refresh Table'. A dropdown menu is open over the 'Action' column, showing various options like 'Approve', 'Remove Ban or Approval', 'Ban', 'File Actions', and 'Delete files from all computers'. The 'Delete files from all computers' option is highlighted with a red box and a mouse cursor. The background table shows a list of files with columns for file name and publisher, with 'Microsoft Corporation' listed as the publisher for several files.

5. Examine the information in the confirmation box, and if you want to proceed with the deletion, click **OK**.



A banner indicating that the delete request was successfully created should appear. All files with the same hash as any of the selected files are scheduled for deletion. You can check the Events page to confirm that the files are actually deleted or see any error messages related to the action, such as failure to delete a file that cannot be deleted or failure to delete a file that no longer exists. When one deletion request affects multiple instances of a file, one event appears to report the outcome of the deletion attempt (success or failure) for each instance.

Requesting File Deletion from Details Pages

Although you can request deletion of files from file table pages, you might want more information about the file before you decide to delete it. For this, you can go to the File Details or File Instance Details page. The Actions menu on these pages allows you to delete the file from the current computer or all computers.

To request deletion of a file via the File Instance Details page:

1. When you locate a file you want to delete the Files on Computers or Find Files page, click the View Details button next to its listing in a table. The File Instance Details page appears.
2. Examine the information on the File Instance Details page to be certain you want to delete all instances of the file from the computer shown.
3. If you have Carbon Black File Reputation enabled, the Carbon Black File Reputation Information panel shows Trust, Threat, and other information about the file, if available. You can click the **View Carbon Black Reputation Data** button to search Carbon Black File Reputation for information if none is shown or to check for updated information.
4. In the Actions menu, choose to delete files from the current computer or all computers:
 - **Delete All Files Instances from <computername>** – This submits a request to delete all instances of the file on the named computer. Any other instances of the file remain in place.
 - **Delete File from All Computers** – This submits a request to delete all instances of the file on all computers reporting to the current server

File Instance Details ⓘ

Details for file on computer: WORKGROUP\APC-79895

File Name:	cygevent_openssl-2-0-5.dll
Date Created:	Jun 29 2020 07:41:16 PM
File Path:	c:\cygwin64\bin\
Computer:	WORKGROUP\APC-79895
Platform:	Windows
User Name:	(none)
Local State:	Approved
Local State Details:	Locally Approved
Detached Publisher:	(none)
Executed:	Yes
Present At Initialization:	Yes
Top-Level File:	No
Deleted:	No
Root File Name:	(none)

General

First Seen Name:	cygevent_openssl-2-0-5.dll
First Seen Date:	Jun 29 2020 07:43:16 PM
Last Updated:	Jun 29 2020 07:43:16 PM
First Seen Path:	c:\cyqwin64\bin\


Related Views

- All File Instances
- File Events
- Computers with this file
- Computers without this file

Actions

- Remove Local Approval
- Approve Globally
- Ban Globally
- Approve by Policy...
- Ban by Policy...
- Delete All File Instances from WORKGROUP\APC-79895**
- Delete File from All Computers
- Add Meter
- Add Alert

5. Examine the information in the confirmation box, and if you want to proceed with the deletion, click **OK**.



A banner indicating that the delete request was successfully created should appear. All instances of the file are scheduled for deletion on one or all computers. Check the Events page to confirm that the files are actually deleted or see any error messages related to the action, such as failure to delete a file that cannot be deleted or failure to delete a file that may have been deleted locally before it could be deleted by App Control.

To request deletion of all instances of a file from all computers via the File Details page:

1. When you find a file you think you want to delete everywhere, click the View Details button next to its listing in the File Catalog, or click its hash or name if it is in the Events table. The File Details page appears.
2. Follow the procedure described in [“To request deletion of a file via the File Instance Details page:”](#) on page 320. The only difference is that on the File Details page, the only choice is **Delete File from All Computers**.

Automating File Deletion Requests

You can create Event Rules that will automatically delete files when certain events occur, such as a report of a malicious file. Unlike deletions done using commands, Event Rule deletions have no confirmation dialog. If an event matches the rule, the files specified in the rule are immediately scheduled for deletion, without feedback on the console. You can, however, create Alerts that inform you when a file deletion is requested or completed.

Event Rules that delete files do offer you the same scope options as delete commands: you can either delete all instances of a file on one computer or you can delete all instances of the file from all computers.

To create an Event Rule that deletes files (example):

1. On the console menu, choose **Rules > Event Rules**.
2. On the Event Rule page, click the **Create Rule** button.
3. In the Rule Name field, provide a unique name for the rule. For example, you might name the rule *Delete Malicious Files*.
4. In the Description field, provide a longer description of the rule if you choose.
5. In the Status field, choose **Simulate only**. This means that actions specified by the rule will be simulated. Events will be generated indicating what the rule would have done if enabled, but the actions specified will not actually be taken.

Important

Simulate only is strongly advised for initial testing of a new event rule, especially for a *Delete file* rule. Rule status can be changed to *Enabled* when you are sure that it the rule does not have any negative effects. See [“Testing a Rule before Enabling”](#) on page 524 for more about this choice.

6. In the Select Event Properties panel, use the Add filter menu to choose one or more event properties, including at least one Subtype filter. This is also the preferred location for specifying file names and paths. For your delete rule, you might choose a filter that specifies **Subtype is Malicious file detected**.
7. In the Select File Properties panel, use the Add filter menu to choose one or more file properties with which to further refine the conditions under which this rule will be triggered. For example, you might add filters that require more evidence that a file is malicious, such as:
 - **Analysis Result: Check Point is Malicious**
8. In the Select Process Properties panel, use the Add filter menu to choose one or more process properties with which to further refine the conditions under which this rule will be triggered. You might not need to specify anything in this field.
9. In the Select Action panel, choose **Delete file** on the Action menu. If you don't see this choice, check that you have permission to delete files – see [“User Role Permissions”](#) on page 106.

Create Event Rule

General

Copy Settings From: (none) ▼

Rule Name: Delete Malicious Files Reported by Multiple Sources

Description: If a file is considered malicious by three sources, delete it from all computers.

Status: Enabled Simulate only Disabled

Select Event Properties

Add filter [] ▼

Subtype is Malicious file detected [] +

Select File Properties

Add filter [] ▼

Analysis Result: Check Point is Malicious [] +

Analysis Result: Microsoft SCEP is Malicious [] +

Select Process Properties

Add filter [] ▼

Select Action

Action: Change global file state ▼

Change Global State: Change global file state

Resolve Related Approval Request: Change global process state

Create For: Change local file state

Upload file

Delete file

Move computer

Create & Exit Save

10. Once you choose *Delete file*, an Action Scope field with two radio buttons appears – choose one:

- **Delete files from the computer from the event** – This option deletes all instances of the matching file on the computer that reported the event.
- **Delete files from all computers** – This option deletes all instances of the matching file on all computers.

Select Process Properties

Add filter [] ▼

Select Action

Action: Delete file ▼

Action Scope: Delete files from the computer from the event
 Delete files from all computers

Create & Exit Save Cancel

11. When you have completed the rule definition, click **Save** to remain on the page or **Create & Exit** to create the rule and leave the Create Event Rule page.
12. If you configured the rule in Simulate Only mode as recommended, monitor the Processed Events section of the Edit Event Rule page for this rule to see what would have been deleted. Adjust the rule if necessary, and when you are satisfied that it will not delete files that shouldn't be deleted, change its Status to **Enabled**.

Monitoring File Deletions

When you submit a request for file deletion, you receive a confirmation in the console that the *request* was successful. This does not mean that the actual deletion of the file has occurred or that it will be successful when attempted. You can monitor the progress of a deletion request by viewing the Events page, and you can use Alerts to notify you when files are deleted, or when deletion fails.

When one deletion request affects multiple instances of a file, one event appears to report the outcome of the deletion attempt (success or failure) for each instance. This can happen when you request deletion from all computers or you request deletion from one computer and there are multiple instances of the file on that computer.

Table 42: File Deletion Events

Event Type	Event Subtype	Description / Example
Computer Management	File deletion requested	<p><i>If the request was to delete a file from one computer:</i></p> <ul style="list-style-type: none"> User 'admin' requested file deletion of all instances of [2488C...558F1] from MYCORP\DESKTOP6. <p><i>If the request was to delete a file from all computers:</i></p> <ul style="list-style-type: none"> User 'admin' requested file deletion of all instances of [FBAD9...34F00] from 100 computer(s). <p><i>If the deletion request came from an Event Rule:</i></p> <p>User 'System' requested file deletion of all instances of [81027...576DA] from MYCORP\DESKTOP6.</p>
Computer Management	File deleted	File 'test123.bat' [FBAD9...34F00] was successfully deleted from MYCORP\LAPTOP3
Computer Management	File deletion processed (file not found)	<i>If a file is in a computer's inventory but not on disk:</i> File deletion processed with file not found for [EDBD7...12F06] from MYCORP\DESKTOP9

Event Type	Event Subtype	Description / Example
Computer Management	File deletion failed	<p><i>If the deletion failed because it was a file from a protected publisher:</i></p> <p>File deletion failure of 'emet_gui.exe' [2024F...41CCD] from MYCORP\LAPTOP3. Error: Microsoft File</p> <p><i>If the deletion failed because the agent version doesn't support server-based deletion:</i></p> <p>File deletion failure of 'emet_gui.exe' [2024F...41CCD] from MYCORP\LAPTOP3 because this Agent version doesn't support it.</p> <p><i>If the deletion failed because the file is no longer present on the computer and not in its inventory:</i></p> <p>File deletion failure of 'tryme.bat' [76C7F...BD915] from MYCORP\DESKTOP8. Error: Delete Error[C0000034]</p>

In addition to actively monitoring the Events page, you can create alerts that will notify you when file deletion events occur. For example, you might create one alert that is triggered when file deletion succeeds and another that is triggered if a requested deletion fails. See [“Creating Alerts”](#) on page 612 for more information.

You can also check the Files on Computers page to determine whether a file deletion was successful, although if there is a large backlog of file activity, this might take a few minutes.

Chapter 10

Reputation Approval Rules

This chapter describes reputation approval rules, which can be used to automatically approve files based on the file and publisher trust ratings provided by the Carbon Black File Reputation.

Notes

Reputation approval rules require activation of Carbon Black File Reputation. See [“Activating Carbon Black File Reputation”](#) on page 763.

Other methods for approving files are described in [Chapter 8, “Approving and Banning Software.”](#)

Sections

Topic	Page
Overview	327
Reputation Approval Strategy	328
Creating Exceptions for Files and Publishers	331
Enabling Reputation Approvals	333
Modifying and Disabling Reputation Approvals	335
Views Related to Reputation Approvals	335

Overview

Carbon Black File Reputation provides a cloud-based database of known files. It pulls file data from a combination of distribution partners, Web crawlers, honeypots, and the Carbon Black user community. For files in the database, Carbon Black File Reputation reputation data provides context information such as who published the file and what product (if any) it is associated with. It also screens software using multiple anti-malware tools, and cross-references it against third-party vulnerability databases.

Using the information it has about a file, the Carbon Black File Reputation assigns a *threat* level and a *trust* rating. It also assigns a trust rating to publishers.

Reputation approval rules allow you to use these trust ratings to approve files automatically, with the following options:

- Approvals can be based on file or publisher reputation, and these options can be enabled together for maximum coverage and benefit.
- You set the trust thresholds at which you want files and publishers to be approved.
- Reputation approvals can be enabled for all agent-managed computers or by policy.
- You can disable reputation approvals for specific publishers and specific files that you don't want to be automatically approved.

If you are concerned about advanced threats, reputation approvals can be a good choice for approving files considered trustworthy. Automatic approval using reputation can give your end users more flexibility and reduce the effort of maintaining the list of approved files. Reputation approvals are based only on a file's *trust* rating (i.e., how *safe* it is believed to be), not on whether it is appropriate for a business environment.

When you enable reputation approvals, any manual file or publisher state assignments you have made remain in effect and take precedence over reputation. For example, if you ban a file by name or hash, that file remains banned even if it would have been approved by reputation. When and how reputation approval rules affect files on computers is described later in this chapter.

Trust Ratings for Files and Publishers

File Trust Ratings

The Carbon Black File Reputation bases a file's trust rating on a proprietary algorithm that takes the following factors into account:

- **Source Trust** – The origin of the file
- **Publisher Trust** – Whether the file has a signed digital certificate and the trust associated with that specific certificate
- **Malware Severity** – Whether anti-virus scanners identify the file as malicious or potentially malicious (e.g., a virus or malware); files in the Carbon Black File Reputation database are scanned by multiple anti-virus products
- **Vulnerability Severity** – Whether there is a known vulnerability for the file (specifically, a Microsoft-reported vulnerability), and if so, how severe
- **Duration Seen** – How long Carbon Black File Reputation has seen this file in the field
- **First Seen** – When this file was first seen in the field by Carbon Black File Reputation
- **Prevalence** – How common this file is in the field, as reported to Carbon Black File Reputation

The combination of these factors is used to calculate the trust rating of a file. Carbon Black File Reputation rates file trust on a scale from **0 (lowest trust)** to **10 (highest trust)**. For example, a signed operating system file with no known vulnerabilities would have a Trust value near 10. An unsigned third-party application not distributed via well-known websites might have a trust value of 3. Known malicious software, or an application distributing known malicious software, would have a Trust value at or near 0.

Publisher Trust Ratings

A publisher's trust rating is based on factors including aggregate experience with files from that publisher and the publisher's general reputation. There are four possible values for publisher trust: High, Medium, Low, and Not Trusted. If a publisher is Not Trusted, either there is no information about it or it is known *not* to have any of the factors that would elevate its trust level.

Reputation Approval Strategy

Reputation approvals allow high-trust software to run on agent-managed computers with little administrative effort. How you choose to implement reputation approvals will depend on your goals, especially the balance between convenience and protection. Although you can enable them separately, you get the maximum benefit of reputation approvals by enabling *both* file and publisher reputation approvals:

- **File reputation approvals** – Not all files are signed by a publisher. By using *file* reputation approvals, you can take advantage of the reputation data for specific files known to Carbon Black File Reputation, regardless of whether a file has a known publisher.
- **Publisher reputation approvals** – By using *publisher* reputation approvals, you ensure that all files signed by trusted publishers, including new files that might not have their own reputation yet, are approved and can run on agent-managed computers. Files from approved publishers are approved locally on connected agent-managed computers.

You can enable reputation approvals for all computers or only for computers in specific policies. There is no performance benefit or penalty for limiting reputation approvals to certain policies, so you should enable reputation approvals for all policies except those in which you want complete control over which specific files can be executed.

Note

When Carbon Black File Reputation is activated, Publisher Trust values are shown on the Publishers tab. This tells you what to expect when you enable Approvals for publishers. If the Trust value for a Publisher is High, then all files from that publisher will be approved when reputation approvals for publishers are enabled.

Setting the Trust Level for Approvals

You can set trust levels for file and publisher approvals any way you choose, but there are two recommended combinations:

Goal	File Trust	Publisher Trust
High Critical Asset Protection – For high protection for intellectual property and other confidential information	8	High
Protection with Flexibility – To protect your computers from risky files but allow automatic approval of more files with relatively low threat	6	Medium

When you enable both file and publisher reputation approvals, a file is approved if *either* its own reputation or its publisher's reputation meets the thresholds you set.

You can adjust these settings to meet your own judgment on the trade-offs, but setting the approval level at a very low trust level is not advisable. One way to see what the effect of approvals at different trust levels will be is to examine the File Catalog and the Publishers list in the console, grouped to show their contents by Trust.

To see files by trust category, choose **Assets > Files** on the console menu, click the **File Catalog** tab, and choose **Trust** on the *Group By* menu.

To see current publishers by trust category, choose **Rules > Software Rules** on the console menu, click the **Publishers** tab, and choose **Trust** on the *Group by* menu. This list includes only those publishers whose files have been inventoried on agent-managed computers or added by importing a certificate from a file on a computer without an agent.

How File Reputation Approvals Work

File reputation approvals rely on the most specific information available for the files known to the Carbon Black File Reputation. A separate reputation approval rule (global or by policy) is created on the App Control Server for each file meeting the reputation threshold. The scope of a reputation approval is determined by the list of policies on which reputation is enabled. As with other file approvals, reputation approvals can behave like per-policy approvals or global approvals, depending on your reputation settings.

File reputation rules are not listed on the App Control Server, but you can view a list of files approved by reputation. See [“Views Related to Reputation Approvals”](#) on page 335.

Unlike other approvals, file reputation approvals are not pushed to endpoints automatically. There are three conditions that cause a reputation-based file approval to be sent to endpoints *on which reputation approval is enabled*:

- If the App Control Server has a record of a file being blocked *on any endpoint* and that file is later approved by reputation, the server begins sending the approvals of the file to agents immediately.
- If a user attempts to execute an instance of a reputation-approved file on a computer connected to the App Control Server, and if the server detects that the file satisfies the reputation trust threshold, the server will allow the agent to run the file immediately, and also will begin sending the approval to other agents.
- If the reputation-approved file is identified as an installer, the App Control Server begins sending the approval of the file to agents immediately.

Even if a file is approved by reputation and not blocked by another rule, until its approval is sent to agents because of one of the cases above, instances of the file may be locally unapproved and may block if the agent computer is disconnected from the server before the approval is distributed.

Removal of Reputation Approval for a File

If the file reputation approval rule changes in a way that removes reputation approval from a file – by disabling reputation approval completely or by policy, by raising the approval threshold, or by lowering the file's own reputation – the global approval for that file is eliminated from connected computers, and the file state in the File Catalog reverts to unapproved. If an instance of this file was executed during the time it was approved by reputation, that instance remains locally approved on the computer where it was executed. Local approval would need to be removed manually unless you chose to ban the file.

Any explicit assignment of a ban or approval state to a file or its publisher takes precedence over a reputation approval.

Note

Approval by file reputation involves a significant initial impact on the App Control Server as files are analyzed to see whether they would be approved according to Trust Level. In addition, disabling file approvals or changing the approval threshold has a similarly significant impact. Avoid unnecessary changes in file reputation rule configuration.

How Publisher Reputation Approvals Work

When approval by publisher reputation is enabled, the list of all trusted publishers that meet the specified threshold is sent down to all computers. This allows App Control Agents that receive this publisher approval to approve new files from approved publishers as soon as they are seen for the first time, even if the computer is disconnected from the server when the new file arrives. In addition, the agent will approve all existing files from these publishers that were previously unapproved, unless they are explicitly banned.

Approval by publisher reputation has a low impact on the App Control Server and network traffic.

As with manual publisher approval, only files whose certificates meet all requirements described in [Approving or Banning by Publisher](#) can be approved by publisher reputation.

Removal of Reputation Approval for a Publisher

Once a file is locally approved because of its publisher's reputation, removal of *publisher* approval at a later time does not remove local approval of that file. Anything that removes approval from the publisher, including a change in reputation, a change in reputation approval settings, or completely disabling reputation approvals, affects only files that are encountered in the future.

If a publisher is no longer approved by reputation, its files can be returned to the unapproved state by manually removing the local approval on the Files on Computers or File Instance Details page for each instance. If a publisher is banned, however, that ban removes the approval from the file that was previously locally approved because of publisher reputation – it is not necessary to manually remove local approval for each instance.

Explicitly banning a file removes a local approval that occurred because of publisher approval.

Reputation Approvals and Other App Control Rules

Reputation rules can be affected by other actions you perform on the console:

- Any explicit file rule that identifies a file by name or hash automatically disables reputation control for that file. This includes global and policy-specific File Rules (bans and approvals), files on imported lists of hash approvals or bans, trusted directories, and manual publisher bans and approvals. Once a file is not controlled by reputation, reputation settings and thresholds have no effect on its approval state.
- To allow reputation to control the state of a file with reputation disabled, you must remove the explicit rule (approval or ban) and then re-enable reputation for the file.
- Custom rules and any other rules that directly block or allow access to a file supersede reputation approvals by file or publisher, even if they do not change file state.

Creating Exceptions for Files and Publishers

In general, you should enable reputation approvals because you want to rely on Carbon Black File Reputation data to eliminate a large number of unnecessary file blocks on trusted files. However, there might be a particular file or publisher that you do not want approved, regardless of its reputation in the Carbon Black File Reputation. You have the option of disabling reputation approvals for an individual file or publisher.

Notes

If you create file or publisher exceptions *before* the reputation feature is enabled for the App Control Server, those files or publishers are unaffected by reputation rules. Exceptions added *after* reputation rules are enabled prevent reputation approval of newly discovered files and remove global approvals based on file reputation, but they do not undo local approval of files whose publisher was approved by reputation.

Disabling Reputation Approvals for a File

You can disable reputation approvals for individual files. If you create the exception before enabling reputation rules on your server, it prevents any approvals by reputation for instances of the file. If you create the exception *after* enabling reputation rules, the reputation-approved file will revert to unapproved (both globally and locally) if no other approvals apply. If the file was already locally approved by some other means, however, (such as publisher approval or a custom rule), it will remain locally approved.

When you disable reputation for a file, it affects only that file, even if it is an installer.

To disable reputation approval for a file:

1. Open the File Details or File Instance Details page for the file.
2. In the Advanced menu to the right of the main page, click on **Disable Reputation Approval for this File**. Reputation approvals are disabled for the file and the menu choice changes to *Enable Reputation Approval for This File*.

File Properties		Advanced
Publisher:	Microsoft Corporation	View Carbon Black Reputation Data
Publisher State:	Approved (Manual)	Upload File to Server
Certificate:	Microsoft Corporation Microsoft Corporation Redmond Washington US	Mark as Installer
Certificate Type:	Embedded Signer	Disable Reputation Approval for This File
Certificate Global State:	Approved	
Company:	Microsoft Corporation	
Product Name:	Microsoft Malware Protection	
Product Version:	1.319.537.0	
File Size:	4,820,912 bytes	
Description:	AntiVirus Definition Update	
File Type:	Application	
SHA-256:	273D54CBC55C94C248F94DD6D9001A8A342ED9E9DB1936BCCA0DD9E5A8C1B44A	
MD5:	BC68A47B4274384CE4CDCAACE9ED32CD	
SHA-1:	6D2A51BB2A28968C72F2233A94FE886DAA29ED7E	

To re-enable reputation approval for a file:

1. Open the File Details or File Instance Details page for the file.
2. In the Advanced menu to the right of the main page, click on **Enable Reputation Approval for this File**. Note that this choice only appears if reputation has been disabled.

Disabling Reputation Approvals for a Publisher

You can disable reputation approvals for individual publishers. If you create the exception before enabling reputation rules on your server, it prevents any approvals by reputation for instances of files from the publisher. If you create the exception after enabling reputation rules, however, any files from an approved publisher found on agent-managed computers prior to disabling the publisher will already be locally approved by reputation, and will not become unapproved if you disable the publisher. Only files first seen by this App Control Server *after* you disable approval for the publisher are unaffected by the publisher's reputation.

To disable reputation approval for a publisher:

1. Open the Publisher Details page for the publisher.
2. Un-check the checkbox next to *Enable reputation approvals for this publisher*.
3. Click the **Save** button. Reputation approvals are disabled for this publisher.

Publisher Details

General

Publisher Name: Microsoft Corporation

State: Enable reputation approvals for this publisher

Acknowledged:

Trust: High

Description:

Platforms: All platforms Selected platforms

To re-enable reputation approval for a publisher:

1. Check *Enable reputation approvals for this publisher* on the Publisher Details page.
2. Click the **Save** button. Reputation approvals are disabled for this publisher.

Enabling Reputation Approvals

This section describes enabling the reputation approvals feature for your App Control Server. Before enabling reputation approvals:

- Consider exceptions you want to create for files and publishers that you do not want approved by reputation. These exceptions should be created *before* you enable the feature. See [“Creating Exceptions for Files and Publishers”](#) on page 331 for details.
- Consider whether you would like reputation approvals to be available for all of your agent-managed computers or only those in certain policies. This choice is covered in the procedure below.

Keep in mind that although you can add file and publisher exceptions after you enable reputation approvals for the App Control Server, the publisher exceptions do not reverse any local approvals that have already occurred due to publisher reputation.

To enable reputation approvals:

1. In the console menu, choose **Rules > Software Rules**. The Software Rules page appears.
2. On the Software Rules page, click the **Reputation** tab. The Reputation Approvals page appears.

Note

Carbon Black File Reputation must be activated before you can enable Reputation Approvals. If no Reputation tab appears on the Software Rules page, Carbon Black File Reputation is not activated. In this case, follow the instructions in [“Activating Carbon Black File Reputation”](#) on page 763 before continuing with this procedure.

- Click to check the box labeled *Enable reputation approvals*. This opens the fields on the page for editing.

- To enable *file* approval by reputation, make sure the box next to *Approve applications with a trust greater or equal to* is checked and then choose a trust level from the menu. File trust choices range from 1 (very low trust) to 10 (highest trust). See [“Setting the Trust Level for Approvals”](#) on page 329 for recommendations.
- To enable *publisher* approval by reputation, make sure the box next to *Approve publishers with trust greater or equal to* is checked and then choose a publisher trust level. Publisher trust has three values: *Low*, *Medium* and *High*.
- Select the policies for which you want to enable reputation approvals:
 - To enable the rules for all policies, click the *All policies* radio button.
 - To enable the rules only for some policies, click the *Selected policies* radio button and check the box next to each policy you want to be affected by these rules.

Note

You also can enable or disable reputation approvals for a policy on its Edit Policy page.

- When you have finished configuring reputation approvals, click the **Save** button at the bottom of the page and choose **OK** in the confirmation dialog. Reputation approvals are activated.

Note

Enabling file reputation approvals can require that very large numbers of file states are re-evaluated. You will not necessarily see changes in file state immediately in the console, but the server continues to process these changes in the background until all are up-to-date with the new approval rules. Full processing of the approvals may take several minutes.

Modifying and Disabling Reputation Approvals

You can modify or disable the reputation approval features in the same place where they were enabled. Modifications include changing the file or publisher trust threshold and changing the policies affected by reputation approvals. If you choose, you also can disable one type of reputation approval (i.e., publisher or file) while leaving the other in place.

The effect of modifying or disabling reputation approvals depends upon what kind of approval you enabled. Changes in reputation approval also have different network impacts as rules are re-evaluated.

- Changing the approval threshold for file reputation approvals can have a very significant one-time impact on server and network traffic while the changes are processed. Evaluation and updating of the File Catalog will take a few minutes, but depending upon the number of agents and the size of the File Catalog, it could take from hours to days to send the new file state information to all agents.
- Disabling file approvals can have a very significant network impact and, as with changing the approval threshold, might require from hours to days before all agents are updated with the changes in file state.
- Changes in publisher approval rules or policy coverage do not have a significant impact.
- Disabling publisher approval does not undo any local file approvals that already occurred because of publisher reputation.

To modify or disable the reputation approvals feature:

1. On the **Rules > Software Rules** on the console menu and click the **Reputation** tab. The Reputation Approvals page appears.
2. Make any needed changes and click **Save**.


Notes

- You also can enable or disable reputation approvals for a policy on its Edit Policy page.
- You can create exceptions for files or publishers you don't want controlled by reputation approvals. See [“Creating Exceptions for Files and Publishers”](#) on page 331.

Views Related to Reputation Approvals

The Related Views menu on the right side of the Reputation Approvals page provides links to additional information related to these approvals that is available in the console:

- **All files approved by reputation** – Clicking on this link shows the File Catalog page filtered to show all unique files globally approved by reputation.
- **All publishers approved by reputation** – Clicking on this link shows the Publishers tab of the Software Rules page filtered to show all publishers approved by reputation.
- **Reputation approval events** – Clicking on this link shows the Events page filtered to show all events related to reputation approvals (publisher and file).



These views can help you understand how reputation approvals are affecting your computers and perhaps point to changes you would like to make in the reputation approvals configuration, or in the state of specific files or publishers.

In other views that show files or publishers, you can see whether a file or publisher has been affected by reputation approvals by looking at these fields:

- **File State Reason** – If the file was approved by file reputation, this field shows *Reputation*. If the file has an approved publisher the File State can be *Approved by Reputation* even when File State Reason is something other than Reputation.
- **Publisher State Reason** – If the publisher for a file is approved by reputation, this field shows *Reputation*.
- **Reputation Enabled (files)** – The File Details and File Instance Details pages include a Reputation Enabled field that shows whether file reputation approvals are enabled for the current file. You can add this same field to the File Catalog and Files on Computers pages. Note that a value of Yes means that the file can be approved by reputation, not that it is approved.
- **Reputation Enabled (publishers)** – On the Publishers tab on the Software Rules page, you can add a column that shows whether reputation approvals are enabled for each listed publisher. As with files, a value of Yes means that the publisher can be approved by reputation, not that it is approved.

Chapter 11

Managing File-Signing Certificates

This chapter describes advanced features for using file-signing certificates in App Control file monitoring and enforcement activities. These features provide the following capabilities:

- **Certificate Discovery and Inventory** – Information about file-signing certificates discovered on agents and all certificates in their chains is collected and stored in the App Control Server database.
- **Enforcement by Certificate State** – Any certificate in a certificate chain may be approved or banned for a specific publisher, and its state can be used to approve or ban files managed by App Control.

Platform Note

These certificate visibility and control features are available for computers running Windows only. The features are not available for systems running Linux or macOS.

Sections

Topic	Page
Overview	338
Summary of Certificate Management Features	339
Viewing Certificate Information	339
Viewing Certificates for a Publisher	346
Certificate Alerts	347
Certificate Events	348
Using Certificates for Enforcement	348

Overview

App Control provides the ability to approve or ban a publisher by its *name*, as identified in a certificate. Files signed with certificates whose publisher name matches an approved publisher are approved unless banned by some other rule; files with certificates whose publisher name matches a banned publisher are banned. All files with a given publisher name in their certificate are affected by that publisher's state as defined on your App Control Server. These rules are described in “[Approving or Banning by Publisher](#)” on page 281.

The certificate management features described in this chapter add another layer of security and information to publisher approvals. While publisher names in certificates are not controlled by any central authority, certificates themselves are. A certificate identifies an individual, a server, a company, or other entity, and associates that identity with a public key. It provides generally recognized proof of identity based on public-key cryptography. Only the public key certified by the certificate will work with the corresponding private key possessed by the entity identified by the certificate; in this case, the entity is a file.

File-signing certificates are the final link in chains or paths of certificates. There is a root certificate, which identifies the entity that conferred the initial trust. That certificate might be used to sign an intermediary certificate, which then confers its trust to the final leaf certificate that specifically identifies the file. There can be more than one intermediary certificate in a path.

The App Control Agent reports all identifiable, valid certificates in the path of trust for signed files it discovers. Any certificate in the path of the signing certificate can be approved or banned. When a certificate is assigned a state of Approved or Banned (or left as Unapproved), that state applies only for a specific publisher of a leaf certificate. If the same certificate happens to appear in the certificate chain for a file signed by a different publisher, a separate certificate approval or ban is needed to affect that file.

Note

In late 2013, Microsoft published security bulletin MS13-098 describing a flaw in the Authenticode signature verification that could allow remote code execution. In response, Microsoft announced availability of an update for all supported releases of Microsoft Windows to change how signatures are verified for binaries signed with the Windows Authenticode signature format.

If this change is enabled, Windows Authenticode signature verification will no longer allow extraneous information in the WIN_CERTIFICATE structure, and Windows will no longer recognize non-compliant binaries as signed. Activation of this new behavior could cause files previously approved by publisher to block on App Control-managed systems.

The change is included with Security Bulletin MS13-098, but (as of July 2014) will only be enabled on an opt-in basis. However, Microsoft states that it may make this a default behavior in a future release of Microsoft Windows.

See <https://technet.microsoft.com/library/security/2915720> for more information on this change.

Summary of Certificate Management Features

App Control Certificate Management includes the following specific features:

- In the console menu, you choose **Assets > Certificates** to open the **Certificates table** page. The Certificates table shows all leaf certificates that have been used to validly sign or cosign files found on agent-managed computers, plus all certificates in the paths for those leaf certificates.
- Clicking on the View Details button next to a certificate in the table opens the **Certificate Details** page for that certificate. The Certificate Details page shows complete details for one certificate and has links to Related Views relevant to the certificate, such as a table of all files signed by the certificate.
- The **Publisher Details** page for each publisher includes an **All Certificates for This Publisher** panel. This panel shows all certificates that have this publisher name as the CN portion of the certificate Subject Name. It also shows the approval/ban state for each certificate in the certificate path for leaf certificates associated with that publisher, and allows you to add or remove approvals or bans for each certificate.
- Certificate-related fields are included on **File Details** and **File Instance Details** pages.
- On the **Advanced Options** tab of the **System Configuration** page, the **Certificate Options** panel includes settings that determine what requirements (such as key length and algorithm) a certificate must meet if it is to be used for approving files. You can configure rules that enable the agents to do their own certificate revocation checks.
- Regardless of whether agent-based certificate revocation checks are enabled, the App Control Server validates certificates in its inventory on a recurring basis to make sure that they have not been revoked. This validation generally occurs on a weekly basis and involves downloading certificate revocation lists (CRLs) from registration authorities or making Online Certificate Status Protocol (OCSP) calls to OCSP responders. If you are monitoring network traffic, keep in mind that these downloads might involve a variety of sites in a variety of countries. Currently, only agent-based revocation information affects enforcement of rules. Server-based validation checks are provided to inform administrators when the status a certificate changes, but they do not affect enforcement of rules. Enable agent-based revocation checks if you want revocations to affect rule behavior.
- Certificate-related **Events** and **Alerts** may appear when triggering conditions occur.

Viewing Certificate Information

Certificate information is available in several locations in the App Control Console. This information can help you make decisions about whether to approve or ban certain certificates.

Certificates Table

The Certificates table shows all leaf certificates that have been used to validly sign or cosign files found on agent-managed computers, plus all certificates in the paths for those leaf certificates. The table also provides access to the Certificate Details page for each

certificate – you can click either the View Details button or the Subject Name in the table to see details for a certificate.

Note

The Certificate table is a read-only page with no Action menu. Certificate state can be changed only in the context of a specific publisher, on the Publisher Details page. See [“Approving or Banning Certificates for a Publisher”](#) on page 351 for more information.

To view the Certificates table:

- On the console menu, choose **Assets > Certificates**.

Certificates

Group By: (none) Ascending

Show Filters | Show Columns | Export to CSV | Refresh Page

Subject Name	Publisher	Unique Signed Files	Path Position	Global State	Certificate State
<input checked="" type="checkbox"/> .NET Foundation Projects Code Signing CA .NET Foundation U...	.NET Foundation Projects Code Signing CA	0	Intermediary	Unapproved	Unapproved
<input checked="" type="checkbox"/> AddTrust External CA Root AddTrust External TTP Network Ad...	AddTrust External CA Root	0	Root	Unapproved	Unapproved
<input checked="" type="checkbox"/> Adobe Inc. Flash Player Adobe Inc. San Jose ca US 2748...	Adobe Inc.	2	Leaf	Unapproved	Unapproved
<input checked="" type="checkbox"/> "Bit9, Inc." "Bit9, Inc." Waltham Massachusetts US	Bit9, Inc	0	Leaf	Approved	Approved
<input checked="" type="checkbox"/> "Bit9, Inc." "Bit9, Inc." Waltham Massachusetts US	Bit9, Inc	0	Leaf	Approved	Approved
<input checked="" type="checkbox"/> "Bit9, Inc." Engineering "Bit9, Inc." Waltham Massachuse...	Bit9, Inc.	0	Leaf	Approved	Approved
<input checked="" type="checkbox"/> "Carbon Black, Inc." "Carbon Black, Inc." Waltham Massach...	Carbon Black, Inc.	0	Leaf	Approved	Approved
<input checked="" type="checkbox"/> "Carbon Black, Inc." "Carbon Black, Inc." Waltham Massach...	Carbon Black, Inc.	0	Leaf	Approved	Approved
<input checked="" type="checkbox"/> "Carbon Black, Inc." "Carbon Black, Inc." Waltham Massach...	Carbon Black, Inc.	13	Leaf	Approved	Approved
<input checked="" type="checkbox"/> Cortado AG Software Development Cortado AG Berlin Berlin...	Cortado AG	20	Leaf	Unapproved	Unapproved
<input checked="" type="checkbox"/> DigiCert Assured ID CA-1 www.digicert.com DigiCert Inc US	DigiCert Assured ID CA-1	0	Intermediary	Unapproved	Unapproved
<input checked="" type="checkbox"/> DigiCert Assured ID Code Signing CA-1 www.digicert.com Dig...	DigiCert Assured ID Code Signing CA-1	0	Intermediary	Unapproved	Unapproved
<input checked="" type="checkbox"/> DigiCert Assured ID Root CA www.digicert.com DigiCert Inc ...	DigiCert Assured ID Root CA	0	Root	Unapproved	Unapproved

The default table includes selected columns with key information about each certificate. As with any App Control table, you also may add or remove columns from the table view using the *Column Settings* panel (See [“Console Tables”](#) on page 66 for more information about customizing a table view.). [Table 43](#) shows the possible fields available on the Certificates table and also the Certificate Details page. Keep in mind that some of these fields are not shown by default in the table.

Table 43: Fields in Certificates Table and Details Pages

Field/Column	Source	Appears	Description
Note: In the Where column, T = Table page, D = Details page			
Subject Name	Cert	T, D	Distinguished name of the subject of the certificate, in this case the signer of the file. In the table, the name is shortened, but a tooltip provides a full length Subject Name. Clicking on the name in the table opens the details page for this certificate.
Publisher	Cert	T, D	Publisher name as identified by the CN portion of the Subject Name in the certificate. If this publisher signed any files in the File Catalog, clicking the name opens the Publisher Details page. Some of the “Publishers” listed are certificate authorities, not actual software publishers, and so do not have linked names.
Unique Signed Files	App Control	T, D	Number of unique files in the File Catalog signed by this certificate. If greater than zero, clicking on the number opens the File Catalog filtered to show these files.
Path Position	Cert	T	Position of this certificate in the certificate path cataloged on the server. The possible values are: Root, Intermediary, Leaf. See “Path Position and Agent Differences” on page 350 for details about certificate path position, variations among agents, and the impact on certificate management.
Root Certificate	Cert	D	Is this a root certificate? The possible values are: Yes, No.
Global State	App Control	T,D	Effective state of this certificate derived from the following: Publisher State of the publisher identified in this certificate; Certificate State; Certificate Path State, and certificate configuration settings. See “Certificate Global State” on page 352 for global certificate state determination, values, and how it interacts with the states of other objects.
Certificate State	App Control	T	State assigned to the certificate for this publisher. The possible values are: Approved, Unapproved, Banned. See “Certificate Global State” on page 352 for a description of how this affects global certificate state and file state.

Field/Column	Source	Appears	Description
Certificate State Details (in details) Global State Details (in table)	App Control & Cert	T,D	Detailed description of all of the factors contributing to Certificate Global State. See “Certificate Global State” on page 352 for more information.
Valid From	Cert	T,D	Date this Certificate is valid from. Format is MMM DD YYYY HH:MM:SS AM/PM (UTC).
Valid To	Cert	T,D	Date this Certificate is valid to. Format is MMM DD YYYY HH:MM:SS AM/PM (UTC).
Signature Algorithm	Cert	T,D	Algorithm used to create the certificate’s signature. Typical values: MD2RSA, MD5RSA, SHA1RSA, SHA256RSA. See “Certificate Approval Configuration Choices” on page 349 for configuration settings related to this field.
Thumbprint	Cert	T,D	SHA1 hash value of this certificate.
Certificate ID	App Control	T,D	Unique App Control-generated hash identifier for this certificate.
First Seen Date	App Control	T,D	Date and time this certificate was first seen and inventoried on this App Control Server.
Last Modified Date (in details) Date Modified (in table)	App Control	T,D	Date and time the record for this certificate was last modified on this App Control Server.
Description	App Control	T,D	An editable field in which console users can add or modify a comment about this certificate.
Last Validation Date	App Control	T,D	Last date and time when this certificate was validated on the App Control Server. Certificates are validated when discovered and periodically re-checked.
Public Key Algorithm	Cert	T,D	Algorithm used to produce the public key.
Public Key Size	Cert	T,D	Size of the public key for this certificate. See “Certificate Approval Configuration Choices” on page 349 for size settings.
Serial Number	Cert	T,D	A field in the certificate containing a number that is unique among certificates from its issuing certificate authority.

Field/Column	Source	Appears	Description
Type	Cert	T,D	<p>Indicates whether a certificate was embedded or detached or both, and whether the signature was used to sign the file or to countersign the signature, usually for timestamp validation. Leaf certificates only.</p> <p>The possible values are: Embedded, Detached, Signer, Cosigner. Each certificate has two or more of these values.</p> <p>See “Certificate Types” on page 350 for details about type and its impact on certificate management.</p>
Validation Error (in Table) Validation Message (in Details)	Cert	T,D	<p>Shows any error messages returned when the certificate is checked. If the certificate check produces no errors, this field will be blank. See http://msdn.microsoft.com/en-us/library/windows/desktop/aa377590(v=vs.85).aspx for a list of possible messages.</p> <p>Many certificates show validation errors for reasons that are not necessarily an indication of significant risk. For example, a certificate authority may stop providing information (and thus validation) for older certificates.</p>
History	App Control	D	<p>Panel includes the following where appropriate:</p> <ul style="list-style-type: none"> • First Seen Date – The date and time this certificate was first seen in your App Control environment. • Last Modified by – The console user that made the most recent change to certificate state (not in table). • Last Modified Date – The date and time when the most recent change to certificate state was made.
Certificate Path	Cert	D	<p>Panel shows this certificate in the context of its path. Each item in the list (except for the current certificate) is a link to the certificate details for other certificates in the path.</p>

Searching, Sorting and Grouping on the Certificates Table

You can use any of the standard table customization methods to show or find specific certificates. For example, you can use the Show Filters menu to search for a particular certificate by Subject Name or Hash, or use the *Group by* menu to organize the Certificates by certain fields. The *Group by* menu includes the following choices:

- Subject Name

- Publisher
- Unique Signed Files
- Path Position
- Global State
- Certificate State
- Valid From
- Valid To
- Signature Algorithm
- Thumbprint

Certificate Details

The Certificate Details page shows complete details for one certificate. It also has links to Related Views relevant to the certificate. [Table 43](#) describes the fields that appear on this page.

To view the Certificates Details page for one certificate:

- In the Certificates table or the certificates section of a Publisher Details page, click on the View Details button or the Subject Name for the certificate.

In other locations in which certificates information is displayed, such as the Events table, you can click on the Subject Name of a certificate to see its details.

Certificate Details ?

General

Publisher: [Google LLC](#)
Subject Name: Google LLC Google LLC Mountain View ca US
Thumbprint: cb7e84887f3c6015fe7edfb4f8f36df7dc10590e
Last Validation Date: Jun 29 2020 07:53:07 PM
Unique Signed Files: [17](#)
Description:

Certificate State For Publishers

Publisher	Certificate Global State	Certificate State Details
Google LLC	Approved	Certificate is Unapproved, Publisher is Approved, Certificate Path is Unapproved

Certificate Properties

Serial Number: 0c15be4a15bb0903c901b1d6c265302f
Signature Algorithm: sha256RSA
Valid From: Nov 6 2018 07:00:00 PM
Valid To: Nov 17 2021 07:00:00 AM
Root Certificate: No
Type: Embedded Signer
Public Key Algorithm: RSA
Public Key Size: 2048

History

First Seen Date: Jun 29 2020 07:51:57 PM
Last Modified By: System
Last Modified Date: Jun 29 2020 07:51:57 PM

Certificate Path

Subject Name

- DigiCert Assured ID Root CA www.digicert.com DigiCert Inc US
- DigiCert SHA2 Assured ID Code Signing CA www.digicert.com DigiCert Inc US
- Google LLC Google LLC Mountain View ca US

Related Views

- All files signed by this certificate
- All unique files signed by this certificate
- Files signed by certificates with this certificate in path
- All events for this certificate

In the Certificate Details, if the Publisher name is highlighted as a link, you can click on it to go to the details page for the publisher of this certificate. You also can click on any highlighted certificate name in the Certificate Path panel to view its details. If this certificate has signed files, clicking on the number next to the Unique Signed Files field displays a File Catalog view filtered to show those files. Note that Publisher names for intermediate and root certificates are *not* links.

Related Views Menu on Certificate Details

The Certificate Details menu includes a Related Views menu that can provide additional information about a certificate and how it is being used in your environment. Not all Related Views choices are available for all certificates. The view options are:

- **All files signed by this certificate** – Displays the Find Files page filtered to show all file instances signed by this certificate (i.e., for which this is the “leaf” certificate).
- **All unique files signed by this certificate** – Displays the File Catalog page filtered to show all unique files signed by this certificate.
- **Files signed by certificates with this certificate in path** – Displays the Find Files page filtered to show all file instances that have this certificate in their certificate path.

- **All child certificates for this certificate** – Displays the Certificates page filtered to show child certificates at any level below this certificate. Does not appear for Leaf certificates.
- **All events for this certificate** – Displays the Events page filtered to show events related to this certificate. This includes creation or deletion of bans and approvals, discovery or addition of certificates, and certificate checks.

Viewing Certificates for a Publisher

The Publisher Details page includes a panel entitled *All Certificates for This Publisher*. Because this panel has the potential to be long, it can be collapsed and expanded on the page by clicking the panel name.

To view certificates in the Publisher Details page:

1. Click on the highlighted name of a publisher.

Note: The publisher name appears in many places, including events, file details, and certificate details. If it is not highlighted, it is not a software publisher that signs files directly but may be a certificate authority that signs certificates that sign files.
2. If the certificates for the publisher are not shown, click on **All Certificates for this Publisher**.

Publisher Details

General

Publisher Name: Google LLC

State: Enable reputation approvals for this publisher

Acknowledged:

Trust: Not Trusted (Unknown)

Description:

Rule Applies To: All Current and Future Policies
 Selected policies

Platforms: All platforms
 Selected platforms

All Certificates For This Publisher (click to hide)

Group By:

[Show Filters](#) | [Show Columns](#) | [Export to CSV](#) | [Refresh Page](#)

Action

<input type="checkbox"/>	Subject Name
<input checked="" type="checkbox"/>	Certificate State: Unapproved
<input type="checkbox"/>	DigiCert Assured ID Code Signing CA-1 www.digicert.com DigiCert Inc US
<input type="checkbox"/>	DigiCert Assured ID Root CA www.digicert.com DigiCert Inc US
<input type="checkbox"/>	DigiCert SHA2 Assured ID Code Signing CA www.digicert.com DigiCert Inc US
<input type="checkbox"/>	Google LLC Google LLC Mountain View ca US
<input type="checkbox"/>	Google LLC Google LLC Mountain View California US

History

The panel shows all leaf certificates for this publisher, and all root and intermediate certificates associated with these leaf certificates. It is similar to the Certificates table, and you can modify it using standard filter, column, and grouping tools for tables. For any

certificate shown, you can go to its details page by clicking on the View Details button or Subject Name.

The table of certificates on the Publisher Details page has an Action menu. Using this menu, you can ban, approve, or remove an approval or ban from a certificate in the context of the current publisher. This is described in more detail in [“Approving or Banning Certificates for a Publisher”](#) on page 351.

Certificate Fields in File/File Instance Details

There are certificate-related fields in the File Catalog, Files on Computers, File Details, and File Instance Details pages in the console. In most cases, the certificate name or hash is a link from the file information to full information about the certificate that has signed the file. Certificate information is also included in the Global State Details in on all of these pages. See [Chapter 7, “File, Publisher, and Application Information,”](#) for more information.

[Table 44](#) shows where certificate-related fields appear on file pages.

Table 44: Certificate-Related Fields in File Table and Details Pages

Field	File Catalog	Files on Computers	File Details	File Instance Details
Certificate			Link	X
Certificate Type			X	X
Certificate Global State	X	X	X	X
Certificate Hash	Link			
Certificate State Reason	X	X		
Certificate Subject Name	X	X		
Detached Certificate Subject Name (Detached Certificate in Details pages)		X		Link
Detached Certificate Type				X
Detached Certificate State				X

Certificate Alerts

There are two certificate-related alerts that may be of particular interest if you are using certificates as part of your security enforcement plan:

- **New Certificate Alert** – Alerts subscribers when a file with a certificate for a publisher not previously listed on this server is discovered, and when a new certificate is imported directly into the App Control Server. By default, this alert is triggered when a new certificate for any publisher is detected. However, it can be configured to trigger only for new certificates for specific publishers.
- **Revoked Certificate Alert** – Alerts subscribers when a certificate known to this App Control Server is revoked. By default, this alert is triggered when a certificate for any

publisher is revoked. However, it can be configured to trigger only for specific publishers.

There is a special mail template for informing users about certificate discovery or revocation.

See [“Using App Control Alerts”](#) on page 608 for more information about configuring, enabling, and responding to alerts.

Certificate Events

App Control reports events associated with file-signing certificates. These events appear on the Events page in the console and are also available in Syslog output. The event description for certificate-related events includes the Subject Name. On the console Events page, Subject Name is a link to the Certificate Details page.

See the separate *VMware Carbon Black App Control Events Guide* to for more on event subtypes (the unique App Control identifier for an event) for certificates. The subtypes fall into two types:

- **Discovery events** – These are events that have to do with the certificates themselves, independent of their App Control state.
- **Policy Enforcement events** – These are events that report addition or removal of a App Control ban or approval for a certificate.

See [“Event Reports”](#) on page 590 for more information about viewing events in the App Control Agent. See the separate document *VMware Carbon Black App Control Events Guide* for a list of events.

Certificates in External Views

App Control provides public views into the database of files and events as an alternative to the console. You can create your own reporting and data analysis solutions through the use of these public views. The certificate-related event subtypes described in the previous section may be included in the ExEvents view. In addition, certificate metadata is included with file information in the following views:

- **ExFileCatalog** – Metadata for all unique hashes
- **ExFileInstances** – Metadata of all file instances on all computers
- **ExDeletedFileInstances** – Metadata of all deleted file instances

See [Appendix A, “Live Inventory SDK: Database Views,”](#) for more information about accessing the external views of the App Control database.

Using Certificates for Enforcement

The previous sections of this chapter focused on information App Control provides about certificates. This section describes certificate approvals and bans, and their effect on file state. The following is a summary of the certificate approval and ban features:

- **Certificate Approval Settings** – The System Configuration page has Advanced Options that affect whether certain certificates can be globally approved.
- **Manageable Certificate Types** – Regardless of configuration choices, not all discovered certificates can be approved or banned.

- **Path Position and Agent Differences** – For the same certificate/publisher combination, different agents can have different certificate paths, and the path on the server may match some or none of those currently on the agents.
- **Certificate State** – Approving or banning a certificate (or removing approvals and bans) determines Certificate State for a specific certificate for a specific publisher.
- **Certificate Global State** – Other factors interact with Certificate State to determine the Certificate Global State, which is its effective state.
- **Impact on File State** – Certificate Global State interacts with other rules and states to determine the state of a file signed by a particular certificate or one of its children.
- **Certificate Ban Setting** – Each computer's Policy has an Advanced Setting that determines whether certificate bans are effective.

A key point to keep in mind when preparing to approve or ban certificates is that you must specify the state *in each publisher for which you want it to be effective*.

Note

To be effective for approving a file on Windows agents (the only agents these features currently apply to), all certificates in the certificate chain for that file must be considered valid by Windows. For example, current root certificates must be installed for a certificate to be accepted.

Certificate Approval Configuration Choices

There are configuration settings on the Advanced tab of the System Configuration page that determine whether a certificate approval is effective in determining the state of a file signed by that certificate. [“Determining Which Certificates Can Approve Files”](#) on page 286 describes these configuration options in detail.

Remember that certificates can be approved and banned themselves, and also can be used to approve or ban a publisher by name. Keep the following in mind when setting or viewing Certificate Options on the Advanced Options page:

- You can approve a certificate that does not meet these configuration requirements, and the certificate itself will show a Certificate State of Approved. However, the Certificate Global State (the effective state) of such a certificate cannot be Approved.
- Certificate Options choices have no effect on cosigner certificates.
- Certificate Options choices do not prevent any certificate from being banned, or prevent the value of Certificate Global State from being Banned. See [“Certificate Global State”](#) on page 352 for more information.
- The Expired Certificates option on the System Configuration/Advanced Options tab does not affect the ability to globally approve a *certificate*; it determines whether an expired certificate can be used to approve a file by *publisher*. If the box is checked, then if a file has a certificate that has expired but was used to sign the file during the valid period, the certificate may be used for approval by publisher. If not checked, expired certificates may not be used to approve files by publisher. This setting does not affect Certificate Global State.

Certificate Types

The Certificate Details page includes a Certificate Type field, which has a value for leaf certificates only. Certificate type indicates what the leaf certificate is being used for and how it is associated with a file. Type is some combination of the following terms:

- **Embedded** – The digital signature for a file is embedded in a non-executable part of the file itself.
- **Detached** – The file to be signed is hashed into a digest and the digital signature is applied to the digest and included in a separate catalog file, which can contain certificates for multiple files.
- **Signer** – The certificate is the code-signing certificate for files it signs.
- **Cosigner** – The certificate is a cosigner (also called “countersigner”) certificate for files it signs. Cosigner certificates are normally used for time stamping.

Each instance of a leaf certificate must be either embedded or detached, and it must be a signer or a cosigner, so the minimum number of descriptors in the Type field for any certificate is two. There could be more than two since the same certificate can be used in different ways and so can have different types. One certificate in the Certificates Table may display its Type as Embedded Detached Signer, for example, or some other combination of these terms.

Important

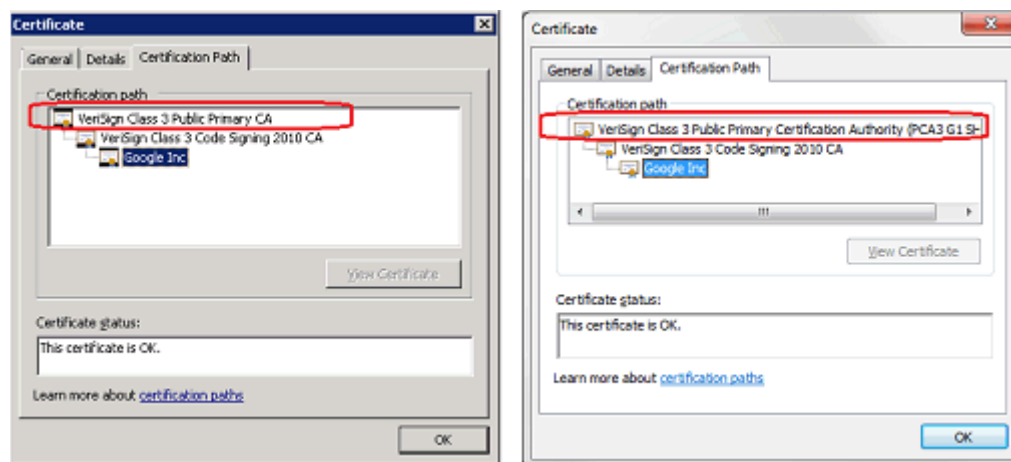
Only certificates identified and used as the Signer for a file may be approved or banned. Cosigner certificates are not assigned a state by App Control.

Path Position and Agent Differences

When you view the Publisher Details page, the *All Certificates for This Publisher* panel shows all certificates in the path. You can approve or ban any of these certificates. Before doing that, however, consider the different impact of approving or banning at different points on the path.

Certificate paths for the same leaf certificate may vary on different agents, or between an agent and the server. This could occur when the same file is received from different sources, or when one computer has updaters enabled and another does not. Agents update their certificate paths over time to minimize these differences.

Because of the potential for path differences, approving or banning intermediary or root certificates might not have the results you expect. The following example shows the same leaf certificate (same Issuer and Serial Number) with different root certificates:



If you approved one of these roots and expected that to take care of all instances of the leaf, you would not see the desired results on all agents. Path differences might be less of an issue for internally signed certificates for which you control the entire certificate path.

To reduce certificate path variation, keep your certificate stores on agents and the server current. Also, make sure that operating system updaters and other key application updaters are allowed to run so that you have the latest versions of signed files.

Approving or Banning Certificates for a Publisher

You approve or ban a certificate in the context of a publisher. The certificate state is effective only within that publisher; if a certificate is used by multiple publishers, you must assign its state in each.


Approving or banning a certificate defines its Certificate State, which can be Approved, Banned, or Unapproved. The effective certificate state, called Certificate Global State, is what is applied to files signed by that certificate. A certificate's Global State depends on the following: the Certificate State, the state of other certificates in its path, the Publisher State, and (for approvals) certificate configuration choices. See [“Certificate Global State”](#) on page 352 to see how different Certificate Global States are produced.

To approve or ban a certificate:

1. Make sure you have set the appropriate Certificate Options on the System Administration Advanced Options tab. These determine which certificates may be used for approvals. See [“Determining Which Certificates Can Approve Files”](#) on page 286.
2. Locate the certificate(s) you want to approve or ban, and then open the Publisher Details page for the publisher to which you want this approval or ban applied.

Note: You may locate a certificate first and click on its publisher name (e.g., from the File Details page, the Events page, or the Certificates table); or if you know the publisher for the certificate, open its details page directly.
3. If the certificates for the publisher are not already showing, click on **All Certificates for this Publisher**.

- In the *All Certificates for this Publisher* panel, check the boxes next to any certificates you want to approve or ban. Note that all of the checked certificates must have the same state applied to them – that is, you cannot simultaneously approve some certificates and ban others.



- On the Action menu, choose **Approve Certificates** or **Ban Certificates**. The Certificate State of the checked certificates is changed to the state you selected.

See “[Certificate Global State](#)” to see how a Certificate State of Approved, Unapproved, or Banned interacts with other states and rules to produce Certificate Global State.

To remove a certificate approval or ban:

- On the Publisher Details page, check the certificates whose state you want to change. Note that you can select a combination of banned and approved certificates for this operation.
- On the Action menu, choose **Remove Approval or Ban**. The Certificate State of all checked certificates becomes Unapproved.

Certificate Global State

The Certificate Global State is the effective state of a certificate. The possible values for Certificate Global State:

- Unapproved
- Approved
- Banned
- Approved By Policy
- Banned By Policy
- Mixed

Certificate Global State is determined by the following factors:

- Certificate State** – Values are Unapproved, Approved, or Banned.
- Publisher State** – Values are Unapproved, Approved, Banned, Approved By Policy, or Banned By Policy.
- Certificate Path State** – Values are Unapproved, Approved, Banned, or Mixed (some certificates in the chain are Approved and some are Banned).

- **Certificate Key Length and Algorithm** – Does this certificate meet System Configuration/Advanced Settings requirements.

For any certificate, you can view the factors that contribute to Certificate Global State on the Certificates (table) page or the Certificate Details page. The *Certificate State for Publishers* panel on the details page summarizes the relevant factors.

Certificate Details

General

Publisher: Google LLC
Subject Name: Google LLC Google LLC Mountain View California US
Thumbprint: a3958ae522f3c54b878b20d7b0f63711e08666b2
Last Validation Date: Jun 29 2020 07:53:06 PM
Unique Signed Files: 68
Description:

Certificate State For Publishers

Publisher	Certificate Global State	Certificate State Details
Google LLC	Approved	Certificate is Unapproved, Publisher is Approved, Certificate Path is Unapproved

Certificate Properties

Even if other elements of certificate state are approved, Certificate Global State might be Unapproved if the certificate does not meet the certificate specifications on the System Configuration/Advanced Options page. When a certificate fails to meet more than one configuration requirement (e.g., both the minimum key size and the allowed algorithm specifications), only one of the two reasons appears in Certificate State Details.

The following examples may help clarify the way these values interact with each other to produce Certificate Global State. All possible combinations are shown in [Table 45, “Determining Certificate Global State”](#), on page 356.

Example 1: All States and Configuration Allow Approval

Condition	Example/Comments
If the certificate meets the minimum key size configuration...	Minimum Certificate Key Size: 1024 Key length of this certificate: 2048
...and its algorithm type is not configured to be ignored...	Certificate Signature Algorithms to Ignore: Only MD2RSA is checked Signature Algorithm of this certificate: SHA1RSA
...and the certificate has a countersignature if required...	.
...and the configured revocation checks have not found the certificate revoked...	
...and the leaf Certificate State is Approved...	A console user chose to approve the certificate.
...and the Publisher State is Approved...	The publisher was approved by a console user or by reputation.
...and no other certificate in this certificate's path is Banned...	The state of the Certificate Path is shown in the Certificate Details
...then the Certificate Global State is Approved.	This is the state that will affect files signed by the certificate.

Example 2: Certificate Does Not Meet a Configuration Requirement

Condition	Example/Comments
If the certificate meets the minimum key size configuration...	Minimum Certificate Key Size: 1024 Key length of this certificate: 2048
but its algorithm type is configured to be ignored...	Certificate Signature Algorithms to Ignore: Has MD2RSA and SHA1RSA checked Signature Algorithm of this certificate: SHA1RSA
...and the Certificate State is Approved...	
...and the Publisher State is Approved...	
...and no other certificate in this certificate's path is Banned...	
...then the Certificate Global State is Unapproved.	Although all other approval criteria were met, the certificate algorithm is not allowed for approvals.

Example 3: Banned Certificate in the Path

Condition	Example/Comments
Whether or not the certificate meets the minimum key size...	
...and no matter whether it meets any of the other Advanced Options requirements...	
...and if the Publisher State is Approved or Unapproved and does not have any policy restrictions...	
...if this certificate or any certificate in the certificate path is Banned...	
...then the Certificate Global State is Banned.	Although Certificate Global State is Banned, the ban's effectiveness on each agent depends upon <i>Block files with banned publishers or certificates</i> on the Advanced Settings of the agent's policy. This setting is active by default.

Example 4: Mixed Global State

Condition	Example/Comments
If the Publisher State is Approved by Policy...	
...and if this certificate or any certificate in it the certificate path is Banned...	
...then the Certificate Global State is Mixed.	Certificate Global State acts as Unapproved for policies with publisher approval. Certificate Global State acts as Banned for policies not included in the publisher approval if banning by certificates is allowed in the policy.

[Table 45](#) shows how different combinations of Certificate, Publisher, and Certificate Path states produce different Certificate Global states. All of these outcomes assume that all certificates in the path meet the configuration requirements specified on the System Configuration Advanced Options page. Where “(by Policy)” appears in parentheses in the table, the Certificate State shown is not *specified* as being by policy but is *effectively* “by Policy” because Publisher State is Approve by Policy or Ban by Policy.

Table 45: Determining Certificate Global State

#	Certificate State	Publisher State	Certificate Path State	Certificate Global State
1	Unapproved	Unapproved	Unapproved	Unapproved
2	Approved	Unapproved	Unapproved	Approved
3	Banned	Unapproved	Unapproved	Banned
4	Unapproved	Approved	Unapproved	Approved
5	Approved	Approved	Unapproved	Approved
6	Banned	Approved	Unapproved	Banned
7	Unapproved	Banned	Unapproved	Banned
8	Approved	Banned	Unapproved	Banned
9	Banned	Banned	Unapproved	Banned
10	Unapproved	Approved By Policy	Unapproved	Approved By Policy
11	Approved (by Policy)	Approved By Policy	Unapproved	Approved By Policy
12	Banned (by Policy)	Approved By Policy	Unapproved	Mixed
13	Unapproved	Banned By Policy	Unapproved	Banned By Policy
14	Approved (by Policy)	Banned By Policy	Unapproved	Mixed
15	Banned (by Policy)	Banned By Policy	Unapproved	Banned By Policy
16	Unapproved	Unapproved	Approved	Approved
17	Approved	Unapproved	Approved	Approved
18	Banned	Unapproved	Approved	Banned
19	Unapproved	Approved	Approved	Approved
20	Approved	Approved	Approved	Approved
21	Banned	Approved	Approved	Banned
22	Unapproved	Banned	Approved	Banned
23	Approved	Banned	Approved	Banned
24	Banned	Banned	Approved	Banned
25	Unapproved	Approved By Policy	Approved (by Policy)	Approved By Policy
26	Approved	Approved By Policy	Approved (by Policy)	Approved By Policy
27	Banned (by Policy)	Approved By Policy	Approved (by Policy)	Mixed
28	Unapproved	Banned By Policy	Approved (by Policy)	Mixed
29	Approved (by Policy)	Banned By Policy	Approved (by Policy)	Mixed
30	Banned (by Policy)	Banned By Policy	Approved (by Policy)	Mixed
31	Unapproved	Unapproved	Banned	Banned
32	Approved	Unapproved	Banned	Banned

#	Certificate State	Publisher State	Certificate Path State	Certificate Global State
33	Banned	Unapproved	Banned	Banned
34	Unapproved	Approved	Banned	Banned
35	Approved	Approved	Banned	Banned
36	Banned	Approved	Banned	Banned
37	Unapproved	Banned	Banned	Banned
38	Approved	Banned	Banned	Banned
39	Banned	Banned	Banned	Banned
40	Unapproved	Approved By Policy	Banned (by Policy)	Mixed
41	Approved (by Policy)	Approved By Policy	Banned (by Policy)	Mixed
42	Banned (by Policy)	Approved By Policy	Banned (by Policy)	Mixed
43	Unapproved	Banned By Policy	Banned (by Policy)	Banned By Policy
44	Approved (by Policy)	Banned By Policy	Banned (by Policy)	Mixed
45	Banned (by Policy)	Banned By Policy	Banned (by Policy)	Banned By Policy
46	Unapproved	Unapproved	Mixed*	Banned
47	Approved	Unapproved	Mixed*	Banned
48	Banned	Unapproved	Mixed*	Banned
49	Unapproved	Approved	Mixed*	Banned
50	Approved	Approved	Mixed*	Banned
51	Banned	Approved	Mixed*	Banned
52	Unapproved	Banned	Mixed*	Banned
53	Approved	Banned	Mixed*	Banned
54	Banned	Banned	Mixed*	Banned
55	Unapproved	Approved By Policy	Mixed* (by Policy)	Mixed
56	Approved (by Policy)	Approved By Policy	Mixed* (by Policy)	Mixed
57	Banned (by Policy)	Approved By Policy	Mixed* (by Policy)	Mixed
58	Unapproved	Banned By Policy	Mixed* (by Policy)	Mixed
59	Approved (by Policy)	Banned By Policy	Mixed* (by Policy)	Mixed
60	Banned (by Policy)	Banned By Policy	Mixed* (by Policy)	Mixed

Mixed and By-Policy States

As [Table 45](#) shows, Certificate Global State can be Mixed, Approved by Policy or Banned by Policy in certain cases. The table shows each case, but these general rules apply:

- **Mixed Path State** – Certificate Path State is considered Mixed when some certificates are Approved and some are Banned. The Mixed state for path is strictly informational, however. In terms of its contribution to Certificate Global State, a Mixed path is equivalent to a Banned path.
- **By Policy Publisher State** – When a Publisher State is Banned by Policy or Approved by Policy, any ban or approval on a certificate in the path is filtered through the publisher policy choices. For example, if a certificate is Banned and it's publisher is Approved by Policy, the Certificate Global Policy is Mixed.

Certificate Ban Setting in Policies

The Advanced Settings tab for each policy includes a *Block files with banned publishers or certificates* setting. This setting must be Active (the default) for certificate bans to have affect file blocking. The certificate setting is effective only in High, Medium and Low Enforcement policies. It affects only *enforcement* of certificate bans, not whether you can assign a ban to a certificate. Also, your choice here does not prevent any certificate from being *approved* or for an approval to be effective on a file.

Advanced						
File Rules						
Custom Rules						
Memory Rules						
Registry Rules						
Publisher Rules						
Rapid Configs						
Computers						
Device Control Settings						
Name	Status	Notifiers				
Block unanalyzed scripts and executables	Active	<default>: Block unanalyzed scripts and execut			Add	Edit
Block unapproved scripts	Active	<default>: Block unapproved scripts			Add	Edit
Block unapproved executables	Active	<default>: Block unapproved executables			Add	Edit
Block banned file names	Active	<default>: Block banned file names			Add	Edit
Block banned file hashes	Active	<default>: Block banned file hashes			Add	Edit
Block executables run from a network drive	Off	<default>: Block executables run from a network			Add	Edit
Block files with banned publishers or certificates	Active	<default>: Block files with banned publishers or			Add	Edit
Enforce memory rules	Active	<default>: Enforce memory rules			Add	Edit
Enforce registry rules	Active	<default>: Enforce registry rules			Add	Edit
Enforce custom (file and path) rules	Active	<default>: Enforce custom (file and path) rules			Add	Edit
Enforce tamper protection	Active	<default>: Enforce tamper protection			Add	Edit
Terminate processes with banned images	Report Only	<default>: Terminate processes with banned in			Add	Edit

Locally approve unapproved files on transition from Visibility or Low Enforcement Level to Medium or High

Interactions with Other Rules

Certificate Global State interacts with other rules and states to contribute to the state of a file. The configuration rules contributing to Certificate Global State were described “[Certificate Global State](#)” on page 352. The following are other rules of potential interest:

- **Enforcement Level** – If Certificate Global State is Banned, it can have an effect on whether files can be executed in High, Medium, and Low Enforcement Levels. If Certificate Global State is Approved, it can have an effect on file execution in High and Medium Enforcement Level.
- **Reputation Rules** – Reputation Rules can affect Publisher State, which can affect Certificate Global State. Keep this in mind if you have already assigned state to individual certificates and then enable or change Reputation Rules. See [Chapter 10](#), “[Reputation Approval Rules](#),” for more details.

How Certificate Global State Affects Global File State

Global File State is a combination of File State, Publisher State, and Certificate Global State. If all certificates in a path are Unapproved, certificates do not contribute to Global File State. If a certificate has a Certificate Global State other than Unapproved, it can play a part in determining Global File State. The two simplest cases are:

- If there are no “by Policy” state settings, then if File State or Publisher State or Certificate Global State is Banned, the Global File State is Banned.
- If there are no “by Policy” state settings, then if none of the three components contributing to Global File State is Banned and at least one is Approved, Global File State is Approved.

Agent Version and Global File State

For agents at v7.0.1 Patch 3 and later, including all v7.2 and 8.0 agents, Global File State is effectively a combination of Certificate Global State and File State – Publisher State is already considered in the calculation of Certificate Global State.

For agents prior to v7.0.1 Patch 3, Global File State remains a combination of Publisher State and File State. Certificate Global State is not involved in Global File State determination.

See [Chapter 7, “File, Publisher, and Application Information,”](#) for more information on how Global File State is determined.

Chapter 12

Managing Devices

This chapter describes features for tracking and control of storage devices detected on computers running the App Control Agent.

Sections

Topic	Page
Overview	361
Devices Managed by App Control	361
Enabling Per-Policy Device Control	362
Managing Specific Devices	365
Viewing Device Information	365
Managing Devices by Model	366
Managing Device Instances	371
Managing Computer-Device Attachments	375

Overview

App Control enables you to track fixed and removable storage devices on agent-managed Windows and Mac computers, and to control file operations that users can perform on those removable devices. App Control device management consists of the following:

- **Policy-specific device control settings** determine whether App Control rules control write and execute operations on devices connected to computers in a policy, and whether this control applies to unapproved devices, banned devices, or both.
- **Device-specific rules** allow you to explicitly approve or ban specific removable devices, either by *model* or by *individual device*, so that files can be written or executed on approved devices while banned or unapproved devices may be restricted by your policy settings. The behavior of these approval and ban rules is similar to the behavior of file approvals and bans in App Control.
- **Device inventory** tables show each device discovered by a App Control Agent, and make it possible for you to implement the device-specific rules. This inventory includes a list of device models, a list of individual devices, and a list of unique *attachments* of an individual device and an individual computer. You can drill down on any instance in these lists.

Throughout this chapter, the term *individual device* means one specific device that can only be attached to one computer at a time. Generally, this means a specific model plus a unique serial number (at least unique for that model).

Platform Note

All device visibility and control features are available for computers running Windows agents, including the agents built into all recent server releases.

Beginning with agent release 7.2.3 Patch 10, device visibility and control features are available on the Mac platform. See the [Carbon Black User Exchange](#) for the release announcement, release notes (including information about Mac-specific differences and known issues), and downloadable installer for App Control 7.2.3 Patch 10 Mac Agent (and later versions when available).

Device management features are not currently available for Linux computers.

Devices Managed by App Control

The App Control Agent can detect several different kinds of devices on computers. In general, if a device has an identifiable file system, it is added to the Devices tables. How a detected device is managed depends upon whether it is identified as fixed or removable:

- **Fixed devices** are included in the device inventory, but they cannot be approved, banned, or blocked by App Control rules.
- **Removable devices** are included in the device inventory, and they can be approved, banned, and blocked by App Control rules.

App Control must rely on the information provided by a device to determine whether it is fixed or removable, and there are some cases in which the information is incorrect.

Specific categories of devices detected by App Control Agents include:

- IDE Devices
- SATA Devices
- SCSI Devices
- USB Devices
- FireWire (IEEE 1394) Devices
- Serial Bus Protocol 2 Devices
- Floppy Disk Drives

The USB devices detected may include solid-state “stick”-type drives, CD/DVD drives, and media card readers. For all CD/DVD drives and for card readers on macOS, the drive itself, not the media it reads, appears in the devices table. On Windows agents, it is possible that both a media card and the card reader may be reported on the Devices page.

Note

In addition to the device settings and rules described here, you can create custom path rules that affect what a device can or can't do. See [“Specifying Devices in Paths in Windows Rules”](#) on page 409 in [Chapter 14, “Custom Software Rules,”](#) for more information.

Enabling Per-Policy Device Control

For any of the device control features in App Control to be enabled, you must activate device control settings on policies. Each policy can have its own device control configuration. These settings allow you to activate blocking for any combination of the following:

- banned devices and/or unapproved devices
- write and/or execute operations

You cannot block read operations on devices, but you can enable reporting so that when a file is read on a banned or unapproved device, an event is generated.

You enable device control on the Edit Policy page for policies that have already been created. Device Control Settings do not appear on the Add Policy page for a new policy you are creating.

For policies in Visibility mode, you can choose any device control setting, but no device operations are blocked. To block device activity, a policy must be in Control mode.

Note

The effect of the settings on drives with removable *media*, such as CD/DVD drives, differs from the effect on devices with non-removable media. Burning a CD or DVD does not constitute a “Write” operation. If you want to block burning of CD/DVD media, ban the media-burning software application.

[Table 46](#) shows the effects of specific choices for Device Control settings.

Table 46: Device Control Setting Behavior

Setting	Active	Off	Report Only
Block writes to unapproved removable devices	Tracks write operations to unapproved removable devices and blocks them in all Control mode policies (High, Medium and Low Enforcement). Notes: <ul style="list-style-type: none"> All devices are unapproved by default, so be certain you want to block everything you haven't explicitly approved before activating this setting. Blocking writes to removable devices does not block writes to CD/DVD media. 	Permits write operations to removable devices; does not report the event.	Permits write operations and reports them as events.
Block writes to banned removable devices	Tracks write operations to banned removable devices and blocks them in all Control mode policies (High, Medium and Low Enforcement). Note: Blocking writes to removable devices does not block writes to CD/DVD media.	Permits write operations to banned removable devices; does not report the event.	Permits write operations and reports them as events.
Report reads from unapproved removable devices	Choice not available.	Permits reads from unapproved removable devices; does not report the event.	Permits reads and reports them as events.
Report reads from banned removable devices	Choice not available.	Permits reads from banned removable devices; does not report the event.	Permits reads and reports them as events.
Block execution from unapproved removable devices	Tracks execution of files on unapproved removable devices and blocks them in all Control mode policies (High, Medium and Low Enforcement). Note: All devices are unapproved by default, so be certain you want to block all devices not explicitly approved before activating this setting.	Permits files on unapproved removable-device to execute unless the file itself is banned by another rule; does not report the event.	Permits executions and reports them as events.

Setting	Active	Off	Report Only
Block execution from banned devices	Tracks execution of files on banned removable devices and blocks them in all Control mode policies (High, Medium and Low Enforcement).	Permits execution of files on banned removable-device unless the file is banned by another rule; does not report the event.	Permits executions and reports them as events.

In the Default, Template and Local Approval policies, device controls are all set to *Off* (no blocking or reporting) except for the settings that block writes and executions to banned devices, which are *Active*. You can change this for all except the Local Approval Policy. Changing the settings in the Template Policy *before* you create other policies can save time in policy configuration.

To enable device control for a policy:

1. On the console menu, choose **Rules > Policies**. The Policies page opens.
2. On the Policies page, click the View Details button next to the name of the policy whose device settings you want to edit. The Edit Policy page opens.
3. Click the **Device Control Settings** tab.

Edit Policy Help Desk

Policy Name: Help Desk
 Description: Policy specific to Help Desk employees.

Mode: Visibility Control Disabled

Options: Allow Upgrades Track File Changes
 Load Agent in Safe Mode Suppress Logo in Notifier

Total Computers: 0
 Connected Computers: 0

Advanced File Rules Custom Rules Memory Rules Registry Rules Publisher Rules Rapid Configs Computers **Device Control Settings**

Name	Status	Notifiers
Block writes to unapproved removable devices	Off	<default>: Block writes to unapproved removab
Block writes to banned removable devices	Active	<default>: Block writes to banned removable de
Report reads from unapproved removable devices	Off	<none>
Report reads from banned removable devices	Off	<none>
Block executions from unapproved removable devices	Off	<default>: Block executions from unapproved n
Block executions from banned removable devices	Active	<default>: Block executions from banned remo

Save & Exit Save Cancel Reset Policy

4. On the Device Control Settings panel, choose **Active** for any setting you want to enable, **Off** for any setting you want to disable, and **Report Only** for any setting for which you want the App Control Server to report file activity on devices but not enforce the setting. Note that you cannot block Read access to devices, so Active is not a choice for the two Read settings. See [Table 46, “Device Control Setting Behavior,”](#) on page 363 for details about the effects of each setting.
5. You can change (or eliminate) the notifier that appears when a device setting blocks file access. To do this, make a choice on the Notifier menu next to each setting whose notifier you want to change. See [Chapter 20, “Endpoint Notifiers and Approval Requests,”](#) for more options and more information.

6. When the Device Settings and their notifiers are edited to your preferences, click the **Save** button (to remain on the page) or the **Save & Exit** button. Your changes are saved for that policy.
7. Repeat this procedure for each policy whose Device Settings you want to change.

Managing Specific Devices

App Control collects many different kinds of information about the devices it detects on your computers. You can use this information to make decisions about how you want to treat file activities on devices.

By default, all devices are in an unapproved state (neither approved nor banned). You can explicitly approve or ban specific removable devices, either by model or by serial number. Files not blocked by other rules are always allowed to execute and be written on approved devices. Treatment of unapproved and banned files varies depending upon the Device Control Settings for each policy.

Note

Banned devices do not block in policies that are set to Visibility mode, but you can choose Report Only for the Device Settings to generate events for device-related activity that would have blocked in Control mode.

Similarly, device-specific bans and approvals do not block or allow access in policies that do not have Device Settings set to *Active*.

Viewing Device Information

Device information is presented in table form on the Devices page, which you access by choosing **Assets > Devices** on the console menu. From each device table, you can drill down to a details page for any single item on the page (model, device instance, or attachment) by clicking on the View Details button next to the item. The following table shows the type of information available in each of these views:

This Device information...	...is listed in this Table	...and this Details page for each Table row
Device Models found (vendor plus name)	Device Catalog (<i>Show Individual devices</i> box not checked)	Device Model Details (for one model)
Individual Devices found (unique serial number)	Device Catalog (<i>Show Individual devices</i> box checked)	Device Details (for one serial number)
Individual Devices attached to Individual Computers	Devices on Computers	Device Attachment Details (for one device-computer pair)

The Device tables do not have Saved Views, but the *Group By* menu allows you to group information by different fields. For example, you might want to see all of the devices grouped by *vendor*, or view all devices models for which certain serial numbers have rules that are an *exception* to the rule for the model. The Group By menu provides options for

each of these cases. If you have not already become familiar with modifying views, see [“Console Tables”](#) on page 66.

Managing Devices by Model

You can monitor and manage devices attached to computers by their model. Managing devices by model provides a way to control many devices with a single rule. You can:

- View the full list of device models in the Device Catalog.
- View complete information about one device model on the Device Model Details page. You can view other information *related* to a device model by using the Related Views menu.
- Approve, ban, and remove approvals or bans from either the Device Catalog or the Device Model Details page.

Viewing Device Models in the Device Catalog

Device models are identified as a specific pairing of vendor and product name. The Device Model table provides general information about the types of devices connected to your computers, and allows approving or banning all instances of a device model.

To view all device models detected by App Control:

1. On the console menu, choose **Assets > Devices**. The Devices page appears.
2. If it is not already the active tab, click on the **Device Catalog** tab. The Device Catalog table appears on the page.
3. Scroll to the bottom of the page, and if the *Show individual devices* checkbox is checked, click on it to remove the checkmark. The Device Catalog shows the table of device models.

Devices: Storage Device Catalog

Device Catalog Devices on Computers

Group By: State | Ascending

Show Filters | Show Columns | Export to CSV | Refresh Page

Action

<input type="checkbox"/>	State	Exceptions	Vendor	Name	Friendly Name	Acknowledged	Computer Count
State: Unapproved 2 items							
<input type="checkbox"/>	Unapproved	No	VMWARE	VIRTUAL_DISK	VMware Virtual disk SCSI Disk Device	No	2
<input type="checkbox"/>	Unapproved	No	NECVMWAR	VMWARE_IDE_CDR10	NECVMWar VMware IDE CDR10 ATA Device	No	1

2 items in 1 group Page 1/1 25 rows per page

Show individual devices

See [Table 47, “Device Model Details,”](#) on page 368 for a description of the columns that can be displayed in this table.

The Action menu in the Device Catalog for models acts on checked table rows. It includes the following commands:

- Globally Approve
- Globally Ban
- Remove Approval or Ban
- Acknowledge

The approval and ban commands are described in [“Approving and Banning Device Models”](#) on page 369. You can use the Acknowledge command to indicate that you have reviewed a particular model and perhaps taken any action you intend to take on its status. You can then sort or filter the table so that device models you haven’t yet acknowledged are more visible.

Viewing Details for One Device Model

The Device Model Details page provides information about the model. [Table 47, “Device Model Details,”](#) describes the fields shown on this page.

Device Model Details

General

Vendor: NECVMWAR
 Name: VMWARE_IDE_CDR10
 Class: IDE Device
 Friendly Name: NECVMWar VMware IDE CDR10 ATA Device
 Removable Device: Yes
 Acknowledged: No
 Description:
 Device Count: There is 1 individual device for this model.
 Computer Count: This device model was attached to 1 computer

Rule

State: Select the default state for this device model...

Approved Serial Numbers: Approve only serial numbers that match...

Banned Serial Numbers: Ban only serial numbers that match...

Rule Applies To: All Current and Future Policies
 Selected policies

History

Jun 29 2020 08:58:22 PM This device model was first seen on WORKGROUP\APC-79895

Related Views

- All devices of this model
- All computers with this device model
- All events for this device model

The Device Model Details page is also where you configure the rule for how devices of this model are treated. This is done on the page itself rather than on a menu. The rule includes the overall state of the model and any exceptions for specific serial numbers.

The Related Views menu provides links to the following information:

- **All devices of this model** – Filters the Device Catalog to show all instances of this device model that have been attached to agent-managed computers.
- **All computers with this device model** – Filters the Devices on Computers table to show all computers to which devices of this model have been attached.
- **All events for this device model** – Goes to the Events page and filters it to show all events related to this device model, including initial discovery of each instance and any time a device of this model has been attached or detached from a computer.

Table 47: Device Model Details

Field	Description
Vendor	The brand of the device (e.g., "SanDisk"). If the device does not have detectable vendor information, this field might show something like "USB DISK" or "Flash".
Name	The name of the device model, which might be a trade name (e.g., "Jumpdrive Pro") or a model number (e.g., "c30w"). If the device does not have detectable model name, this field might show something like "USB Storage Device" or "Unnamed Product".
Class	This is primarily a description of the interface for the device. The choices are IDE Device, SATA Device, SCSI Device, USB Device, FireWire (IEEE 1394) Device, Serial Bus Protocol 2, Floppy Disk, and Unknown.
Removable Device	Whether the device is removable or not removable. Values are Yes or No . Note that some devices might not provide accurate information for this field.
Friendly Name	The common name for this device, for example, as you would see it in Windows Explorer or Mac Finder when the device is connected. Platform Note: MacOS relies on the Disk Arbitration Framework to deliver the friendly name. This uses partition block rather than device information, and may yield a different friendly name on MacOS than the same device shows in Windows.
Acknowledged*	You may Acknowledge a device to indicate that you have seen it and perhaps do not need to track it as closely. Acknowledging a device does not change its approval state. You can choose Yes or No for this field using the Action menu or a menu on the details page.
Description	Editable text providing any information you would like to include with the record of this device model.
Device Count	The number of unique devices (i.e., unique serial numbers) of this model detected by App Control on your computers.
Computer Count	The number of computers to which a device of this model has been attached.
First Seen Platform	The first platform (Windows or Mac) on which this device model was seen. Device management is not currently available for Linux.
State	The default state for this device model. The choices are Approved, Banned, and Unapproved. Specific instances (serial numbers) of a model can have a state that differs from the default model state.
Approved Serial Numbers	If the default state of the device model is Unapproved or Banned, you can specify serial numbers that are Approved. You can enter one or more specific serial numbers, or a pattern that uses wildcards to include a range of numbers.
Banned Serial Numbers	If the default state of the device model is Unapproved or Approved, you can specify serial numbers that are Banned. You can enter one or more specific serial numbers, or a pattern that uses wildcards to include a range of numbers.

Field	Description
Rule Applies To	You can make a device model rule apply to computers in all policies or only certain policies.
History	Records the date and time when the device was discovered and when rules affecting it were applied or changed.

Approving and Banning Device Models

There are two options for managing device model approvals and bans:

- In the Device Catalog, you can check one or more device models in the table and use the Action menu to approve, ban, or remove the approval or ban for all of the checked items.
- On the Device Model Details page, you can approve, ban, or remove the approval or ban for the device model listed on the page. You also can view, add and delete exceptions (by serial number) to the default rule for the model, and you can make the rule apply to all policies or only certain policies.

To approve one or more device models from the Device Catalog:

1. On the console menu, choose **Assets > Devices**. The Devices page appears.
2. Click on the Device Catalog tab, and in the lower right corner of the catalog page, make sure the *Show individual devices* box is *not* checked. This displays the table entitled *Devices: Storage Device Catalog*.
3. Check the box next to each device model you want to approve and then choose **Globally Approve** on the Action menu.
4. Choose **OK** on the confirmation dialog. The device models will be approved, and all instances of the device model will be approved by default.

To ban one or more models, use the procedure above and substitute **Globally Ban** for the Action menu choice in Step 3.

To remove approvals or bans from one or more models, use the procedure above and substitute **Remove Approval or Ban** for the Action menu choice in Step 3.

Notes

- Only devices identified as removable can be approved or banned. If any *fixed* devices are checked when you attempt to approve or ban models from the Device Catalog, you will see an error message and the non-removable drives will not be affected. If any *removable* devices are included in the selection, they will be affected by the command even if other devices are not. You can determine whether a device can be approved or banned by checking the *Removable Device* column in the table.
- All approval and ban actions taken from the Device Catalog are global, affecting all device instances and computers in all policies. If you want to limit an approval or ban to devices on computers in particular policies, or if you want to add exceptions to the rule for specific device serial numbers, use the Device Model Details page.
- You can select combinations of Banned and Approved models when you use the Remove Approval or Ban command – all will be moved to the Unapproved state.

To approve one device model from the Device Model Details page:

1. On the console menu, choose **Assets > Devices**. The Devices page appears.
2. Click on the Device Catalog tab, and in the lower right corner of the catalog page, make sure the *Show individual devices* box is *not* checked. This displays the table entitled *Devices: Storage Device Catalog*.
3. Click on the View Details button next to the device model you want to approve. The Device Model Details page appears.
4. If you want to limit this approval to certain policies, click the **Selected policies** radio button and check the boxes next to the policies you want enabled.
5. On the State menu, choose **Approved**.
6. If you want to ban certain instances of this device model even though you are approving the model itself, enter one or more serial numbers (or a serial number pattern with wildcards) into the *Banned Serial Numbers* field.

You also can add exceptions later by approving or banning device instances in the Device Catalog or Devices on Computers tables, or by using the approve or ban commands in the Device Instance Details or Device Attachment Details page.

7. Click the **Save** button at the bottom of the page and click **OK** on the confirmation dialog. The device model will be approved, and all instances except those you created exceptions for will be approved.

To ban a model from its details page, use the procedure above and choose **Banned** for the State menu choice in Step 5. If you want to create exceptions and you know their serial numbers, enter the numbers or a pattern to match in the *Approved Serial Numbers* field.

To remove a model approval or ban using the details page, use the procedure above and substitute **Unapproved** for the Action menu choice in Step 3.

Managing Device Instances

You can monitor and manage individual devices, as identified by their serial number. Managing devices by instance provides a way to control specific devices for which you might want different treatment than others devices of the same model. You can:

- View the full list of device instances in the Device Catalog.
- View complete information about one device instance on the Device Details page. You also can see other information related to a device through Related Views.
- Approve, ban, and remove approvals or bans from either the Device Catalog or the Device Details page.

Viewing Instances in the Device Catalog

A device instance is identified by its serial number, vendor and name. The device instance view can be useful for information about the number of devices on your computers, and for approving or banning specific device instances.

To view all unique device instances detected by App Control:

1. On the console menu, choose **Assets > Devices**. The Devices page appears.
2. Click on the **Device Catalog** tab. The Device Catalog table appears on the page.
3. On the bottom of the page, make sure the *Show individual devices* checkbox is checked. The Device Catalog shows the device instances with unique serial numbers.

	State	Vendor	Name	Friendly Name	Device Class	Serial Number	First Seen Computer	Computer Count
State: Approved 2 items								
<input type="checkbox"/>	Approved	SanDisk	Cruzer Fit	SanDisk Cruzer Fit U...	USB Device	44D40017611...	MYCORPLAPTOP-4	1
<input type="checkbox"/>	Approved	IronKey	IronKey	IronKey Secure Drive...	USB Device	860301065a69...	MYCORPDESKTOP-1	4
State: Approved by Policy 2 items								
<input type="checkbox"/>	Approved by Policy	HTC	Android Phone	HTC Android Phone ...	USB Device	HT12ABC0123...	MYCORPLAPTOP-5	1
<input type="checkbox"/>	Approved by Policy	HTC	Android Phone	HTC Android Phone ...	USB Device	HT19DAP1197...	MYCORPLAPTOP-3	1
State: Unapproved 3 items								
<input type="checkbox"/>	Unapproved	(Standard...		SDHC Card	IDE Device	448222011&08...	MYCORPDESKTOP-1	1
<input type="checkbox"/>	Unapproved			USB Device	USB Device	447DA080&0.0...	MYCORPLAPTOP-12	1
<input type="checkbox"/>	Unapproved	(Standard...		SDHC Card	IDE Device	272ABC057&0...	MYCORPDESKTOP-7	1
State: Unapproved by Policy 5 items								

12 items in 4 groups Page 1/1 25 rows per page

See [Table 48, “Device Details \(unique serial number\) and Device Attachment Details,”](#) on page 373 for a description of the columns that can be displayed in this table.

The Action menu in the Device Catalog for instances acts on checked table rows. It includes the following commands:

- Globally Approve
- Globally Ban
- Remove Approval or Ban
- Acknowledge

The approval and ban commands are described in [“Approving or Banning Device Instances”](#) on page 374. You can use the Acknowledge command to indicate that you have reviewed a particular device instance and perhaps taken any action you intend to take on its status. You can then sort or filter the table so that device models you have not yet acknowledged are more visible.

Viewing Details for One Device Instance

The Devices Details page shows the information about one unique device (with a unique serial number). [Table 48, “Device Details \(unique serial number\) and Device Attachment Details,”](#) on page 373 describes the fields shown on this page.

Device Details ?

General

Vendor: NECVMWAR
Name: VMWARE_IDE_CDR10
Class: IDE Device
Friendly Name: NECVMWar VMware IDE CDR10 ATA Device
Removable Device: Yes
Serial Number: 5&2A3267F0&0&1.0.0
Default State: Unapproved
Device State: Unapproved
First Seen Computer: WORKGROUP\APC-79895
Platform: Windows
First Seen Date: Jun 29 2020 08:58:22 PM
Computer Count: This individual device was attached to 1 computer

Close

Related Views

- Model details
- All computers with this device
- All events for this device

Actions

- Ban Serial Number
- Approve Serial Number

The Device Details page includes an Actions menu and a Related Views menu.

The Actions menu includes commands for approving and banning this device, and for removing approvals or bans. The commands that appear depend on the current state of the device. See [“Approving or Banning Device Instances”](#) on page 374 for more information about using these commands.

The Related Views menu provides links to the following information:

- **Model details** – Goes to the Device Model Details page for this device, which shows both information about the model itself and the default rule definitions for the model.
- **All computers with this device** – Filters the Devices on Computers table to show all computers to which this device instance has been attached.
- **All events for this device** – Goes to the Events page and filters it to show all events related to this device instance (by serial number), including its initial discovery and the dates and times it has been attached or detached from a computer.

Table 48: Device Details (unique serial number) and Device Attachment Details

Field	Description
Vendor	The brand of the device (e.g., "SanDisk"). If the device does not have detectable vendor information, this field might show something like "USB DISK" or "Flash".
Name	The name of the device model, which might be a trade name (e.g., "Jumpdrive Pro") or a model number (e.g., "c30w"). If the device does not have detectable model name, this field might show something like "USB Storage Device" or "Unnamed Product".
Class	Mainly a description of the interface for the device. The choices are IDE Device, SATA Device, SCSI Device, USB Device, FireWire (IEEE 1394) Device, Serial Bus Protocol 2, Floppy Disk, and Unknown.
Removable Device	Whether the device is removable or not removable. Values are Yes or No . Note that some devices might not provide accurate information for this field.
Friendly Name	The common name for this device, for example, as you would see it in Windows Explorer or Mac Finder when the device is connected. This is often some combination or variant of the Vendor and Name. Platform Note: MacOS relies on the Disk Arbitration Framework to deliver the friendly name. This uses partition block rather than device information, and may yield a different friendly name on MacOS than the same device shows in Windows.
Serial Number	The serial number that identifies this unique individual device.
Default State	The default state for this device (which is the state for its <i>model</i>). The choices are Approved, Banned, and Unapproved. Note that this specific instance might have a state that differs from the default.
Device State	The actual state for this individual device (as identified by serial number). The choices are Approved, Banned, and Unapproved.
Platform	For Device Details, the platform (Windows or Mac) of the computer on which the device was first detected. For Device Attachment Details, the platform (Windows or Mac) of the computer for this attachment. Platform Note: Device management is not currently available for Linux.
Computer Count	The number of different computers to which this individual device has been connected.
Fields on the Device Details page only	
First Seen Computer	On the Device Details page, the computer on which this individual device was first detected by a App Control Agent.
First Seen Date	On the Device Details page, the date and time when this individual device was first detected by a App Control Agent.

Field	Description
Fields on the Device Attachment Details page only	
Current Status	On the Device Attachment Details page, whether the device and computer that define this attachment are currently Attached or Detached. Note: Device attachment status for computers disconnected from the App Control Server is the last known status when the computer was connected.
First Attach Date	On the Device Attachment Details page, the date and time when the device and computer were first attached.
Last Attach Date	On the Device Attachment Details page, the date and time when the device and computer were last attached.
Last Detach Date	On the Device Attachment Details page, the date and time when the device was last detached from the computer.

Approving or Banning Device Instances

There are two options for managing device instance (serial number) approvals and bans:

- In the Device Catalog or Devices on Computers page, you can check one or more device instances in the table and use the Action menu to approve, ban, or remove the approval or ban, for all of the checked items.
- On the Device Details page or Device Attachment Details page, you can approve, ban, or remove the approval or ban for the device instance listed on the page.

You only need to approve, ban, or remove approvals or bans from an instance if you want it to have a state other than the default state for its device model. Instance-specific exceptions appear on the Device Model Details page for the device model.

To approve one or more device instances from the Device Catalog:

1. On the console menu, choose **Assets > Devices**. The Devices page appears.
2. Either:
 - Click on the **Device Catalog** tab, and in the lower right corner of the catalog page, make sure the *Show individual devices* box is checked. This displays the table *Devices: Individual Storage Devices*.
 - **- or -**
 - Click on the **Devices on Computers** tab.
3. Check the box next to each device instance you want to approve and then choose **Globally Approve** on the Action menu.
4. Choose **OK** on the confirmation dialog. The device will be approved by serial number.

To ban one or more instances, use the procedure above and substitute **Globally Ban** for the Action menu choice in Step 3.

To remove approvals or bans from one or more instances, use the procedure above and substitute **Remove Approval or Ban** for the Action menu choice in Step 3.

Notes

- Only devices identified as removable can be approved or banned. If any *fixed* devices are checked when you attempt to approve or ban devices, you will see an error message and the non-removable drives will not be affected. If any *removable* devices are included in the selection, they will be affected by the command even if other devices are not. You can determine whether a device can be approved or banned by checking the *Removable* column in the table.
- All approval and ban actions taken on device *instances* become exceptions within the rule for their device *model*, and are applied to all policies or selected policies as specified in the model rule.
- You can select combinations of Banned and Approved devices when you use the Remove Approval or Ban command – all will be moved to the Unapproved state.

To approve or ban one device instance (Device or Attachment Details page):

1. On the console menu, choose **Assets > Devices**. The Devices page appears.
2. Either:
 - Click on the **Device Catalog** tab, and in the lower right corner of the catalog page, make sure the *Show individual devices* box is checked. This displays the table *Devices: Individual Storage Devices*.
 - **- or -**
 - Click on the **Devices on Computers** tab.
3. Click on the View Details button next to the device instance you want to approve. The Device Details or Device Attachment Details page appears.
4. In the Actions menu on the right side of the page:
 - To **approve** the device instance, choose **Approve Serial Number**. The device will be approved, and its serial number will be added as an exception on the Device Model Details page for its model.
 - To **ban** the device instance, choose **Ban Serial Number** as the Actions menu choice.

To remove a device instance approval or ban using the details page, use the procedure above and substitute the appropriate removal command.

Managing Computer-Device Attachments

You can monitor attachments between a specific device instance and a specific computer, and manage the individual devices. You can:

- View the full list of device-computer attachments in the Devices on Computers table.
- View complete information about an attachment between one specific device and one specific computer on the Device Attachment Details page. You also can see other information related to this attachment or the individual device through Related Views.
- Approve, ban, and remove approvals or bans from either the Devices on Computers table or the Device Attachment Details page.

Viewing Devices on Computers

The Devices on Computers tab provides a table of individual devices that have been connected to individual computers. The relationship between one device and one computer counts as a single “attachment” in the table, regardless of how many times the two have been connected and disconnected. If you are concerned about the use of removable devices on a particular computer, the Devices on Computers page provides a way to find out if any such connections exist. You can approve and ban individual devices from this table.

To view all attachments between a specific device and a specific computer:

1. On the console menu, choose **Assets > Devices**. The Devices page appears.
2. Click on the **Devices on Computers** tab. The Devices on Computers page shows each pairing of a device instance (unique serial number) and a specific computer.

	State	Vendor	Name	Friendly Name	Serial Number
<input type="checkbox"/>	Unapproved	VMWARE	VIRTUAL_DISK	VMware Virtual disk SCSI Disk Device	5&1EC51BF7&0&000000
<input type="checkbox"/>	Unapproved	VMWARE	VIRTUAL_DISK	VMware Virtual disk SCSI Disk Device	5&1EC51BF7&0&000100
<input type="checkbox"/>	Unapproved	NECVMWAR	VMWARE_IDE_CDR10	NECVMWar VMware IDE CDR10 ATA Device	5&2A3267F0&0&1.0.0

3 items Page 1/1

See [Table 48, “Device Details \(unique serial number\) and Device Attachment Details,”](#) on page [373](#) for a description of the columns that can be displayed in this table.

The Action menu in the Devices on Computers table instances acts on checked table rows. It includes the following commands:

- Globally Approve
- Globally Ban
- Remove Approval or Ban
- Acknowledge

The approval and ban commands on both the Devices on Computers table and the Device Catalog for *instances* affect the instance, as defined by serial number, in the checked rows. You are not approving or banning a particular *attachment*. See [“Approving or Banning Device Instances”](#) on page [374](#) for more details.

You can use the Acknowledge command to indicate that you have reviewed a particular device instance and perhaps taken any action you intend to take on its status. You can then sort or filter the table so that device models you have not yet acknowledged are more visible.

Viewing Details for One Computer-Device Attachment

The Devices Attachment Details page shows information about the history of attachment between one device instance and one computer. [Table 48, “Device Details \(unique serial number\) and Device Attachment Details,”](#) on page 373 describes the fields shown on this page.

Device Attachment Details
?

General

Vendor: NECVMWAR
Name: VMWARE_IDE_CDR10
Class: IDE Device
Friendly Name: NECVMWar VMware IDE CDR10 ATA Device
Removable Device: Yes
Serial Number: 5&2A3267F0&0&1.0.0
Default State: Unapproved
Device State: Unapproved
Computer: WORKGROUP\APC-79895
Platform: Windows
Current Status: Detached
First Attach Date: Jun 29 2020 08:58:38 PM
Last Attach Date:
Last Detach Date:
Computer Count: This individual device was attached to 1 computer

Close

Related Views

Model details

All computers with this device

All events for this device

Actions

Ban Serial Number

Approve Serial Number

The Device Attachment Details page includes an Action menu and a Related Views menu.

The Action menu includes commands for approving and banning this device instance, and for removing approvals and bans. The commands that appear depend on the current state of the device. See [“Approving or Banning Device Instances”](#) on page 374 for more information about using these commands.

The Related Views menu provides links to the following information:

- **Model details** – Goes to the Device Model Details page for this device, which shows both information about the model itself and the default rule definitions for the model.
- **All computers with this device** – Filters the Devices on Computers table to show all computers to which this device instance has been attached.
- **All events for this device** – Goes to the Events page and filters it to show all events related to this device instance (by serial number) *on this computer*, including its initial discovery and any time it has been attached or detached from a computer.

Chapter 13

Script Rules

This chapter describes Script Rules, which identify files to be tracked and managed as scripts by App Control. The App Control Server includes built-in script rules, and you can create custom rules to identify other scripts.

Sections

Topic	Page
Overview	379
Script Rules and Other App Control Rules	382
Policy Settings for Script Rules	383
Creating a Custom Script Rule	384
Editing a Script Rule	387
Disabling or Deleting a Script Rule	388
Viewing Rule Status on Computers	389
Script Rule Examples	389

Overview

App Control tracks and manages two categories of files: *executables* and *scripts*. Executables are identified based on App Control's analysis of their content. Scripts are identified by name, with the exception of certain non-Windows shell scripts.

What is a Script?

A *script* is a file that contains executable or interpretable content that has meaning only in the context of a *script processor*. This dependency on a specific host process is what differentiates a script from typical executables. Script rules require two specifications:

- a *script type* file pattern definition to allow identification of the script file.
- a *script processor* specification that identifies the file that processes the script identified by the script type. You can either specify a string to match for the processor or, for Windows computers, let the File Association list on each agent computer determine the default processor for a file matching the script type. Only one processor may be specified for a script type, even if there are multiple compatible processors.

Examples of script files include Visual Basic scripts (*.vbs), batch scripts (*.bat and *.cmd), and shell scripts (*.sh, *.csh, etc.). Scripts might also be add-ons or extensions for browsers, such as Firefox XPI plug-ins, or application data files such as Word documents (*.docx). Some files, such as [Chrome *.crx](#) extensions, are not scripts by definition; however, they are compressed files that may contain [.JS](#), [.JSON](#), and executable programs, and are thus considered "script files" for the purpose of tracking and management.

Examples of script processors include cmd.exe (batch scripts), bash (shell scripts), wscript.exe (Visual Basic scripts), and processes that are not obviously script processors such as firefox.exe, chrome.exe and word.exe.

Certain scripts are identified by their content and may be subject to executable rules rather than the script rules. See ["Shell Scripts Identified by Content"](#) on page 383.

Notes

- App Control monitors and controls scripts that use script and processor file names that can be identified and defined in a rule. Script processing that takes place in browser memory, such as with JavaScript, is not a candidate for control by App Control script rules.
- You can configure and enable a set of rules that ensure that Windows script processors only run from expected locations. See [Chapter 18, "Rapid Configs,"](#) for more information.
- Any file smaller than 4 bytes will not be inventoried or tracked. Therefore files shorter than 4 bytes cannot be blocked by a Script Rule. However, name based rules, such as a Custom Rule, will still apply to these files and can be used to block them.

What App Control Script Rules Do

Script rules implement two types of action for files matching the rules:

- **Visibility:** When a file matching the script type in a rule is discovered, either because it is newly present on an agent computer or because a new rule was created, the file is added to the File Catalog and Files on Computers tables, and is tracked from that point forward. Although identified by name, a script file is hashed like other identified as “interesting” by App Control, and its hash is stored in the file database.
- **Control:** When a file matching a script processor attempts to access a file identified as a script type in the same rule, that is considered a *script execution*. For enabled rules, script executions are controlled according to the policy settings for the computer on which the execution is attempted and any other applicable App Control rules.

The file state of a script identified by App Control depends upon when it was discovered and on the state of the setting *Rescan Computers: Check to approve all existing scripts matching this definition*. If the Rescan Computers box is *not* checked, all scripts of the type identified by the rule are treated as unapproved when executed. If the Rescan Computers box *is* checked, script files currently on agent-managed computers at the time of the rescan are *locally approved* and (unless banned by a rule) allowed to execute under all Enforcement Levels. Script files discovered after the rescan are considered Unapproved, and their execution will be blocked at High or Medium Enforcement Levels.

Note

Because of the rescanning requirement described above, it is best to create and enable Script Rules before you install agents on computers, whenever possible.

Pre-configured Script Rules

App Control includes several standard script rules, some of which are enabled by default. On the Script Rules page, you can enable or disable existing rules, modify the rules, and create new custom script rules.

Software Rules						
[Updates] [Rapid Configs] [Publishers] [Users] [Directories] [Files] [Custom] [Memory] [Registry] [Scripts] [Reputation]						
Saved Views:		Group By:				
(none) [Add]		(none) [Ascending]				
[Show Filters] [Show Columns] [Export to CSV] [Refresh Table]						
[Add Script Rule] Search: [] <input type="checkbox"/> Automatically apply Showing 18 out of 18 item(s)						
Name	Type	Process	Enabled	Date Modified	Last Modified By	
[] [] Batch	*.cmd, *.bat	<System>\cmd.exe, <Systemx...	Yes	Jun 29 2020 07:00:12 PM	System	
[] [] Registry	*.reg	<System>\reg.exe, <Systemx8...	Yes	Jun 29 2020 07:00:12 PM	System	
[] [] Visual Basic	*.vb, *.vbe, *.vbs, *.wsh, *.wsf	<System>\cscript.exe, <Syste...	Yes	Jun 29 2020 07:00:12 PM	System	
[] [] Java	*.class, *.jar	*\java.exe, *\javaw.exe, <Yara...	Yes	Jun 29 2020 07:00:12 PM	System	
[] [] Perl	*.pl, *.pm	*\perl.exe, <YaraTags:perl_int...	No	Jun 29 2020 07:00:12 PM	System	
[] [] Python	*.py, *.pyc, *.pyo, *.pyw	*\python.exe, *\pythonw.exe, ...	No	Jun 29 2020 07:00:12 PM	System	
[] [] PowerShell	*.ps1, *.psm1	*\powershell.exe, <YaraTags:p...	Yes	Jun 29 2020 07:00:12 PM	System	

Table 49 shows the standard script rules. Where the file extension is the same for different rules, a different process (or process path) is paired with the file extension. If you are upgrading from a pre-8.1.6 release, please see “[Windows Script Rules Changes on Upgrade](#)” on page 382 in addition to the table.

Table 49: Standard Script Rules and File Extensions

Application or Category	Script Extensions	Processes	Platform	Default State
Linux Shell	.sh, .csh, .zsh, .ksh	/bin/bash, /bin/csh, /bin/ksh, /bin/sh, /bin/tcsh, /bin/zsh /bin/dash, /bin/static-sh, /bin/busybox	Linux	Enabled
Linux Perl	.pl	/usr/bin/perl	Linux	Enabled
Linux Python	.py	/usr/bin/python	Linux	Enabled
Mac Shell	.sh, .csh, .zsh, .ksh	/bin/bash, /bin/csh, /bin/ksh, /bin/sh, /bin/tcsh, /bin/zsh	Mac	Enabled
Mac Perl	.pl	/usr/bin/perl	Mac	Enabled
Mac Python	.py	/usr/bin/python	Mac	Enabled
Batch	.cmd, .bat	<System>\cmd.exe <Systemx86>\cmd.exe <YaraTags:cmd_interpreter>*	Windows	Enabled
Registry	.reg	<System>\reg.exe <Systemx86>\reg.exe <System>\regedt32.exe <Systemx86>\regedt32.exe <Windows>\regedit.exe <Systemx86>\regedit.exe <YaraTags:reg_interpreter>*	Windows	Enabled
Visual Basic	.vbs, .vb, .vbe, .wsf, .wsh	<System>\cscript.exe, <Systemx86>\cscript.exe <System>\wscript.exe, <Systemx86>\wscript.exe <YaraTags:vb_interpreter>*	Windows	Enabled
Java	.jar, .class	*\java.exe, *\javaw.exe <YaraTags:java_interpreter>*	Windows	Disabled
Perl (or Perl using Yara)¹	.pl, .pm	*\perl.exe, <YaraTags:perl_interpreter>*	Windows	Disabled
Python (or Python using Yara)²	.py, .pyc, .pyo	*\python.exe, *\pythonw.exe, <YaraTags:python_interpreter>*	Windows	Disabled
PowerShell	.ps1, .psm1	*\powershell.exe, <YaraTags:powershell_interpreter>*	Windows	Disabled
TCL	.tcl	*\wish.exe, *tclsh.exe	Windows	Disabled

Application or Category	Script Extensions	Processes	Platform	Default State
Ruby	.rb	*\ruby.exe	Windows	Disabled
Chrome Extensions	.crx	*\chrome.exe	Windows	Disabled
Mozilla Extensions	.xpi	*\firefox.exe	Windows	Disabled
HTML Application	.hta	*\mshta.exe, <YaraTags:mshta_interpreter>*	Windows	Enabled

Notes

1,2 -- For upgrades from pre-8.1.6 releases, the updated Perl and Python rules that include Yara are named "Perl using Yara" and "Python using Yara". The pre-Yara versions remain in place with their old names. For new 8.1.6 installations, the rules "Perl" and "Python" include Yara and there is no version without Yara.

Windows Script Rules Changes on Upgrade

Beginning with App Control 8.1.6, the script processors in certain pre-configured Windows script rules are identified by Yara in addition to path and name. The processes identified in this way include: cmd.exe, regedit.exe, reg.exe, regedt32.exe, cscript.exe, wscript.exe, java.exe, javaw.exe, mshata.exe, perl.exe, python.exe, and pythonw.exe.

The Script rules that reflect this change are: **Batch, Registry, Visual Basic, Java, Powershell, and HTML Application.**

The Perl and Python rules will differ depending upon whether the server is new or upgraded from a pre-8.1.6 version:

- When a pre-8.1.6 server is upgraded, two new Windows Script rules are added that include the previous extensions but also use Yara to identify the processors. These new rules are: **Perl using Yara** and **Python using Yara**.

The existing **Perl** and **Python** script rules on an upgraded server remain unchanged since they do not incorporate a process in the rule but rather rely on file associations for the extensions pl, pm, py, pyc, pyo, and pyw.

- New (non-upgrade) server installations do not have **Perl using Yara** and **Python using Yara** rules. Because there are no non-Yara versions of these rules on new servers, the rules using Yara are identified as **Perl** and **Python**.

Because they do not support Yara, the Linux and Mac script rules are unchanged.

Script Rules and Other App Control Rules

A script file defined by a Script Rule is also subject to any matching (non-script) Custom Rules, including those with actions that would Ignore writes, Block, Prompt or Report execution or writing, or Allow execution. For example, if a script file matches a Custom Rule with a Write Action of *Ignore*, the file state of the script will be Unapproved, and execution will be blocked at High or Medium Enforcement Levels. Also, if a script file and its processor match a Custom Rule with an Execute Action of Allow, the script will be allowed to execute regardless of its file state.

In addition, script files can be banned or approved by hash.

There is a Rapid Config called Script Processors that allows you to limit the locations in which script processors can be run. See [Chapter 18, “Rapid Configs,”](#) for more details.

Shell Scripts Identified by Content

The Script Rules table includes rules for native Mac and Linux shell script files, and these are enabled by default. Although scripts are generally identified by file extension and processor in an explicit rule, there is an exception for Mac and Linux shell scripts.

Some shell scripts contain special markup in their first line that identifies the default interpreter that should be used to process them. This markup is usually referred to as *hashbang* or *shebang*, and consists of the “pound” or “hash” symbol (#) followed by an exclamation point (!). For example:

```
#!/bin/bash
```

indicates that the /bin/bash interpreter should be used to process this script file.

Because the shebang markup clearly identifies a file as “interesting” to App Control, shell scripts with this markup are identified by content and tracked, regardless of whether there is a script rule for them. In effect, the markup creates an invisible script rule with the file as the script and the shebang markup identifying the processor.

Enforcement of rules on Mac and Linux shell scripts with the shebang pattern depends on how the script is run and whether any matching Custom Script Rule remains in effect:

- **Use the script as the command** – If a script file is run as a command, it will use the processor identified in the shebang, and will be subject to the policy settings that control *executables*, not scripts. An example of this might be:

```
$ ./foo.sh
```

Note that to run the script this way, the script itself must have execute permission in the operating system.

- **Use a defined processor/script combination as the command** – If a script file is run with the processor as the command and the script file as the argument, and if this combination is defined in the shebang or a Custom Script Rule, the action will be subject to the policy settings that control *scripts*. An example of this might be:

```
$ csh ./foo.sh
```

In this case, execution permission is not necessary for the script file.

- **Use an undefined processor/script combination** – If a script file is run with the processor that is not defined in a shebang pattern for the file nor in a Custom Script Rule, the script action is not controlled by the policy settings for scripts, even if the file itself has been identified as an script to track. This includes the case in which a script file includes a shebang pattern but a different processor is used to run it.

Policy Settings for Script Rules

Unlike custom, registry, and memory rules, script rules do not specify an action. They function primarily to include files in a category already subject to tracking and action rules in App Control. Each policy has two Advanced Settings that specify how script files are controlled on computers in that policy:

- **Block unanalyzed scripts and executables:** This setting determines whether scripts and executables not yet analyzed by App Control are blocked (e.g., in cases where initialization has not yet completed on a computer). It also provides a menu and links through which you can change or disable the notifier that appears if such files are blocked.
- **Block unapproved scripts:** This setting determines whether execution of scripts whose file state is Unapproved is blocked on computers with High or Medium Enforcement. It also provides a menu and links through which you can change or disable the notifier that appears if such files are blocked.

Also keep in mind that scripts are sometimes subject to the policy settings for executables instead of scripts. See [“Shell Scripts Identified by Content”](#) for more details.

Related Topics

See [Table 21, “Advanced Setting Behavior,”](#) on page 189 for information on script-specific settings in policies.

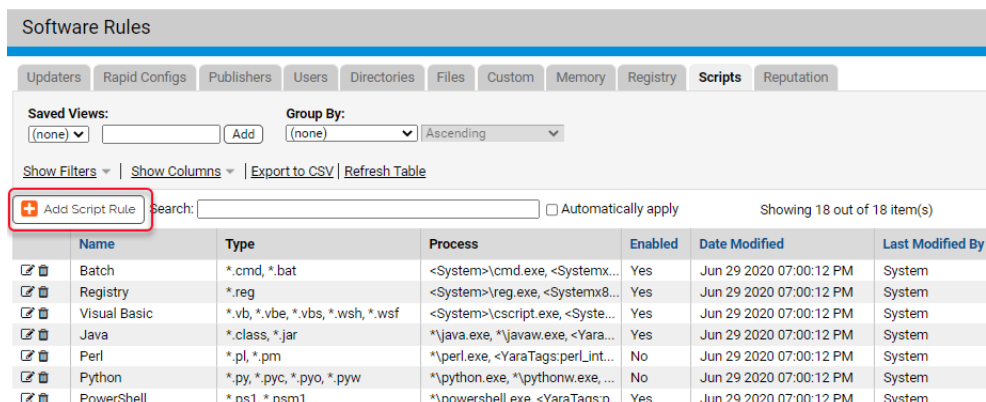
See [Chapter 20, “Endpoint Notifiers and Approval Requests,”](#) for information on configuring notifiers for blocked scripts.

Creating a Custom Script Rule

The procedure below describes how to create a custom script rule. The rule parameters are shown in [Table 50](#).

To add (create) a custom script rule:

1. In the console menu, choose **Rules > Software Rules**, and on the Software Rules page, click the **Scripts** tab. The Custom Script Rules table appears:



Name	Type	Process	Enabled	Date Modified	Last Modified By
Batch	*.cmd, *.bat	<System>\cmd.exe, <Systemx...	Yes	Jun 29 2020 07:00:12 PM	System
Registry	*.reg	<System>\reg.exe, <Systemx8...	Yes	Jun 29 2020 07:00:12 PM	System
Visual Basic	*.vb, *.vbe, *.vbs, *.wsh, *.wsf	<System>\cscript.exe, <Syste...	Yes	Jun 29 2020 07:00:12 PM	System
Java	*.class, *.jar	*\java.exe, *\javaw.exe, <Yara...	Yes	Jun 29 2020 07:00:12 PM	System
Perl	*.pl, *.pm	*\perl.exe, <YaraTags:perl_int...	No	Jun 29 2020 07:00:12 PM	System
Python	*.py, *.pyc, *.pyo, *.pyw	*\python.exe, *\pythonw.exe, ...	No	Jun 29 2020 07:00:12 PM	System
PowerShell	*.ps1, *.psm1	*\powershell.exe, <YaraTags:p...	Yes	Jun 29 2020 07:00:12 PM	System

2. Click the **Add Script Rule** button. The Add Script Rule page appears.

3. In the Rule Name field, enter the name you want to appear on the list of rules. You may also provide a longer, optional Description below the name field.
4. By default, a new script rule is **Enabled** when you configure it and click **Save**. If you want to enable the rule later, click **Disabled** in the Status field.
5. Choose a Platform: **Windows**, **Mac** or **Linux**. All script rules are platform-specific.
6. Choose a Script Definition, which determines how the script processor will be identified. See [Table 50](#) for the choices.
Platform Note: For Mac or Linux scripts, only **Script Type and Process** is allowed in this field.
7. For all Script Rules, enter one or more Script Types. A Script Type is the file name definition for this script type, usually the asterisk followed by a dot and the file extension. To add more script types, after entering the first pattern in the Script Type field, clicking the down arrow to the right of the field. This displays the **Add** and **Remove** buttons, which allow you to add to (or remove from) the list of patterns shown.
8. For Script Type and Process rules (Windows only), you must also add one or more Script Processes. To enter more than one process, use the **Add** button to the right of the field as described in the previous step.
9. If you want to make sure all existing scripts matching this definition are added to the list of files tracked and controlled by App Control, check **Rescan Computers** box.
10. Click the **Save** button to save the rule. The rule now appears on the Script Rules page.

Table 50: Script Rule Parameters

Field	Description
Rule Name (Name in the table)	Name by which this rule is listed in the Script Rules table. (Required)
Description	Additional information about the rule. This can be any text you choose to enter. (Optional)
Status	Radio buttons that make this rule Enabled or Disabled. This allows you to create a rule that you use only at certain times, or to temporarily disable a rule without losing its definition.
Platform	Platform (Windows, Mac, or Linux) for which this script rule is defined. Each script rule must be specific to one platform.
Script Definition (Add/Edit page only)	<p>How you want to define the script rule. The menu choices are:</p> <p>File Association – Choose this definition to allow the file association list on the agent computer to determine the Script Process. You still must provide the Script Type (file name).</p> <p>File Association might be a good choice for a common script type if your environment includes computers with different versions of the script engine for that type (for example, different versions of Perl). However, it is not necessarily the best choice when individual computers have multiple versions of the script engine; only the one identified in the File Association will be managed by App Control. Consider your environment before making this choice.</p> <p>Platform Note: Only Windows scripts can use File Association.</p> <p>Script Type and Process – Choose this definition if you want to specify both the file patterns that define the script and the process that runs the script.</p>
Script Type (Type in the table)	The file name pattern that determines whether a file matches this rule and is therefore considered a script. In most of the standard rules, the script type is defined by the file extension you want identified as a script (for example, *.vbs). You can use paths, wildcards, and macros in the script type. See “Specifying Paths and Processes” on page 407 for a general description of pattern definitions options in App Control rules.
Script Process (Process in the table)	The executable whose behavior you want to control when it processes files matching the Script Type. Examples of script processors include wscript.exe (Visual Basic scripts), cmd.exe (batch scripts), ps.exe (PowerShell scripts) as well as processes that are not obviously script processors such as firefox.exe, chrome.exe and word.exe. You can use paths, wildcards, and macros in the script process. See “Specifying Paths and Processes” on page 407 for a general description of pattern definition options in App Control rules

Field	Description
Rescan Computers	<p>If checked, rescans all connected computers running the App Control Agent to discover any files matching the script rule. If a matching file is found, it is added to the File Catalog with a file state of Approved. If not checked, all script files matching the rule are Unapproved. If a computer is disconnected, it will get the “rescan” rule once it reconnects, and will be re-scanned. Keep the following in mind:</p> <ul style="list-style-type: none"> • Enabling Rescan Computers in a new or existing rule causes a delay during which existing local scripts might not be approved. • If a script file matches a custom rule that instructs the App Control Agent to ignore rules, it will continue to be ignored.
History	<p>For existing rules, a History panel on the Edit Rule page appears showing some or all of the following fields. In addition, these fields can be added as columns on the rules table page.</p> <ul style="list-style-type: none"> • Created By – If the rule was created on this server, the user who created it. Rules created during server installation or upgrades show “System” in this field. • Date Created – If the rule was created on this server, when it was created. • Last Modified By – If the rule has been modified since creation or import, the user who modified it. • Date Modified --If the rule has been modified since creation or import, when it was modified. • CL Version – Rules created after server installation also show the CL (config list) number that first contained the rule so that you can compare an agent CL number to determine whether the agent has received the rule.

Important

Use of very broad definitions for either the Script Type or Script Process field is not recommended because of negative performance impact. If either field in a rules uses * or *.* , a warning will be displayed on the page. Be as specific as possible in defining the file patterns in a Script Rule.

Editing a Script Rule

You might choose to edit a script rule for a variety of reasons, including:

- enabling or disabling the script (see [“Disabling or Deleting a Script Rule”](#) on page 388 for more on the effects of enabling or disabling a script)
- adding, removing, or modifying patterns used to identify the script, or its processor
- changing the Script Definition to use File Association to identify the Script Processor, or to change from File Association to a specified processor pattern or patterns.

To edit a script rule:

1. In the console menu, choose **Rules > Software Rules**. The Software Rules page appears.
2. On the Software Rules page, click the **Scripts** tab. The Custom Script Rules table appears.
3. Click on the View Details button to the left of the rule you want to edit. The Edit Custom Script Rule page appears.
4. Edit the rule as you choose (see [Table 50](#) for a description of the parameters) and then click **Save**. The Edit Custom Rule page closes and the Custom Script Rules page is displayed.

Disabling or Deleting a Script Rule

If you do not want a script rule to be effective anymore, you can either disable it, which leaves it in the table of script rules, or delete it from the table. In either case, the script rule stops affecting newly discovered files. However, any script file that was discovered while the rule was effective continues to be tracked by App Control and retains any file state assigned to it during the time the rule was enabled.

Disabling a script definition does not immediately remove the matching files from the inventory of files tracked by App Control. This prevents loss of information if an action such as a rule change is taken accidentally. However, the exact amount of time a script file matching a disabled rule remains in inventory depends factors such as whether it is actually deleted from the agent or modified.

If a disabled definition is subsequently enabled with rescan enabled, only newly discovered scripts will be locally approved. Scripts that remained in the inventory will retain their previous state.

If you think you might use a rule again, disabling it temporarily is the best choice.

To disable a script rule:

1. In the console menu, choose **Rules > Software Rules**, and when the Software Rules page appears, click the **Scripts** tab. The Custom Script Rules table appears.
2. Click the View Details button next to the rule you want to disable. The Edit Script Rule page appears.
3. In the Status line, click the **Disabled** radio button, and then click the **Save** button at the bottom of the page. The rule is now disabled.

Deleting a rule eliminates it permanently – there is no undo or retrieval for a deleted rule. Because of that, be sure you actually want to delete the rule. Deletion of the rules that were pre-configured in App Control is not recommended.

To delete a script rule:

1. In the console menu, choose **Rules > Software Rules**, and when the Software Rules page appears, click the **Scripts** tab. The Custom Script Rules table appears.
2. Click the Delete button (red circle with X) next to the rule you want to delete, and click **OK** on the confirmation dialog. The rule is now deleted.

Viewing Rule Status on Computers

Depending upon the number of agents managed by your App Control Server and the number that are disconnected, not all agents might receive new or updated rules in a short amount of time. The Related Views menu on the Edit page for an enabled rule provides links to two different filtered views of the Computers page to help determine the status of the rule on agent-managed computers. The choices are:

- **All Computers that have received this rule**
- **All Computers that have not yet received this rule**

This menu does not appear for rules that have never been enabled.

Script Rule Examples

App Control includes several pre-configured Script Rules. These are useful as examples for creation of other rules.

Example: Windows Batch Scripts

App Control includes a script rule to identify and control executions of Windows batch scripts. On the Scripts tab of the Software Rules page, you can click on the View Details button next to the Batch rule to see how it is defined.

Edit Script Rule

General

Rule Name:

Description:

Status: Enabled Disabled

Definition

Platform:

Script Definition:

Script Type:

*.cmd
*.bat

Script Process:

<System>\cmd.exe
<Systemx86>\cmd.exe
<YaraTags:cmd_interpreter>*

Rescan Computers:

History

Created By: System
Date Created: Jun 29 2020 07:00:12 PM
Last Modified By: System
Date Modified: Jun 29 2020 07:00:12 PM
CL Version: 161

The *Script Type* field for the Batch rule includes two patterns – *.cmd and *.bat. Any file ending in either of these extensions will be identified as a batch script file, and will be tracked by App Control once discovered.

The *Script Definition* field shows *Script Type and Process*, so it is necessary to provide at least one pattern to match for the Script Process. In this case, there are two processes listed so that *cmd.exe* is identified as the processor for this script for both 32-bit and 64-bit systems.

When this rule is enabled, any time the *cmd.exe* (in the locations shown) attempts to access a file with a .cmd or .bat extension, the agent will control execution based on the current approval state of the script file, the policy settings for the computer on which the execution attempt occurs, and any other rules affecting the files.

Because *Rescan Computers* is checked in this rule, as soon as the rule is enabled, all computers managed by this App Control Server will be rescanned, and any files matching the Script Type for the rule will be locally approved and added to the File Catalog and Files on Computers list.

Example: Linux Shell Scripts

The App Control Server includes a script rule to identify and control executions of native shell scripts on Linux computers. On the Scripts tab of the Software Rules page, you can click on the View Details button next to the Linux Shell rule to see how it is defined.

Edit Script Rule

General

Rule Name:

Description:

Status: Enabled Disabled

Definition

Platform:

Script Definition:

Script Type:

- *.sh
- *.csh
- *.zsh
- *.ksh

Script Process:

- /bin/bash
- /bin/csh
- /bin/ksh
- /bin/sh

Rescan Computers:

History

Created By: System

Date Created: Jun 29 2020 07:00:12 PM

Last Modified By: System

Date Modified: Jun 29 2020 07:00:12 PM

CL Version: 153

The *Script Type* field for the Linux Shell rule includes several patterns – *.sh, *.csh, *.zsh, *.ksh. Any file ending in one of these extensions will be identified as shell script file, and will be tracked by the App Control Server once discovered.

The *Script Definition* field shows *Script Type and Process*, which is the only choice usable for Mac and Linux rules. There is a long list of processes in the rule, which support native script processing on the supported Linux platforms. If you choose you can add or remove processors (or script types) for this rule.

When this rule is enabled, any time a listed processor, such as /bin/bash, attempts to access a file with a listed extension, such as .sh, the App Control Server will control execution based on the current approval state of the script file, the policy settings for the computer on which the execution attempt occurs, and any other rules affecting the files.

Because *Rescan Computers* is checked in this rule, as soon as the rule is enabled, all computers managed by this App Control Server will be rescanned, and any files matching the *Script Type* for the rule will be locally approved and added to the File Catalog and Files on Computers list.

Chapter 14

Custom Software Rules

Custom Rules define actions you want the agent to take in response to file, directory, or process activity that matches conditions you specify. They may be used to optimize performance, protect file integrity, create a trusted file path for software distribution, or meet other special needs.

Unified Management

Custom Rules can be centrally managed for multiple servers through the Unified Management feature. This is described in [“Unified Management of Rules”](#) on page 791.

Sections

Topic	Page
Overview	393
Creating a Custom Rule	396
Custom Rule Fields	401
Specifying Paths and Processes	407
Using Macros in Rules	410
Rule Ranking	424
Rule Ranking and Internal Rules	427
Enabling and Disabling Custom Rules	428
Exporting and Importing Rules	431
Custom Rule Types and Examples	439

Overview

Custom Rules allow you to customize the actions you want the agent to take in response to specified file, directory, or process activity. They may be used to take actions such as:

- blocking or allowing file modifications and executions
- controlling which files get tracked
- determining whether or not tracked files are approved
- reporting events when specified activity is seen
- creating exceptions to other types of rules, such as approvals or bans

Notes

Standard methods for approving and banning files are described in [Chapter 8, “Approving and Banning Software.”](#)

Rule Types

App Control provides several partially configured custom rule types for the following specific purposes:

- **File Integrity Control** – Prevents or reports changes to specified folders or files.
- **Trusted Path** – Defines folders or files for which file execution is always allowed.
- **Execution Control** – Creates a rule to control behavior when an attempt is made to *execute* a file matching the rule.
- **File Creation Control** – Creates a rule to control behavior when an attempt is made to *write* a file matching the rule.
- **Performance Optimization** – Specifies folders or files for which file creation, modification, and deletion are ignored (execution will still be monitored).

Two additional rule types allow more detailed configuration of a rule:

- **Advanced** – Provides a menu-based user interface in which you choose all fields yourself. The menu choices are standard actions (such as Write) that might actually involve combinations of actions or permissions internally.
- **Expert** – Provides a checkbox-based interface in which you can select one or more of the internal actions underlying the other rules types. These rules are described in [Chapter 17, “Expert Rules.”](#)

Important

Users enabling Advanced and Expert Rules should be certain they have a thorough understanding of rule behavior to achieve the desired goal and avoid unwanted outcomes. You may want to consult with Carbon Black Support or Services representatives before enabling one of these rules.

Custom rules can be used to enable network login scripts or software deployment systems, or to designate an area for software developers to run executables without the App Control Server tracking file activity or enforcing rules. You also can use a custom rule to prevent users from uninstalling an application by blocking any changes to that application's directory.

Rule Scope

You can create custom rules that apply on all computers on a platform (e.g., all Windows computers) under all conditions. If you have Unified Management enabled, you can even have a rule apply to agents on more than one App Control server. On the other hand, you can focus the scope of a rule by specifying one or more of the following criteria (not all of these options are available for all rule types):

- **Process-specific** – You can choose to make a rule effective only when certain *processes* attempt to write or execute files in the specified location.
- **User- or group-specific** – You can make the rule apply only to a specific *user or group of users*.
- **Policy-specific** – You can choose to limit a rule to *computers in specified policies*.
- **Server-specific** – If you have Unified Management enabled, you can choose to limit a rule to *computers reporting to specified servers* in the management group. This is described in [“Unified Management of Rules”](#) on page 791.
- **Rule ranking** – Custom rules are evaluated in order of *Rank*, a column that is displayed by default on the Custom Rules table. The rule ranked ‘1’ has the highest rank, ‘2’ is next, and so on. If a rule blocks, allows, or prompts the user to make block or allow, that rule stops processing of other rules, so rank is important in these cases. You can change the order of rules, for example, putting a rule applying to *one specific file in a folder* higher on the list, while putting another rule for *all the files in the same folder* lower – because the first rule is higher, it takes precedence.
- **Conditional Macros** – You can use certain macros to restrict the conditions under which specific parameters in rules are applied. Only agents meeting the “test” described in the macro will attempt to match the parameter prefixed with the macro. Most of these macros are *OnlyIf* macros with different arguments, such as `<OnlyIf:OSVersionIs:10.6.8>` and `<OnlyIf:HostName:*SMITH-1*>`.

All user-created custom rules are platform-specific; they apply to only one of the platforms – Windows, Mac, or Linux – that App Control Agents can be installed on.

File and Process Matching

To determine whether a file or process attempting an action matches a custom rule, a string comparison is done between the file or process name and the specifications in the rule. Hash values are not used for custom rule processing.

You can include *wildcards* and special *macros* in a path or process specification to broaden the rule scope or allow the rule to match files or processes in locations that vary from one agent computer to another. See [“Specifying Paths and Processes”](#) on page 407 for additional details.

Note

If you are specifying a file extension for a script in a rule that controls execution, the rule will not recognize the script and control its execution unless there is a corresponding Script Rule for the file extension and the process that executes the script. See [Chapter 13, “Script Rules,”](#) for more details.

Pre-configured Rules

A new installation of the App Control Server is pre-configured with several custom rules found to improve performance and/or prevent unnecessary tracking. These rules are enabled by default. You can remove or disable them if you choose. For upgrades from previous releases, these rules are added *below* (i.e., with a lower rank than) rules that already existed.

The table of rules also includes rules labeled **[Sample]**, which are disabled by default. In general, these are application-specific rules that allow files needed for certain common applications or suites to be executed or written. You may enable these, with or without modifications of your own.

Internal Rules in the Custom Rule Table

The Custom Rules table includes rules labeled *Internal*. These are the rules you enable in other parts of the console, mostly in the Device and Advanced Settings on the Edit Policy page. For example, *Block banned file hashes*, which is on the Advanced Settings table for a policy, is listed as an Internal rule on the Custom Rules page.

An internal rule shows its status as Enabled in the rules table if it is enabled in *any* policy. You cannot enable, disable, modify or move Internal Rules in the Custom Rules table, but you can move other, non-internal Custom Rules, relative to the Internal Rules to better control how and when different rules are enforced. See [“Rule Ranking and Internal Rules”](#) on page 427 for more details.

Internal Custom Rules apply to all platforms.

Specifying the Notifier for a Custom Rule

App Control provides *notifiers* that can be displayed when a rule blocks an action or prompts the user for a decision to allow or block an action. For each custom rule, you can choose from two sources for the notifier:

- **Use Policy Specific Notifier** – Each Policy includes an Advanced Setting, “Enable custom (file and path) rules”, which is always on. This setting has a Notifier field in which you can choose the notifier that appears on agent computers when custom rules block an action.

If you choose Use Policy Specific Notifier for a rule, it is possible that the policy specifies <none> as the Notifier for Enforce custom (file and path) rules. In this case, a notifier will not be shown, even for a Prompt rule. Unless you are certain that you never want to prompt the user for a response to a rule, choosing <none> for the custom rule notifier in a policy is not recommended. See [“Advanced Settings”](#) on page 187 for more information.

- **Custom Notifier** – If you do not choose the policy-specific notifier, you can choose (or create) a notifier specifically for a custom rule. The choices appear on a menu on the Add/Edit Custom Rule page.

When you choose Block as the rule action, you can choose <none> on the Custom Write Notifier menu since it is possible you want the rule to block actions without notification. A Prompt rule requires a user choice, so when you choose Prompt as the rule action, the Custom Notifier menu does not include <none>.

See [Table 51](#) below for the custom rule notifier settings. See [Chapter 20, “Endpoint Notifiers and Approval Requests,”](#) for more on notifiers.

Unified Management

If you are using Unified Management and create a Custom Rule that applies to more than one server, client servers will use default notifiers, even if a custom notifier is specified on the management server.

Custom Rules in Visibility Mode

In Visibility mode policies, the effect of custom rules depends on the type of rule:

- Custom rules that would block a file have no effect in Visibility mode, but they still generate App Control events.
- Custom rules configured to prompt the user for a response in do not interrupt the action, but a "would have prompted" event is generated.
- Custom rules that approve a file *do* change the file state, but in Visibility mode this has no effect on file execution.
- Custom rules that specify "Ignore" on the Write menu (see below) *are* effective in Visibility mode.

Creating a Custom Rule


To create a custom rule from scratch, you would need to provide the information shown in bold in the left column:

General Description	Field on Add/Edit Custom Rule Page
If this/these source process(es) ...	Process
...and/or this/these user(s) ...	User or Group
... attempts to perform this/these operation(s) ...	Operation (Execute, Write or Both)*
... on this/these file(s) ...	Path or File
... on computers in this/these policy(ies) ...	Rule applies to/Policies
... on computers reporting to this/these App Control server(s) ...	Rule applies to/Servers (if Unified Management is enabled)
... on computers running on this platform ...	Platform
... then this/these action (s) should be taken.	Execute Action and/or Write Action*

* Additional operations and actions are available in Expert Rules.

One rule can match one or more processes, users, paths, files, policies and servers. It is always specific to a single platform, however. Also, instead of the descriptions above, you could make the rule function when any process *except* the ones you specify attempts the action, or an action is attempted on any file *except* the ones you specify.

On the Add Custom Rule page, your choice of Rule Type modifies the displayed fields:

- Some *fields* are eliminated from the page if they are not relevant (or have only one sensible value) for the rule type you choose.
- Some *menu choices* are eliminated so that only choices relevant to the rule type are available.
- *Inline Help buttons*  open text boxes with assistance in choosing values appropriate to the rule type for many configurable fields.

The following procedure describes the process of creating a custom rule on one App Control Server.

Unified Management

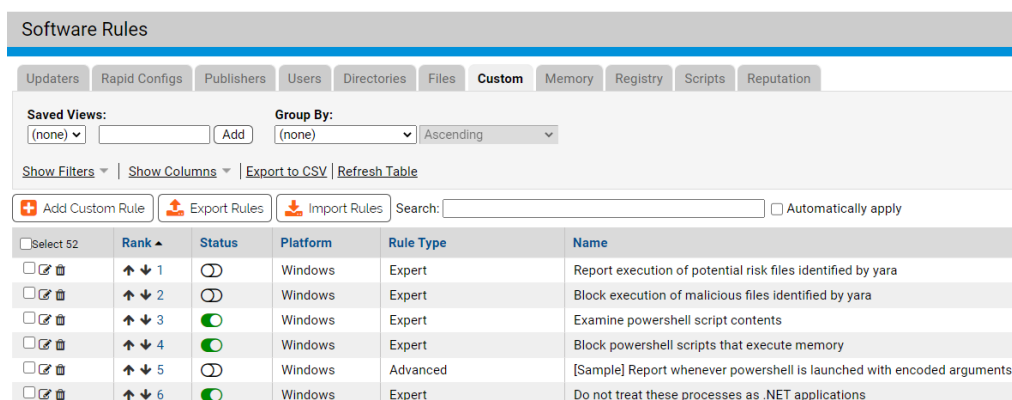
If you are using Unified Management, you can also:

- apply a new rule to multiple servers when you create it
- copy one or more existing rules from the management server to one or more client servers

See [“Unified Management of Rules”](#) on page 791 for more details.

To add (create) a custom rule:

1. In the console menu, choose **Rules > Software Rules**. The Software Rules page appears.
2. On the Software Rules page, click the **Custom** tab. The Custom Rules table appears:



	Rank	Status	Platform	Rule Type	Name
<input type="checkbox"/>	1		Windows	Expert	Report execution of potential risk files identified by yara
<input type="checkbox"/>	2		Windows	Expert	Block execution of malicious files identified by yara
<input type="checkbox"/>	3		Windows	Expert	Examine powershell script contents
<input type="checkbox"/>	4		Windows	Expert	Block powershell scripts that execute memory
<input type="checkbox"/>	5		Windows	Advanced	[Sample] Report whenever powershell is launched with encoded arguments
<input type="checkbox"/>	6		Windows	Expert	Do not treat these processes as .NET applications

3. Click the **Add Custom Rule** button. The Add Custom Rule page appears.

Add Custom Rule

General

Rule Name:

Description:

Status: Enabled Disabled

Definition

Platform:

Rule Type:

Write Action: Use Policy Specific Notifier

Path or File:

Process Exclusion:

Rule Applies To

Policies: All Current and Future policies
 Selected policies

4. In the Rule Name field, enter the name with which you want to identify this rule.
5. If you want to add other comments about the rule, such as its purpose or its relationship to other rules, you may provide an optional Description.
6. By default, a new custom rule is **Disabled** as soon as you define it and click **Save**. If you want the rule to take effect immediately, click **Enabled** in the Status field.
7. Choose the Platform you want this rule applied to (Windows, Mac, or Linux). Each rule applies to one platform only.
8. Choose the Rule Type from the menu. File Integrity Control is the default choice. Specific rule types are partially configured for you. If none of the specific types appears to fit your needs, choose **Advanced** on the Rule Type menu to see a greater number of configuration options. [Table 51](#) describes the different rule types as well as all of the other custom rule fields.

Note: You also can choose **Expert** as the Rule Type. Expert rules require much more detailed configuration and have a slightly different user interface than other rules. They are intended for use under special circumstances as directed by Carbon Black Support or Services representatives. See [Chapter 17, "Expert Rules,"](#) for more details.

- Enter the remaining fields you want for this custom rule (see [Table 51](#)) and then click the **Save** button if you need to remain on the page or **Save & Exit** if to go to the Custom Rules table. By default, a new Custom Rule is **Disabled** and ranked #1, listed at the *top* of the Registry Rules table.

Unified Management

If Servers and Override Permissions fields appear on the page, you are on a Unified Management server, and have the option of applying this rule to multiple App Control Servers. See [“Unified Management of Rules”](#) on page 791 for more details.

- Before you enable a rule, change its rank unless you want it to take precedence over (and perhaps preempt) all other rules. You can change rank using the arrows in the Rank column or drag-and-drop (if the table is sorted by rank), or you can click on the rank number and enter a new rank in the dialog box. See [“Rule Ranking and Internal Rules”](#) on page 427 for more details.
- When you are satisfied with the rank and want to enable the rule, click the toggle switch in the Status column of the Registry Rules table. The button in the switch moves to the right and the background turns from white to green.

Editing a Custom Rule

Editing a Custom Rule is very similar to creating one. If you have permission to edit the rule, you can edit any field, including the rule name.

To edit a custom rule:

- In the console menu, choose **Rules > Software Rules**. The Software Rules page appears.
- Click **Custom** on the left menu or the tab on the page. The Custom Rules table appears.
- On the Software Rules page, click the View Details button for the rule you want to edit.

<input type="checkbox"/> Select 52	Rank ▲	Status	Platform	Rule Type	Name
<input type="checkbox"/>	↑ ↓ 1	<input type="radio"/>	Windows	Expert	Report execution of potential risk files identified by ya
<input checked="" type="checkbox"/>	↑ ↓ 2	<input type="radio"/>	Windows	Expert	Block execution of malicious files identified by yara
<input type="checkbox"/>	↑ ↓ 3	<input checked="" type="checkbox"/>	Windows	Expert	Examine powershell script contents

- On the Edit Custom Rule page for that rule, make your changes. [Table 51](#) describes the custom rule fields that you might choose to edit.
- Click either **Save** (to remain on the Edit Custom Rule page) or **Save and Exit** (to return to the Custom Rules table page).

6. If you clicked Save, a confirmation message appears on the page. Click the message to clear it from the page. If an error occurs, review the error message and correct the conditions that caused the error before saving again.

Unified Management

If you are using Unified Management, and you edit a unified rule shared with other servers, a “wizard” shows the progress as the edited rule is saved on each server.

See [“Unified Management of Rules”](#) on page 791 for more details.

Copying a Custom Rule

There is a *Copy this rule* command on the right menu on the Edit Rules page for Custom, Registry, and Memory Rules. This is for making copies of the rule on the same server. You might do this so that you can customize a sample rule while preserving the original settings as a template. It also allows you to make slightly different rules for different policies without having to manually provide all of the settings for each one.

To copy a custom rule:

1. On the Custom Rules table page, click the View Details button to open the details page for the rule you want to copy.
2. On the details page, click **Copy this rule...** in the Actions menu on the right side of the page. This opens a dialog box. By default, the copy is named using the original rule name plus “(copy)”.
3. In the dialog box, change the rule name if you want something more descriptive that what is there.
4. If you want the new rule enabled immediately, check the **Enable copied rule** box. Otherwise, uncheck the box.
5. Click **OK**. The copied rule is created and its details page replaces the details page for the original rule.
6. Make any changes you would like to make in the new (copied) rule and **Save** or **Save & Exit**.

Custom Rule Fields

Table 51 shows the fields available on the standard Add/Edit Custom Rule page. Column headings on the rule table page are shown when they differ from the Add/Edit page. See “Expert Rule Definitions” on page 492 for the “Expert” rule type fields.

Table 51: Custom Rule Fields

Field	Description
Rule Name (Name in table)	Name by which this rule is identified. (Required)
Description	Additional information about the custom rule. This can be any text you choose to enter. (Optional)
Rank (Table only)	The rank of this rule in order of evaluation. The rule ranked ‘1’ in the table is evaluated before the rule ranked ‘2’, and so on.
Status	Radio buttons that make this rule Enabled or Disabled. This allows you to create a rule that you use only at certain times, or to temporarily disable a rule without losing its definition.
Platform	Platform (Windows, Mac, or Linux) for which this rule is effective. Except for built-in “internal” rules, each custom rule is specific to a single platform.
Rule Type	The Rule Type choice changes other options and defaults on the Add/Edit Rule page to partially pre-configure rules for certain common scenarios. Options are File Integrity Control , Trusted Path , Execution Control , File Creation Control , Performance Optimization , Advanced and Expert . See “Custom Rule Types and Examples” on page 439 for descriptions and examples.
Execute Action (Add/Edit page only)	The action to take when there is a file execution attempt matching this rule. The menu appears when the Operation choice is <i>Execute</i> or <i>Execute and Write</i> . See Table 52 for options.
Write Action (Add/Edit page only)	The action to take when there is an attempt to create, modify or delete a file matching this rule. The menu appears when Operation choice is <i>Write</i> or <i>Execute and Write</i> . See Table 53 for options.
Action (Rule table only)	The type of action the rule takes. The possible values include all of those shown for Execute Action and Write Action plus other actions made available in Advanced and Expert rules.
Operation	The type of operation the rule affects. Menu choices of Execute , Write , or Execute and Write appear for this field on the Create/Edit Rule page for an Advanced rule. Other operations are available for Expert rules.

Field	Description
Action (Legacy) (Rule table only)	This column shows actions and operations for the rule as shown in the Action column in pre-8.1.6 versions, or it shows “Expert Action(s)” in cases where expert rule information was not previously shown. Note: This field is present strictly for continuity with older versions – you should use the separate Action and Operation columns for the most accurate description of the rule.
Send Approval Event	For Advanced rule types that specify <i>Approve</i> or <i>Approve as Installer</i> , when this box is checked (the default), an event is recorded when a file is approved because of the rule.
Use Policy Specific Notifier	If you choose Block or Prompt as the Action, this checkbox appears to the right of the Action choice and is checked by default. If the box is checked, when a custom rule blocks an action, the notifier that appears is the one specified for the Enable Custom (file and path) Rules setting in the policy for the computer on which the action was blocked. If not checked, you can choose a custom notifier from the Custom Notifier menu.
Custom Execute/Write Notifier	If you choose Block or Prompt as the Action, and check the Use Policy Specific Notifier box, this menu appears. When Block is the Action, you can choose any notifier from the menu. The menu also includes a <none> option so that you can disable the notifier for this rule. When Prompt is the Action, you can choose any notifier on the menu. However, Prompt rules <i>must</i> display a notifier, so there is no <none> choice in this case. Note: If you use Unified Management to create a rule that applies to more than one server, client servers will use default notifiers, even if a custom notifier is specified on the management server.
Path or File (Path in table)	Path to which this rule applies. This can be a folder or a specific file. You can use a local path or a UNC path, but not mapped drives (for example, Z:\application). See “Specifying Paths and Processes” on page 407 for details on specifying a path.
Process	This field allows you to limit the rule so that it is applied only when certain processes attempt to execute or write files matching the path specification. See “Specifying Paths and Processes” on page 407 for details and Table 58 for process menu options.
Process Exclusion (Add/Edit page only)	This field allows you to specify one or more processes for which a File Integrity Control rule will <i>not</i> be applied. See “Specifying Paths and Processes” on page 407 for details.
User or Group	The users or groups to which this rule applies. See “Specifying Users or Groups” on page 424 for details.

Field	Description
Rule Applies To: Servers (Add/Edit page only)	<p>Radio buttons allow you to apply the rule to the current server, All Servers or Selected Servers. If you choose Selected Servers, a checklist that includes the current server and of all App Control servers managed by this server appears. In addition, policies for the servers you include appear in the Selected policies list.</p> <p>Unified Management: This field appears only if Unified Management is configured on the server you are logged into.</p>
Unified Server Source (Table only)	<p>If this is a unified rule, the name of the unified management server that created or copied the rule.</p>
Rule Applies To: Policies (Policy in the table)	<p>Radio buttons allow you to apply the rule to All Current and Future Policies or Selected policies. If you choose Selected policies, a checklist of all policies on your App Control Server appears.</p> <p>Unified Management: If Unified Management is configured on the server you are logged into, and if you applied the rule to additional servers, policies for all selected servers appear in this list.</p>
Is Global (Table only)	<p>Indicates whether the rule applies to all policies (Yes) or only selected policies (No).</p>
Rule Applies To: Override Permissions (Add/Edit page only)	<p>Radio buttons allow you to specify whether administrators on other servers can modify rules sent via Unified Management on their own server. The options are No Override, Partial Override (allows changing rank) and Full Override (allows editing and changing rank).</p> <p>Unified Management: This field appears only if Unified Management is configured on the server you are logged into and this rule is applied to more than the current server in the Rule Applies To:Servers field.</p>
History	<p>For existing rules, a History panel on the Edit Rule page appears showing some or all of the following fields. In addition, these fields can be added as columns on the rules table page.</p> <ul style="list-style-type: none"> • Created By – If the rule was created on this server, the user who created it. Rules created during server installation or upgrades show “System” in this field. • Date Created – If the rule was created on this server, when it was created. • Last Modified By – If the rule has been modified since creation or import, the user who modified it. • Date Modified --If the rule has been modified since creation or import, when it was modified. • CL Version – Rules created after server installation also show the CL (config list) number that first contained the rule so that you can compare an agent CL number to determine whether the agent has received the rule. • Imported – (In the table only) indicates whether the rule was imported (Yes/No). • Imported By – If the rule was imported to this server, the user who imported it. • Imported Date – If the rule was imported to this server, when it was imported.

Specifying Execute and Write Actions

You can control two types of action with a custom rule: Execute Action and Write Action.

Execute Action is the action you want to take when there is a file execution attempt matching a rule. The Execute Action menu appears when the Operation choice is *Execute* or *Execute and Write*. [Table 52](#) shows the choices.

Table 52: Execute Action Choices

Menu Choice	Description
Default	Apply existing policy settings and other non-custom rules to file execution attempts matching this rule, <i>and do not process other custom rules</i> .
Allow	Allow a file matching the rule to execute in the specified path, even if execution would otherwise be blocked. Note: The promotion state (whether the file is treated as an installer) depends on the process attempting the action (e.g., if that process is promoted, the newly created process will also be promoted).
Block	Prevent a file matching the rule from executing. When Block is chosen, the Use Policy Specific Notifier checkbox appears and is checked by default. You also can uncheck this box to choose a Custom Notifier to alert the user when the rule blocks an action. See Table 51 for more details.
Promote	Promote (treat as an installer) a file matching this rule. Even if a file is promoted, whether it can <i>run</i> or not depends on its existing file state and the Enforcement Level of the machine on which the execution is attempted. If the file is allowed to run, any files written by it will be locally approved unless already banned, and the written files also will be promoted if the process that <i>wrote</i> them attempts to <i>execute</i> them.
Allow and Promote	Allow a file matching the Path or File specification to execute regardless of its state, and promote it (treat it as an installer). Files written by a file matching an Allow and Promote rule will be locally approved unless already banned. See the section " Trusted Paths " for more on choosing to trust execution of files by path name.

Menu Choice	Description
Prompt	<p>Display a notifier dialog to users when an attempt is made to execute a file matching this rule.</p> <p>When Prompt is chosen, the Use Policy Specific Notifier checkbox appears and is checked by default. You also can uncheck this box to choose a Custom Notifier to alert the user when the rule blocks an action. See Table 51 for more details.</p> <p>The user can Block execution, Allow execution (and locally approve the file if allowed), or Promote (and allow execution of) the file. The behavior for the choice the user makes is the same as the behavior if the rule itself specified Block, Allow, or Allow and Promote. If the user chooses Allow or Promote, subsequent actions that are identical to the one Allowed or Promoted are completed without prompting.</p> <p>Note: Blocking or allowing execution from a Custom Rule prompt does not change the global approval or ban state.</p>
Report	Report (as an event) and allow execution of a file matching this rule, regardless of file state.
Report Process Create	Report (as an event) creation of a process matching the file and path specified by this rule by the process specified by the rule.
Block Silently	Prevent execution of a file when the execution conditions match this rule. Do not display a notifier, and do not generate a App Control event.
Report Process Exit	Report (as an event) the exit of a process matching the file and path specified by this rule that was started by the process specified in the rule.
Report Image Load	Report (as an event) loading of a DLL or EXE matching the file and path specified by this rule when loaded by the process specified in the rule.

Write Action is the action to take when there is an attempt to create, modify or delete a file matching a rule. The Write Action menu appears on the Add/Edit Custom Rule page when Operation choice is *Write* or *Execute and Write*. [Table 53](#) shows the choices.

Table 53: Write Action Choices

Menu Choice	Description
Silence	For an action that matches this rule and one or more additional rules (built-in or user-created), prevent notifications, meters, and events without preventing enforcement of the other matching rule(s) For example, if another rule would ban or block an action, the ban or block is still effective. If an action matching a Silence rule would have displayed a prompt (allow or block) notifier, the action will be blocked. Available for Advanced and Expert rule types only.
Default	Apply existing policy settings and non-custom rules when an attempt is made to write a file matching this rule. <i>Do not process any other Custom Rules for matching files.</i>
Ignore	Do not track creation, modification or deletion of a file matching this rule. Although not tracked, files matching an ignore rule are still blocked from writing if the file state and Enforcement Level would normally enforce a block. Ignore does not stop rule processing. If a write attempt matches both an Ignore rule and another rule lower in rank, the second rule is processed.
Track	Track creation, modification or deletion of a file matching this rule. This action allows creation of exceptions to Ignore rules. Appears only for Advanced and Expert rule types.
Block	Prevent writing of a file matching this rule. This prevents file creations, file deletions and file modifications. When Block is chosen, the Use Policy Specific Notifier checkbox appears and is checked by default. You also can uncheck this box to choose a Custom Notifier to alert the user when the rule blocks an action. See Table 51 for more details.
Approve	Allow a file matching this rule to be created (written) and locally approve it <i>if possible</i> (if it is not banned globally or by policy).
Approve as Installer	Allow a file matching this rule to be created (written) in the named directory, and locally approve and mark it as an installer <i>if possible</i> (i.e., if it is not banned globally or by policy). Note: “Approve as installer” by a custom rule is a local and transient action only. It has no impact on any other instance of the file, and is not effective on this instance if the file is globally flagged as “Not an installer” because the initial state was overridden. The rule <i>is</i> effective if a file is marked as “Not an installer” because of the initial App Control analysis of the file. Use this option with caution since it allows a file to be identified by <i>name</i> as an installer without confirming the file hash.

Menu Choice	Description
Prompt	<p>Present users who attempt to write a file matching the rule with a notifier dialog letting them block or allow writing.</p> <p>When Prompt is chosen, the Use Policy Specific Notifier checkbox appears and is checked by default. You also can uncheck this box to choose a Custom Notifier to alert the user when the rule blocks an action. See Table 51 for more details</p> <p>If the user selects Approve on the notifier, the file is written, and if it is an executable, it is approved. Subsequent identical operations (i.e., the same file and path, not a different matching file) are approved without prompting. Note, however, that global bans by name or hash still control whether the file can be executed.</p>
Allow	<p>Allow a file matching this rule to be created, modified, or deleted. This choice has no effect on the state of the file being written.</p>
Report	<p>Report (as an event) writing of <i>any</i> file matching this rule, even if the file is not normally tracked by the App Control Server. This includes files not analyzed as executable and files that are not the first seen instance of a hash.</p>
Never Report	<p>Never report actions matching this rule to the server. A record of the action will still be maintained on the agent.</p>

Specifying Paths and Processes

When you specify Path or File in a Custom Rule, you have several options for defining the string for that field. The same string options can be used for either of the two Process options that require entry of a path (*Specific Process...* or *Any Process Except ...*).

Definition

Platform: Windows

Rule Type: Performance Optimization

Path or File: Any Path

Process: Any Process

These options are:

- **Specify a directory or a file/process** – You can enter a path or process specification that exactly identifies a file by path and name so that only that file matches the rule. You also can enter a specification that identifies a directory, and so affects all files or processes in that directory and its subdirectories.
- **Specify a local drive or UNC path (Windows only)** – You can use a local drive name, such as *C:\folder1\subfolder\application.exe*, to identify a local path or process. For a remote path or process, use a UNC path, such as *\\computer\dir\application.exe*. Mapped drives in a path or process specification are not recognized.
- **Use wildcards** – You can use wildcards ('?' for any one character and '*' for zero or more characters) to expand the scope of a path or process specification, or to help

- you match a file or folder whose exact location you don't know. Wildcards may be used at the beginning, end or middle of a path.
- **Use macros** – You can use special macros to identify certain well known folders, even if you don't know their exact location on agent computers. Macros are platform-specific, and in the current release, most apply to Windows only.
 - **Specify multiple paths or processes** – For both paths and processes, you can add more than one path definition per rule.

Specifying a File or Directory

You can enter a directory or a specific file as your path. When you specify a directory, you are instructing the rule to operate on files in that directory and any of its subdirectories (unless there are higher-ranked rules specific to certain files or subdirectories).

To indicate that a Path or File definition or a Process definition is a directory, you must end it with the folder delimiter (slash or backslash) for the rule platform or with the delimiter and an asterisk. If you do not include the delimiter, the rule will attempt to match a *file* by the name you provided, not a directory. For example, either of the following correctly identifies a directory in a Windows path definition:

```
c:\folder1\subfolder2\  
c:\folder1\subfolder2\*
```

However, the following is *not* recognized as a directory:

```
c:\folder1\subfolder2
```

If you use path macros in a path or process definition, the App Control Server automatically processes the macro so that it is treated as a directory, even if you don't follow the macro with a backslash. See [Using Macros in Rules](#).

Platform-Specific Syntax

The path you provide for a rule will be interpreted according to the path rules for the platform you choose for the rule. Specifically:

- The case sensitivity of paths and file name in rules usually depends on the operating system. Rules normally are *not* case sensitive for Mac and Windows. They normally *are* case sensitive for Linux. However, if a file system with different case-sensitivity rules is attached to a system – for example by connecting an external drive or mounting a network file system – the case sensitivity of the file system determines whether a rule is effective.
- Path and file name case are preserved in the form you enter them, even for case insensitive platforms.
- Paths must use the correct directory delimiter for the rule platform: forward slash (/) for Mac and Linux and backslash (\) for Windows. Delimiters will not be converted if you change the platform for a rule, and you cannot enter the incorrect delimiter in a rule.
- Paths must meet other requirements of the chosen platform, including not using characters that are illegal in that file system (e.g., no colons (:)) in Mac paths) and not exceeding length limits.
- Any macros used in a path must be specific to the rule's platform (i.e., Windows, Mac, or Linux).

Using Wildcards in Rules

You can use wildcard characters in the Path and Process fields. Asterisk (*) indicates zero or more characters and question mark (?) indicates one character. You can use wildcards to specify partial paths or multiple paths for directories that appear in different locations on different computers (although macros might be a more effective way to accomplish this – see [Using Macros in Rules](#)).

In most cases, wildcards are not allowed *inside of* macros. However, they are allowed in cmdline macros and in certain parts of OnlyIf macros. If you are using a metadata-based OnlyIf macro, you cannot use wildcards in paths within that macro. You can, however, use them in with other parts of the macro, for example, to match a company name.

The number of wildcards in a path or process specification is not restricted. For example, you could define a path as: `*\Win*\folder?\`

Caution

When you use wildcards, do not create a rule so broad that it will interfere with activity in a directory that is required for legitimate use by another application or the operating system. Don't use the asterisk wildcard by itself in the path field, especially with rules that block all executions or writes, unless you are certain it will not interfere with necessary operations on agent computers. Use similar caution with wildcards when creating exceptions to restrictions created by other rules.

Automatic Path Conversions

When a rule is processed, file paths in a process field undergo several automatic path conversions if they contain certain symbols:

- Any path that ends with a backslash (Windows) or forward slash (Mac and Linux) has the "*" wildcard added at the end of the path.
- Any path that has no slash or drive letter has "*" (for Windows) or "*" (for Mac and Linux) added at the beginning of the path.
- In Windows rules, drive letters may be used in a path as long as they are for local fixed volumes. Mapped drive letters should not be used because there is no guarantee that the mapping exists on all computers.
- In Windows rules, the string "*" applies to all attached storage volumes except for floppy disks and CD/DVD-ROMs.

Specifying Devices in Paths in Windows Rules

In Windows rules, you can create rules that apply to processes on some or all devices on the agent computer by including `\device\` in the path. For example:

`\device*\` specifies all devices.

`\device\harddisk*\` specifies attached storage volumes except for floppy disks and CD-ROMs.

`\device\cdrom*\` specifies CD-ROM devices.

Using Macros in Rules

Custom rules support certain macros in the Path and Process fields. You can see a menu of macros by typing the left angle bracket (<) character in either of these fields.

The following types of macros are supported in Custom Rules, and many are supported in other rules where appropriate, including Memory Rules, Registry Rules, Script Rules, File (name) Bans and Rapid Configs:

- **Path macros** – There are two types of path macro:
 - Built-in App Control path macros (e.g., <AppData>) are a subset of the well known folders for Windows – some may operate across platforms.
 - You can insert a macro for any location that corresponds to a Windows known folder by using its GUID.

Both types always identify a location rather than a specific file. Path and Registry macros can be an effective for dynamically identifying a path on all computers for the specified platform, even when the files you want to affect are in different locations on different computers.

- **Registry macros** – These macros allow you to extract file paths and names from keys in the Windows Registry.
- **OnlyIf macros** – These macros specify conditions required for a given parameter in a rule to be expanded. If you want rules to affect only a subset of your agents, you can apply OnlyIf macros to all parameters to dictate which agents accept the rule, independent of which policy they are in.
- **Command-line macros** – These macros specify that a rule will only match if the process command line matches the specified criteria. They can always be used in the process field, and can be used in the target field if the target is a process (for example, process create).
- **Certificate/hash-matching macros** – These macros behave much like the OnlyIf macros, making rules affect an agent only if conditions match specified criteria. For example, you can specify that a rule is applied only if the issuer name for the certificate for a file in the specified path matches a specified string.

The console displays an error message if you enter an invalid macro.

Notes

A *path* macro can be used only at the beginning of a *Path* or *File* in a rule (i.e., with no other text before it in the string). *OnlyIf* and *Registry* macros can be used anywhere in the Path or File specification for a Windows rule.

Most macros are for the Windows platform.

Common Path Macros

App Control provides built-in path macros that are based on a subset of the well-known folders for a platform (CSIDLs for pre-Vista Windows versions and KNOWNFOLDERIDs for Vista and later). Each path macro consists of a unique string surrounded by angle brackets. For example, the macro <MyDocuments> in a Windows rule identifies the My Documents folder for each user on each Windows computer, regardless of its actual location on an individual computer. Full descriptions of CSIDLs are provided at:

[http://msdn.microsoft.com/en-us/library/windows/desktop/bb762494\(v=vs.85\).aspx](http://msdn.microsoft.com/en-us/library/windows/desktop/bb762494(v=vs.85).aspx)

Full description of KNOWNFOLDERIDs are provided at:

[http://msdn.microsoft.com/en-us/library/windows/desktop/dd378457\(v=vs.85\).aspx](http://msdn.microsoft.com/en-us/library/windows/desktop/dd378457(v=vs.85).aspx)

Note

In addition to using the shorter, built-in macros for known folders, you can use the GUID string for any known folder documented in the link above. See “Using Windows KNOWNFOLDER GUIDs in Macros” on page 414.

Because a path macro always represents a directory, it is processed as if it is followed by the directory delimiter (slash or backslash), even if you have not added one. For example, <AppData> in a Windows rule is interpreted as “<AppData>” before it is expanded, and it applies to the Application Data directory and all of its files, subdirectories, sub-subdirectories, etc. Similarly, <AppData>myapp\ is interpreted as “<AppData>\myapp\”. If you add a backslash yourself, the rule processor does not add a second one.

To see the menu of macros, type a left angle bracket (<) as the first character in the Path or File box or the Process box on the add rule page. As you type, the auto-complete menu adjusts to show only those choices matching the string you have typed so far for the platform you have chosen. Table 54 shows the available path macros for Windows rules.

The table includes a “Per User” column to indicate which macros are expanded based upon the logged in user. There is a delay after a user is logged in before rules tied to that user take effect, and this delay varies by how many rules you have and how long it takes to expand macros or group membership in them. Because of this, rules with user-specific macros or that specify a user-group may not take effect immediately after a user logs on.

Important

If you need a rule to be effective as soon as possible after a user logs on, do not use any of the “Per User” macros shown in the table, and do not specify a user *group* in the rule. Rules that specify a *username* or *SID* are always active and won't be affected by this delay.

Table 54: Windows Path Macros for Rules

Macro	Per User	Description
<AppData>	Yes	Directory that serves as a common repository for application-specific data. Maps to: <ul style="list-style-type: none"> • CSIDL_APPDATA • FOLDERID_RoamingAppData
<CommonAppData>	No	Directory that contains application data used by and accessible to <i>all</i> users. This folder is used for application data that is not user specific. For example, an application can store a spell-check dictionary, a database of clip art, or a log file here. Maps to: <ul style="list-style-type: none"> • CSIDL_COMMON_APPDATA • FOLDERID_ProgramData

Macro	Per User	Description
<CommonDesktopDirectory>	No	Directory that contains files and folders that appear on the desktop for all users. Maps to: <ul style="list-style-type: none"> • CSIDL_COMMON_DESKTOPDIRECTORY • FOLDERID_PublicDesktop
<CommonDocuments>	No	Directory that contains documents that are common to all users. Maps to: <ul style="list-style-type: none"> • CSIDL_COMMON_DOCUMENTS • FOLDERID_PublicDocuments
<CommonPrograms>	No	Directory that contains the directories for the common program groups that appear on the Start menu for all users. Maps to: <ul style="list-style-type: none"> • CSIDL_COMMON_PROGRAMS • FOLDERID_CommonPrograms
<CommonStartMenu>	No	Directory that contains the programs and folders that appear on the Start menu for all users. Maps to: <ul style="list-style-type: none"> • CSIDL_COMMON_STARTMENU • FOLDERID_CommonStartMenu
<CommonStartup>	No	Directory that contains the programs that appear in the Startup folder for all users. Maps to: <ul style="list-style-type: none"> • CSIDL_COMMON_STARTUP • FOLDERID_CommonStartup
<Cookies>	Yes	Directory that serves as a common repository for Internet cookies. Maps to: <ul style="list-style-type: none"> • CSIDL_COOKIES • FOLDERID_Cookies
<DesktopDirectory>	Yes	Directory used to physically store file objects on the desktop (not the desktop folder itself). <ul style="list-style-type: none"> • CSIDL_DESKTOPDIRECTORY • FOLDERID_Desktop
<InternetCache>	Yes	Directory that serves as a common repository for temporary Internet files. Maps to: <ul style="list-style-type: none"> • CSIDL_INTERNET_CACHE • FOLDERID_InternetCache
<LocalAppData>	Yes	Directory that serves as a data repository for local (non-roaming) applications. Maps to: <ul style="list-style-type: none"> • CSIDL_LOCAL_APPDATA • FOLDERID_LocalAppData
<MyDocuments>	Yes	Virtual folder that represents the My Documents folder. The file system directory used to physically store a user's common repository of documents. Maps to: <ul style="list-style-type: none"> • CSIDL_PERSONAL • FOLDERID_Documents

Macro	Per User	Description
<Profile>	Yes	User's profile folder. Maps to: <ul style="list-style-type: none"> • CSIDL_PROFILE • FOLDERID_Profile
<ProgramFiles>	No	Program Files folder. Maps to: <ul style="list-style-type: none"> • CSIDL_PROGRAM_FILES • FOLDERID_ProgramFiles
<ProgramFilesx86>	No	32-bit Program Files folder. Maps to: <ul style="list-style-type: none"> • CSIDL_PROGRAM_FILESX86 • FOLDERID_ProgramFilesX86
<ProgramFilesCommon>	No	Folder for components shared across applications. Maps to: <ul style="list-style-type: none"> • CSIDL_PROGRAM_FILES_COMMON • FOLDERID_ProgramFilesCommon
<ProgramFilesCommonx86>	No	32-bit Program Files folder. Maps to: <ul style="list-style-type: none"> • CSIDL_PROGRAM_FILES_COMMONX86 • FOLDERID_ProgramFilesCommonX86
<Programs>	Yes	Directory that contains the user's program groups (which are themselves file system directories). Maps to: <ul style="list-style-type: none"> • CSIDL_PROGRAMS • FOLDERID_Programs
<RecycleBin>	Yes	Directory for the Recycle Bin. The location depends on the type of operating system and file system. Maps to: <ul style="list-style-type: none"> • CSIDL_BITBUCKET • FOLDERID_RecycleBinFolder
<StartMenu>	Yes	Directory that contains Start menu items. Maps to: <ul style="list-style-type: none"> • CSIDL_STARTMENU • FOLDERID_StartMenu
<Startup>	Yes	Directory that corresponds to the user's Startup program group. Maps to: <ul style="list-style-type: none"> • CSIDL_STARTUP • FOLDERID_Startup
<System>	No	The platform-specific Windows System folder. Maps to: <ul style="list-style-type: none"> • CSIDL_SYSTEM • FOLDERID_System
<Systemx86>	No	32-bit "System" folder on both 32-bit and 64-bit operating systems. Allows you to specify that a rule applies only to 32-bit versions of utilities. Maps to: <ul style="list-style-type: none"> • CSIDL_SYSTEMX86 • FOLDERID_SystemX86

Macro	Per User	Description
<Windows>	No	The Windows directory or SYSROOT. This corresponds to the %windir% or %SYSTEMROOT% environment variables. Maps to: <ul style="list-style-type: none"> • CSIDL_WINDOWS • FOLDERID_Windows

Using Windows KNOWNFOLDER GUIDs in Macros

The macros in [Table 54](#) provide shortcuts to a significant number of the folders that you might want to include in a path for a rule. However, they do not cover all of the possible folders in the list of Windows known folder IDs. If necessary, you can use Microsoft's unique identifier (GUID) for a known folder in a rule macro, whether or not that folder has a built-in App Control shortcut. The GUIDs for known folders are shown at:

[http://msdn.microsoft.com/en-us/library/windows/desktop/dd378457\(v=vs.85\).aspx](http://msdn.microsoft.com/en-us/library/windows/desktop/dd378457(v=vs.85).aspx)

For example, if you wanted to include the Applications folder (FOLDERID_AppsFolder) in a rule path, you would look it up at the location above and use its GUID in a macro as shown here: `<KnownFolder:{1e87508d-89c2-42f0-8a7e-645a0f50ca58}>`

where `{1e87508d-89c2-42f0-8a7e-645a0f50ca58}` is the GUID.

On systems where the known folder doesn't exist, the rule will simply not "expand"; that is, the path will not be converted from the macro to an actual path on the endpoint's, and so the rule will not do anything. This could be the case in the following situations:

- Windows versions prior to Vista cannot process known folder GUIDs.
- Some known folders were introduced in later versions of Windows. For example, FOLDERID_AccountPictures was introduced in Windows 8, and so rules that use this macro will never be applied to endpoints running earlier versions.
- Some known folders might have been deleted on specific systems.

Some known folders are user-specific, and macros using their GUIDs must specify the user to be effective. User-specific folders are listed as Folder Type PERUSER on the Microsoft page for Known Folders. If you use one of these folders in a macro, you must specify the intended users in the rule. Otherwise, the rule will only be expanded for the SYSTEM user's version of the known folder. In addition to making the rule ineffective for other users, this often means it will never be triggered at all.

For example, to receive a report whenever a file is written to a user's Downloads directory (FOLDERID_Downloads), you could create a rule that included the following fields:

- Rule Type: Advanced
- Operation: Write
- Write Action: Report
- Path Or File: `<KnownFolder:{374DE290-123F-4565-9164-39C4925E467B}>*`
- Process: *
- User Or Group: Authenticated Users

If you left out the "User: Authenticated Users" specification, the rule would not expand at all since the "SYSTEM" user does not have a Downloads directory.

OnlyIf Macros

The Custom Rules page includes controls that specify conditions under which a rule is applied, for example, by restricting the rule to computers in certain policies. As a supplement to these user interface controls, you can also use OnlyIf macros to specify other conditions for a rule.

OnlyIf macros may be used in Custom Rules, Registry Rules, Memory Rules, Script Rules and File Rules (by name). You enter OnlyIf macros in rule fields that accept paths. In the case of Custom Rules, this includes the Path or File and Process fields.

Note

OnlyIf macros work with App Control Agents beginning with v7.2.0. They do not function on earlier agents. In addition, some OnlyIf macros may require the most recent agent.

The syntax of OnlyIf macros is:

<OnlyIf:condition:value>

For example, **<OnlyIf:Hostname:Laptop-7>** instructs the agent to apply the parameter that follows it only if the system on which the action is attempted is Laptop-7. This could be entered into the Path or File field of a Custom Rule.

Add Custom Rule

General

Rule Name:

Description:

Status: Enabled Disabled

Definition

Platform:

Rule Type:

Path or File: i

i

Process: i

Rule Applies To

Policies: All Current and Future policies Selected policies

If a rule uses a single path and a single process, you can put the OnlyIf macro in either field and it will have the same effect. For rules with multiple path or multiple processes, the location of the OnlyIf macro could be significant. For example, if you have a rule with three different paths but a single process, adding an OnlyIf clause to the process field makes it apply to all paths. If you wanted the OnlyIf test used only for one path, you would put the OnlyIf macro in that path field only. If you want it applied to multiple paths, it must precede each of them.

OnlyIf macros use the same wildcard comparison logic used for target/process name matching in Custom Rules. Strings in the macros are case-insensitive and some arguments can use * and ? wildcards. You can chain multiple macros together in one path.

[Table 55](#) lists and describes the conditions you can use in OnlyIf macros, and provides examples of each.

Some OnlyIf macros rely on metadata of the files that match the other specifications of a Custom Rule. These can be used to ensure that rules apply only apply if a particular version of a product is installed. Not all files will have all of the metadata types for which macros are available. The format of metadata-based OnlyIf macros is:

<OnlyIf:MetadataType:Pattern:Filename>

where Pattern is the (possibly wildcarded) pattern you want to compare the metadata field against, and Filename is the full path of the file whose metadata you want to test. The pathname argument cannot contain wildcards.

For example, if you include the following in a process or path field:

<OnlyIf:Company:*Carbon Black*:<ProgramFilesX86>\bit9\parity agent\parity.exe>

the rule is expanded only if the binary for parity.exe on the path shown contains “Carbon Black” in the Company field of its metadata. Because of the wildcards, this would match files whose Company field showed only “Carbon Black” but also where the company name included “Inc” and commas or periods in the name string. [Table 55](#) lists the OnlyIf macros that rely on file metadata in a separate section.

Note

Some macros work only on certain platforms. For example, rules based on file metadata are likely to be platform-specific. You can create a rule in Report mode and observe its results if you are unsure of whether the rule you create applies to all of the platforms on which you run agents.

Table 55: OnlyIf Macros in Rules

OnlyIf Condition	Description and Example
ConnectedToServer	Expand only if the agent is connected to the server (Yes) or only if the agent is disconnected (No). Example: <OnlyIf:ConnectedToServer:Yes>
ProcessorArchitecture	Expand only if the processor for the agent computer matches the value specified. The choices are x86 (32-bit) and x64 (64-bit). Example: <OnlyIf:ProcessorArchitecture:x86>
HostName	Expand only if the NETBIOS machine name matches the specified string. Example: <OnlyIf:HostName:*BSMITH-1*> Platform Note: This condition is effective for all OS platforms.
DomainName	Expand only if the agent computer is in the specified domain. Example: <OnlyIf:DomainName:*mycompany.local> Platform Note: This condition is effective for all OS platforms.
HardwareManufacturer	Expand only if the computer manufacturer matches the specified string. Example: <OnlyIf:HardwareManufacturer:*Dell*>

OnlyIf Condition	Description and Example
HardwareModel	Expand only if the computer model matches the specified string. Example: <OnlyIf:HardwareModel:*XPS>
ServiceName	Expand only if a service exists with name that matches the specified string. Example: <OnlyIf:ServiceName:*Parity Server*>
ServiceDisplayName	Expand only if a service exists with a display name that matches the specified string. Example: <OnlyIf:ServiceDisplayName:*Parity Server*>
Driver	Expand only if the specified driver is loaded. Example: <OnlyIf:Driver:mfehick>
Virtualized	Expand only if the agent is running on a VM (if the value is 1), or is not running on a VM (if the value is 0). Example: <OnlyIf:Virtualized:1>
DEPSupported	Expand only if DEP is supported on this system (if the value is 1), or if not supported (if the value is 0). Example: <OnlyIf:DEPSupported:0>
RegKeyExists	Expand only if the specified registry key exists. Example: <OnlyIf:RegKeyExists:HKLM\Software\Foo>
RegValueExists	Expand only if the specified registry value exists. Example: <OnlyIf:RegValueExists:HKLM\Software\Foo>
RegValues	Expand only if the data contained in the specified key matches the specified pattern. In the example, Foo is the key and *Bar* is the pattern. Example: <OnlyIf:RegValues:HKLM\Software\Foo:*Bar*>
HostId	Expand only for a computer whose HostID matches the specified number. Example: <OnlyIf:HostId:5> Platform Note: This condition is effective for all OS platforms.
FileExistsOnDisk	Expand only if a file with the specified name exists on the disk. This can be any file accessible to the local system user, and does not need to be tracked by App Control. In the example below, if c:\windows\system32\foo.txt exists, then the system would create a rule that targets d:\foo.exe. note that a full path must be provided in the OnlyIf clause Example: <OnlyIf:FileExistsOnDisk:<System>\foo.txt>d:\foo.exe
FileIsTracked	Expand the rule only if a file with the specified name is present and is an "interesting" file that is tracked by the App Control Agent. A full path must be used. Example: <OnlyIf:FileIsTracked:<System>\calc.exe>
HashExists	Expand the rule only if a file with the specified hash is present and is an "interesting" file that is tracked by the App Control Agent. Example: <OnlyIf:HashExists:1c94cd9e3ee959ff6002eca3c5e7e7fdb9158657>
Bit9Version:Is	Expand the rule only if the App Control Agent version matches that specified. Example: <OnlyIf:Bit9Version:Is:7.2.0.233> Platform Note: This condition is effective for all OS platforms.

OnlyIf Condition	Description and Example
Bit9Version:Atleast	Expand the rule only if the App Control Agent version is the version number specified or greater. Example: <OnlyIf:Bit9Version:Atleast:7.2.0.233> Platform Note: This condition is effective for all OS platforms.
Bit9Version:AtMost	Expand the rule only if the App Control Agent version number is not greater than the version specified. Example: <OnlyIf:Bit9Version:AtMost:7.2.0.233> Platform Note: This condition is effective for all OS platforms.
OSVersionIs	Expand the rule only if the operation system version (major.minor.point) on the agent system matches that specified. Example: <OnlyIf:OSVersionIs:10.6.8> Platform Note: This condition is effective for all OS platforms.
OSVersionAtleast	Expand the rule only if the operation system version (major.minor.point) on the agent system is the version specified or greater. Example: <OnlyIf:OSVersionAtleast:10.6.8> Platform Note: This condition is effective for all OS platforms.
OSVersionAtMost	Expand the rule only if the operation system version (major.minor.point) on the agent system is the version specified or less. Example: <OnlyIf:OSVersionAtMost:10.6.8> Platform Note: This condition is effective for all OS platforms.
OSVersionString	Expand the rule only if the detailed operating system description on the agent system matches the specified pattern. Example: <OnlyIf:OSVersionString:*Windows Server 2008*> Platform Note: This condition is effective for all OS platforms.
ServerEdition	Expand the rule only if this is a server edition of the operating system (if the value is 1) or if it is not (if the value is 0). Example: <OnlyIf:ServerEdition:1>
OnlyIf Macros based on File Metadata	
BuildAttributes	Expand only if the build attributes metadata for a file matches the specified string.
BuildTime	Expand only if the build time metadata for a file matches the specified string.
PrivateBuild	Expand only if metadata indicates that this is a private build.
SpecialBuild	Expand only if metadata indicates that this is a special build.
Comments	Expand only if the comments metadata for a file matches the specified string.
Company	Expand only if the company metadata for a file matches the specified string.
Copyright	Expand only if the copyright metadata for a file matches the specified string. Example: <OnlyIf:Copyright:Copyright 1984-2015 Adobe Systems Inc.:<ProgramFiles>\Adobe\Acrobat Reader DC\Reader\AcroRD32.exe>

OnlyIf Condition	Description and Example
Description	<p>Expand only if the file description metadata for a file matches the specified string. This is sometimes just the file name.</p> <p>Example: <OnlyIf:Description:"VMware Horizon Media Engine Library":<ProgramFiles>\VMware\VMware Horizon Media Engine\VMWMediaProvider.dll></p>
FileType	<p>Expand only if the file type metadata of a file matches that specified.</p> <p>Example: <OnlyIf:FileType:"Application (.exe)":<ProgramFiles>Dell Webcam\Dell Webcam Central\WebcamDell2.exe></p>
FileVersion	<p>Expand only if the file version metadata of the file matches the specified file version number.</p> <p>Example: <OnlyIf:FileVersion:1.0.38.0:<ProgramFiles>Dell Webcam\Dell Webcam Central\WebcamDell2.exe></p>
Language	<p>Expand only if the language metadata of the file matches the specified string.</p> <p>Example: <OnlyIf:Language:"English (United States)":<ProgramFiles>Mozilla Firefox\plugins\npMeetingJoinPluginOC.dll></p>
Manufacturer	<p>Expand only if the manufacturer metadata of the file matches the specified name.</p> <p>Example: <OnlyIf:Manufacturer:Microsoft Corporation:<CommonPrograms>\newfolder\newfile.exe></p>
OriginalName	<p>Expand only if the original name metadata for a file matches the specified string.</p> <p>Example: <OnlyIf:OriginalName:Extractor.exe:<DesktopDirectory>\e_codec.exe</p>
PackageCode	<p>Expand only if a product matching the specified package code is installed.</p> <p>Example: <OnlyIf:PackageCode:{F1D61F7C-6E4C-4902-9278-0F93131BE2D2}:<ProgramFiles>\myapp\myapp.exe></p> <p>Note: See http://blogs.msdn.com/b/pusu/archive/2009/06/10/understanding-msi.aspx for more about Microsoft package codes.</p>
ProductName	<p>Expand only if a product with the specified name is installed. In the example below, asterisk wildcards indicate that the Product Name contains "Microsoft Office".</p> <p>Example: <OnlyIf:ProductName:*Microsoft Office*:<Program Filesx86>\Microsoft Office\Office16\ACCICONS.EXE></p>
ProductCode	<p>Expand only if a product with a matching GUID is installed.</p> <p>Example: <OnlyIf:ProductCode:{F1D61F7C-6E4C-4902-9278-0F93131BE2D2}:<ProgramFiles>\myapp\myapp.exe></p> <p>Note: See http://blogs.msdn.com/b/pusu/archive/2009/06/10/understanding-msi.aspx for more about Microsoft product codes.</p>
ProductVersion	<p>Expand only if the product version metadata of a file matches the specified product version number.</p> <p>Example: <OnlyIf:ProductVersion:6.0.0170.4:<ProgramFiles>\java\jre6\java.exe></p>
TargetOS	<p>Expand only if the target operating system for the application build matches the specified value.</p>

OnlyIf Condition	Description and Example
UpgradeCode	Expand only if a product matching the specified upgrade code is installed. Example: <OnlyIf:UpgradeCode:{F1D61F7C-6E4C-4902-9278-0F93131BE2D2}:<ProgramFiles>\myapp\myapp.exe> Note: See http://blogs.msdn.com/b/pusu/archive/2009/06/10/understanding-msi.aspx for more about Microsoft upgrade codes.
AboutURL	Expand only if the About URL for the application in the path matches the specified URL.
HelpURL	Expand only if the Help URL for the application in the path matches the specified URL.
UpdateURL	Expand only if the Update URL for the application in the path matches the specified URL. This URL is used to check for application updates.

Additional Macros

In addition to the OnlyIf macros, the macros in [Table 56](#) and [Table 57](#) can be used to specify conditions under which a rule should be expanded and applied.

Notes

- You may use wildcards in cmdline macros, for example, to trigger a rule on a particular string in a path regardless of where that string appears.
- Inserting other macros inside cmdline macros is not supported.

Table 56: Command-Line Macros

Test Condition	Description and Example
<CmdLine:X>	Apply the rule only if the full command line matches specified pattern (X). Example: <CmdLine:*\\svn -vq status myfile> matches Program Files\\svn -vq status myfile
<CmdLineArgumentIdx:X:Y>	Apply the the rule only if the command line contains at least X + 1 arguments, and argv[X] matches the pattern Y. Example: <CmdLineArgumentIdx:3:"get help"> matches findstr /s /o "get help" *.inf
<CmdLineAnyArgument:X>	Apply the rule if any argument in the command line matches pattern X. Example: <CmdLineAnyArgument:bar.exe> matches both copy foo.exe bar.exe and copy bar.exe baz.exe
<CmdLineArgumentName:X:Y>	Apply the rule if the command line contains argument X and the argv[X+1] (the next argument after X) matches Y. Example: <CmdLineArgumentName:/x:*\\foo.msi"> matches msiexec /x c:\\foo.msi and also msiexec /q /x c:\\foo.msi /L* c:\\log.txt

Table 57: Additional Macros

Test Condition	Description and Example
<CertIssuer: <i>name</i> >	Evaluate the rule if the specified name string matches the Issuer name string in the certificate details for the file. Only effective in the Publisher fields for target and process.
<CertSerial: <i>number</i> >	Evaluate the rule if the specified serial number matches the Serial number in the certificate details for the file. Only effective in the Publisher fields for target and process.
<CertSHA1: <i>hash</i> >	Evaluate the rule if the specified hash value matches the SHA1 hash in the certificate details for the file. Only effective in the Publisher fields for target and process.
<CertMD5: <i>hash</i> >	Evaluate the rule if the specified hash value matches the MD5 hash in the certificate details for the file. Only effective in the Publisher fields for target and process.
<Sha256: <i>hash</i> >	Evaluate the rule if a file or process matching the path specified in the rule has the specified Sha256 hash value. This macro normally should not be used in the target field of Write rules because Write rules are evaluated before the write occurs (and therefore before the hash is known).
<HostedService: <i>servicename</i> >	This macro ensures that the agent only applies the rule to the specific svchost.exe instance hosting the specified service rather than all svchost.exe instances. For use only in the process pattern field. Example: <HostedService:EMET_Service>
<SourceNameOnly>	For rename operations, compare the path only to the source name. This macro can be placed in the "target pattern" field of any custom rule. If this macro or the <DestinationNameOnly> macro are not included in a rule specification, the agent compares the pattern against both the source and destination names, and the rule will match if either the source or destination name matches. <DestinationNameOnly> and <SourceNameOnly> cannot be used in the same field. Example: This could be used to approve files that are moved <i>out of</i> the c:\temp folder: c:\temp*<SourceNameOnly>.
<DestinationNameOnly>	For rename operations, compare the path only to the destination name. This macro can be placed in the "target pattern" field of any custom rule. If this macro or the <SourceNameOnly> macro are not included in a rule specification, the agent compares the pattern against both the source and destination names and the rule will match if either the source or destination name matches. <DestinationNameOnly> and <SourceNameOnly> cannot be used in the same field. Example: The following could be used to approve files when they are moved <i>into</i> the c:\foo folder: c:\foo*<DestinationNameOnly>.

Windows Registry Macros

For Windows rules, Registry (Reg) macros represent Windows Registry values, which you can use in a Path or Process specification. Unlike path macros, reg macros have variable content between their brackets. A Reg macro must resolve to a value, not a key.

To enter a Reg macro:

1. Begin by typing a left angle bracket (<) followed immediately by **Reg:**
2. Follow <Reg: with one of the following:
 - a. **HKLM** (or HKEY_LOCAL_MACHINE)
 - b. **HKCU** (or HKEY_CURRENT_USER)
 - c. **HKLM-SoftwareX86**
 - d. **HKLM-SoftwareX64**
 - e. **HKCU-SoftwareX86**
 - f. **HKCU-SoftwareX64**
3. Enter the rest of the path you want in this rule. Specify a value, not a key, with one exception – you can provide a key specification and follow it by a backslash to use the default value for this key.
4. Because reg macros contain variable content, they do not auto-complete. You must provide the whole path in the macro and end the macro with the right angle bracket (>). For example (using HKLM as the top-level Registry node):

```
<Reg:HKLM\YourSpecifiedPath>
```

Reg macros are evaluated on each agent the first time the rule becomes available to that agent. If the rule is valid for that computer, it is enabled. For example, it is possible to create a rule that Promotes an updater for an application called “MyApp” by using the path value written to the registry. On systems with MyApp Update installed, <Reg:HKLM\Software\MyApp\Update\path> might expand to C:\Program Files (x86)\MyApp\Update\MyAppUpdate.exe. On systems that did not include the update program, the rule would not be created.

Once evaluated, rules that use Reg macros are not re-evaluated on a computer unless certain conditions occur. This means that changes to the Registry during a session might not affect rule behavior, even if the change would enable or disable the rule. The conditions that cause "re-expansion" of rules on an agent are:

- the agent is stopped and restarted (e.g., the computer is shut down and restarted)
- a new user logs in
- the agent is reassigned to a policy with different rules
- rules are created, edited or deleted on the server
- the agent detects the end of an MSI install/upgrade
- manual re-evaluation is triggered using a special Carbon Black Support command.

Important

A rule in which you specify an HKCU-based registry macro won't become active for a specific user until a short time after the user is logged on. The delay varies depending upon how many rules you have and how long it takes to expand registry macros and other user-based fields. Avoid the HKCU macro if you need a rule to be effective immediately after login.

Entering Multiple Paths or Processes

For both the Path or File value and the Process value in a rule, you can enter more than one string. When you have entered the first Process for this rule, click the Expand button to the right of the box.

You can then add additional paths or files by typing them in the box and clicking **Add** after each one.

You can remove any file or path by selecting the file or path in the list below the Path or File box, and clicking the **Remove** button.

If you enter multiple paths or processes for a rule, the Custom Rules page shows the first path and then **(multiple)** in the relevant column for this rule. Moving the mouse over the value shows a tooltip with the complete list of paths or processes for the rule.

Specifying Processes

You can specify the Process string using the same options available for Path or File. See [“Specifying Paths and Processes”](#) on page 407 for complete details.

If you specify both a User or Group and a Process for a rule, they work together. For example, if you choose Specific Process, a matching user or group must be running a matching process for the rule to be applied. If you choose Any Process Except, the rule is applied unless *both* the User or Group and the Process match the rule definition.

Table 58: Process Menu Choices

Menu Choice	Description
Any Process	Apply the rule no matter what process attempts to execute or write files matching the rule.
Any Promoted Process	Apply the rule when a process that is <i>promoted</i> at the time the rule is evaluated attempts an action matching the rule. A promoted process is any approved process that is generated by a file marked as an installer, or has been promoted as a consequence of a custom rule, or is an approved process launched by a promoted process.
Any System Process	Apply the rule when a process that is running under the security context of the Local System user attempts an action matching the rule. This choice has the same effect as choosing Local System in the User or Group menu, but may be more efficient.
Specific Process...	Apply the rule when a process matching a string you specify attempts an action matching the rule. You can enter one or more processes in the text box below the menu.

Menu Choice	Description
Any Process Except...	Apply the rule when any process <i>except</i> one matching a string you specify attempts an action matching the rule. You can enter one or more processes in the text box below the menu.

Specifying Users or Groups

For certain rule types, you can create a rule that applies only when specific users or users in specific groups attempt an action. The choices for User or Group on the Add/Edit Custom Rule page are:

- **Any Users** – applies the rule to all users.
- **Specific User or Group...** – opens a text box below the menu, in which you can enter AD users or groups in the format *userorgroupname@domain* or *domain\userorgroupname*
Platform Note: To specify a Mac or Linux group, you must precede it with the word “group” and a colon. For example, you would enter *group:consoleusers* for the “consoleusers” group. Without the prefix, group names will be considered user names.
- For Windows rules only, there are other menu choices that are built-in Windows groups, such as **Authenticated Users** and **Local Administrators**.

Notes

- When running on Windows Vista and later, membership in pre-defined security groups like Administrators requires that the application run as an administrator. If a group definition is necessary for a rule, consider using security groups you have defined rather than the pre-defined groups
- There is a brief delay after a user logs in before group membership is established and group-based rules become effective. This delay may be longer if you have a large number of rules. If a rule must be effective as soon as possible after a user logs on, do not specify a user *group* in the rule. Rules that specify a *username* or *SID* are always active and won't be affected by this delay.
- Specifying a user or group also determines whether macros in a path are expanded. Only paths whose macros match the specified user or group are expanded, and so even if the user or group is attempting the action, if the path includes a user-related macro, paths that would evaluate to a user other than those specified are not expanded and the rule is not effective.

Rule Ranking

Custom rules have a “Rank” number and are evaluated from lowest number to highest number, beginning with the rule ranked ‘1’. By default, rules appear in their rank order, but you can re-sort the table by other columns if you choose. Also, if you filter a table, there will be gaps in the rank of rules because not all rules will be shown.

If a file matches one rule that blocks an action and another rule that allows it, the highest ranking rule (that is, the one with the lowest number), takes precedence and the lower-ranked (higher number) rule has no effect. You can change the ranking of rules if you decide that you want one of your rules to be considered before its current position.

Important

Rule ranking is significant only for rules that Block, Allow, or Prompt the user to block or allow. The highest ranking block, allow, or prompt rule that matches an attempted file action not only takes precedence but stops processing of any lower-ranked rules matching the action.

A rule whose action is Approve, Approve as Installer, Track, Report, Promote or Ignore does not stop processing of lower-ranked rules. For example, if a write attempt first matches an Ignore rule and also matches another rule with a lower rank (higher number) on the list, the second rule will also be processed.

Although not custom rules, *Internal* rules for fundamental actions in App Control, such as blocking banned files, are included in the Custom Rules table. See [“Rule Ranking and Internal Rules”](#) for suggestions about how and when you might change the order of other rules relative to internal rules.

The options for changing the rank of rules depend on how table rows are sorted and whether the table is filtered.

- If a rule table is sorted by rank and not filtered, you can use arrow buttons and “drag-and-drop” methods to change the rank of rules. The arrow buttons appear only in tables that meet these conditions.
- For any table, regardless of how it is sorted or filtered, you can click on its rank number and specify a new rank in a dialog box.

To change the rank of a rule in an unfiltered table sorted by rank:

1. On the rules page, if the rules are not currently sorted by rank, click on the Rank column head to sort them.
2. If you have applied a filter or a saved view filtering to the table, either click **Reset** in the Filters panel or choose **(none)** in the Saved View field to return the table to an unfiltered view.
3. To change the rank of a rule, you have three options:
 - In any table that displays the rank column, you can click on the rank number and enter a new rank number in the dialog box.
 - If the table is sorted by rank and not filtered, arrows appear next to the rank, and you can click the up or down arrow button next to the rule to change its rank.
 - If the table is sorted by rank and not filtered, you can hold down the left mouse button with the cursor over the rule and drag the rule to a new location.

<input type="checkbox"/> Select 52	Rank ▲	Status	Platform	Rule Type	Name
<input type="checkbox"/>	↑ ↓ 1	<input type="radio"/>	Windows	Expert	Report execution of potential risk files identified by ya
<input type="checkbox"/>	↑ ↓ 2	<input type="radio"/>	Windows	Expert	Block execution of malicious files identified by yara
<input type="checkbox"/>	↑ ↓ 3	<input checked="" type="radio"/>	Windows	Expert	Examine powershell script contents
<input type="checkbox"/>	↑ ↓ 4	<input checked="" type="radio"/>	Windows	Expert	Block powershell scripts that execute memory
<input type="checkbox"/>	↑ ↓ 5	<input type="radio"/>	Windows	Advanced	[Sample] Report whenever powershell is launched with
<input type="checkbox"/>	↑ ↓ 6	<input checked="" type="radio"/>	Windows	Expert	Do not treat these processes as .NET applications

Note

When using drag-and-drop, your target location must be visible in the current view (including rows you can scroll to but not rows that have not been loaded). If you need to move a rule to a ranking not currently shown, you can use the Click to Show More bar at the bottom of the rules table to add rows to the current view. You also can use the dialog box described in the next procedure.

To change the rank of a rule in any table:

1. Click the rank number of the rule whose rank you want to change.

<input type="checkbox"/> Select 52	Rank ▲	Status	Platform	Rule Type	Name
<input type="checkbox"/>	↑ ↓ 1	<input type="radio"/>	Windows	Expert	Report execution of potential risk files identified by ya
<input type="checkbox"/>	↑ ↓ 2	<input type="radio"/>	Windows	Expert	Block execution of malicious files identified by yara
<input type="checkbox"/>	↑ ↓ 3	<input checked="" type="radio"/>	Windows	Expert	Examine powershell script contents
<input type="checkbox"/>	↑ ↓ 4	<input checked="" type="radio"/>	Windows	Expert	Block powershell scripts that execute memory

2. In the dialog box, specify where in the rankings you would like to move the selected rule:

Examine powershell script contents

Move Rule rank

- a. Enter a rank number. The row for the rule currently at that rank does not have to be visible.
 - b. Choose whether you would like the selected rule to be before or after the rule currently at that rank. Choosing **Before** means that the rule you are moving takes over the rank number you indicated in the dialog, pushing the rule that was in that position to a lower rank. Choosing **After** means the rule you are moving is ranked one lower than the number you indicated.
3. Click **OK** in the dialog box. The rule moves to its new rank.

Rule Ranking and Internal Rules

The Custom Rules table includes Internal rules related to features presented in other parts of the console. These built-in rules are approximately equivalent to the settings you see when you view the Device and Advanced Settings on the Edit Policy page.

Advanced			
File Rules			
Custom Rules			
Memory Rules			
Registry Rules			
Publisher Rules			
Rapid Configs			
Computers			
Device Control Settings			
Name	Status	Notifiers	
Block unanalyzed scripts and executables	Active	<default>: Block unanalyzed scripts and execut	Add Edit
Block unapproved scripts	Active	<default>: Block unapproved scripts	Add Edit
Block unapproved executables	Active	<default>: Block unapproved executables	Add Edit
Block banned file names	Active	<default>: Block banned file names	Add Edit
Block banned file hashes	Active	<default>: Block banned file hashes	Add Edit
Block executables run from a network drive	Off	<default>: Block executables run from a network drive	Add Edit
Block files with banned publishers or certificates	Active	<default>: Block files with banned publishers or certificates	Add Edit
Enforce memory rules	Active	<default>: Enforce memory rules	Add Edit
Enforce registry rules	Active	<default>: Enforce registry rules	Add Edit
Enforce custom (file and path) rules	Active	<default>: Enforce custom (file and path) rules	Add Edit
Enforce tamper protection	Active	<default>: Enforce tamper protection	Add Edit
Terminate processes with banned images	Report Only	<default>: Terminate processes with banned images	Add Edit
<input checked="" type="checkbox"/> Locally approve unapproved files on transition from Visibility or Low Enforcement Level to Medium or High			

Advanced			
File Rules			
Custom Rules			
Memory Rules			
Registry Rules			
Publisher Rules			
Rapid Configs			
Computers			
Device Control Settings			
Name	Status	Notifiers	
Block writes to unapproved removable devices	Off	<default>: Block writes to unapproved removable devices	Add Edit
Block writes to banned removable devices	Active	<default>: Block writes to banned removable devices	Add Edit
Report reads from unapproved removable devices	Off	<none>	
Report reads from banned removable devices	Off	<none>	
Block executions from unapproved removable devices	Off	<default>: Block executions from unapproved removable devices	Add Edit
Block executions from banned removable devices	Active	<default>: Block executions from banned removable devices	Add Edit

For example, *Block banned file hashes* is listed as an Internal Rule on the Custom Rules page and as a setting in the Advanced Settings section of the Edit Policy page.

Software Rules					
Updateers					
Rapid Configs					
Publishers					
Users					
Directories					
Files					
Custom					
Memory					
Registry					
Scripts					
Reputation					
Saved Views: (The Current View Has Unsavd Changes - Discard)			Group By:		
(none)	Delete	Save		Add	(none) Ascending
Show Filters Show Columns Export to CSV Refresh Table					
<input type="button" value="Add Custom Rule"/> <input type="button" value="Export Rules"/> <input type="button" value="Import Rules"/> <input type="text" value="Search:"/> <input type="checkbox"/> Automatically apply					
Select	Rank	Status	Platform	Rule Type	Name
<input type="checkbox"/>	53	Disabled	All Platforms	Internal	Block executables run from a network drive
<input type="checkbox"/>	54	Enabled	All Platforms	Internal	Block executions from banned removable devices
<input type="checkbox"/>	55	Disabled	All Platforms	Internal	Block executions from unapproved removable devices
<input type="checkbox"/>	56	Enabled	All Platforms	Internal	Block writes to banned removable devices
<input type="checkbox"/>	57	Disabled	All Platforms	Internal	Block writes to unapproved removable devices
<input type="checkbox"/>	58	Enabled	All Platforms	Internal	Block files with banned publishers or certificates
<input type="checkbox"/>	59	Enabled	All Platforms	Internal	Block banned file hashes




You cannot enable, disable, modify or move Internal rules in the Custom Rules table – they do not have delete or edit or buttons or ranking arrows. The order of Internal rules cannot be changed relative to each other. However, you can change the rank of any Internal rule relative to other, non-internal Custom Rules to better control how and when different rules are enforced. You do this by moving the other rule (not the Internal rule).

The following are key situations in which you might want to change the order of Internal rules relative to other rules.

- By default, if a file has been banned but you create a Custom Rule that allows the file to execute, that rule appears higher in rank than the internal rule that blocks executions of banned hashes. Because of this, the Custom Rule takes precedence over a hash ban on that file. However, if you move the Custom Rule that allows the banned file to execute to a rank below the Internal rule *Block banned file hashes*, the file will *not* be allowed to execute. Unless you want to bypass file bans, moving the “allow” rule is recommended.
- By default, if you create a Custom Rule that allows a file to be written, it appears higher in rank than internal rules that block writing, and so the allow rule takes precedence. For example, if you create a new rule that allows writes to a device, it appears before the internal rule that blocks writes to a device. However, if you move the rule that allows device writes to a position after the *Block writes to unapproved removable devices* rule, the block rule takes precedence and a file on an unapproved device is blocked from writing, even if it matches an Allow or Prompt rule below.

To make file hash bans override custom rules that allow execution:

1. On the Custom Rules page, if the rules are not currently sorted by rank, click on the Rank column head to sort them.
2. Find the rule that allows execution of the banned file.
3. Use the down arrow to move the allow rule to a position after the *Block banned file hashes* rule. Here we also moved it immediately after *Terminate processes with banned images* so that it would be shut down if running.

<input type="checkbox"/> Select 53	Rank ▲	Status	Platform	Rule Type	Name
	59	Enabled	All Platforms	Internal	Block banned file hashes
 	↑ ↓ 60		Windows	Execution Control	Allow Mystery App to Execute
	61	Disabled	Windows	Internal	Terminate processes with banned images

Enabling and Disabling Custom Rules

If you do not want a custom rule to be effective anymore, you can either disable it, which leaves it in the custom rules table, or delete it from the table. In either case, the rule stops affecting newly discovered files. However, files that were affected by the rule before it was disabled retain any file state assigned to them by the rule.

If you think you might use the rule again, disabling it temporarily is the best choice. Rules that you are allowed to disable and enable are shown in the Custom Rules table with a toggle switch in the Status column. If the toggle button is on the right and the background is green, the rule is enabled. If the toggle button is on the left and the background is white, the rule is disabled.

Rules that you cannot enable or disable show their status in this field with text instead of the toggle switch. This includes Internal rules that are either always on or set through other configuration pages. Status also appears as text (not the toggle button) when the logged in user does not have permission to edit any rules, and also for rules set by Unified Management. Also, notice that rules that cannot be enabled or disabled do not have buttons for deleting, editing, or changing the rank of the rule.

Select 53	Rank ▲	Status	Platform	Rule Type
<input type="checkbox"/>	↑ ↓ 49	<input checked="" type="checkbox"/>	Linux	Expert
<input type="checkbox"/>	↑ ↓ 50	<input type="checkbox"/>	Mac	Expert
<input type="checkbox"/>	↑ ↓ 51	<input checked="" type="checkbox"/>	Windows	Expert
<input type="checkbox"/>	↑ ↓ 52	<input checked="" type="checkbox"/>	Windows	Expert
	53	Disabled	All Platforms	Internal
	54	Enabled	All Platforms	Internal
	55	Disabled	All Platforms	Internal

Rules may be disabled or enabled either on the Custom Rules table page or on the Add/Edit Custom Rule page.

To disable a custom rule from the Custom Rule table page:

- In the Custom Rule table, if the toggle switch in the Status column indicates that the rule is enabled, click the toggle switch. The button will move to the left, the background will become white, and the rule is disabled.

To disable a custom rule from the Add/Edit Custom Rule page:

1. In the console menu, choose **Rules > Software Rules**, and when the Software Rules page appears, click the **Custom** tab. The Custom Rules table appears.
2. Click the View Details button next to the rule you want to disable. The Edit Custom Rule page appears.
3. In the Status line, click the **Disabled** radio button, and then click the **Save** button at the bottom of the page. The rule is now disabled.

To enable a custom rule from the Custom Rule table page:

- In the Custom Rule table, if the toggle switch in the Status column indicates that the rule is disabled, click the toggle switch. The button will move to the right, the background will become green, and the rule is enabled.

To enable a custom rule from the Add/Edit Custom Rule page:

- On the Edit Custom Rule page, go to the Status field and click the **Enabled** radio button, and then click the **Save** button at the bottom of the page. The rule is now enabled.

Unified Management

If you are using Unified Management, rules that have been applied to multiple servers do not have Enabled and Disabled radio buttons. Instead, there are two commands in the Action menu on the right menu bar: Enable on All Unified Servers and Disable on All Unified Servers. See [“Unified Management of Rules”](#) on page 791 for information.

Deleting Custom Rules

Deleting a rule eliminates it permanently – there is no undo or retrieval for a deleted rule. Because of that, be sure you actually want to delete the rule. Deletion of the rules that were pre-configured in the console is not recommended.

To delete a custom rule:

1. In the console menu, choose **Rules > Software Rules**, and when the Software Rules page appears, click the **Custom** tab. The Custom Rules table appears.
2. Click the Delete button (trash can) next to the rule you want to delete, and click **OK** on the configuration dialog. The rule is now deleted.

-or-

If you are on the Edit Rule page, you can delete that rule by clicking the **Remove Rule** button at the bottom of the page and click Yes to confirm the deletion.

Unified Management

If you delete a unified rule that is currently applied to more than the server you are logged into, a different confirmation dialog allows you to choose whether to delete the rule from all servers or just the current server. See [“Unified Management of Rules”](#) on page 791 for information.

Viewing Rule Status on Computers

Depending upon the number of agents managed by your App Control Server and the number that are disconnected, not all agents might receive new or updated rules in a short amount of time. The Related Views menu on the Edit page for an enabled rule provides links to two different filtered views of the Computers page to help determine the status of the rule on agent-managed computers. The choices are:

- **All Computers that have received this rule**
- **All Computers that have not yet received this rule**

This menu does not appear for rules that have never been enabled.

Unified Management

The results of these commands show only those agent-managed computers reporting to one server. If you are using Unified Management, you can click on the rule under each server on the Custom Rules page to see the computers with the rule on that server. See [“Unified Management of Rules”](#) on page 791 for information.

Exporting and Importing Rules

Certain rules created on one App Control Server may be exported to a file and then imported from the file to another App Control Server. The following rule types are exportable:

- Custom Rules
- Registry Rules
- Memory Rules

Rule export and import can be useful in several different situations:

- **Transfer from Test to Production Environments** - You may want to create, test and perfect rules in a lab environment before you apply them to your production server. With rule export and import, once you are satisfied with rule behavior, you can export rules from the test server and import the rule file to the production server, eliminating the effort and error potential of manual re-entry of the rule fields.
- **Rule Sharing in the Carbon Black Community** - Users who have created a rule or set of rules they consider particularly useful can make their rule(s) available in a file that may be imported by other members of the community.
- **Solutions from Carbon Black Support** - If you need assistance in creating a rule to accomplish a particular outcome and have not found an appropriate rule on the User Exchange, Carbon Black Support or Services may be able to provide an appropriate rule for you to import it to your server.

Software Rules

Updaters Rapid Configs Publishers Users Directories Files **Custom** Memory Registry Scripts

Saved Views: (none) Delete Save Add Group By: (none) Ascending

Show Filters Show Columns Export to CSV Refresh Table

+ Add Custom Rule Export Rules Import Rules Search:

Select	Rank	Status	Platform	Rule Type	Name
<input type="checkbox"/>	1	<input checked="" type="checkbox"/>	Windows	Advanced	Ignore EXEMaker Temp f
<input type="checkbox"/>	2	<input checked="" type="checkbox"/>	Windows	Execution Control	Allow EXEMaker executi

You export and import rules on the Software Rules page tab that shows the rule table for each of the exportable rule types. One or more rules of the same type may be included in an export file, but rule types are not mixed in the same file; for example, you cannot mix custom and registry rules in the same file.

Exported rule files are downloaded using the standard download mechanism and target location for the browser in which the console is displayed. They have the extension **.rules**. As new rules are created or existing rules changed, new export files may be generated as needed.

Rule files are encrypted to prevent tampering. When a file is exported, it can be further protected with an optional password.

Notes

- Rules may be exported and imported only from and to App Control Servers at version 7.2.1 and later.
- If you are using Unified Management, rules can be copied directly to managed client servers without using a file as an intermediary. See [“Unified Management of Rules”](#) on page 791.

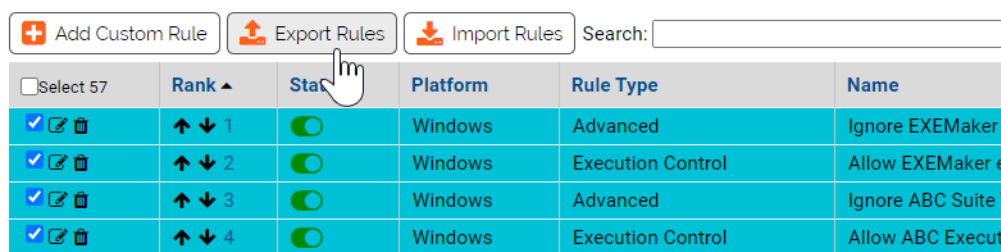
Exporting Rules

When exporting rules, consider the destination of the rules. You might export one set of rules for internal use and another to share with other members of the Carbon Black community. The following are some points to keep in mind when exporting rules:

- **Proprietary Information** – It is possible that a rule could reveal information that you would prefer not to share outside your organization. This might include path or user names, or comments in the Description field of a rule. Note that you can choose not to export user and group specifications that are not well known SIDs.
- **Environment Dependencies** – For rules shared outside your environment, hard paths could limit the usefulness of a rule. Rules using macros might be more portable.

To export rules to a file:

1. In the console menu, choose **Rules > Software Rules** and click on the tab for the type of rules (Custom, Registry, or Memory) you want to export.
2. All of the rules you want to export to one file must be showing on the current page. On dynamically scrolling pages such as the Custom Rules page, if the current “page” is long enough that you must scroll the browser view to see rules at the bottom, you may select any rules that you can scroll to. To make it easier to see all of the rules you want to export at the same time, you can use filters, grouping, or a Saved View to change the page content.
3. Check the box next to each rule you want to export and then click the **Export Rules** button.



The screenshot shows the top navigation bar with three buttons: 'Add Custom Rule', 'Export Rules', and 'Import Rules'. Below the buttons is a search field. The main content area displays a table of rules with columns for selection, rank, status, platform, rule type, and name. A mouse cursor is hovering over the 'Export Rules' button.

<input type="checkbox"/> Select 57	Rank ▲	Status	Platform	Rule Type	Name
<input checked="" type="checkbox"/>	1	<input type="checkbox"/>	Windows	Advanced	Ignore EXEMaker e
<input checked="" type="checkbox"/>	2	<input type="checkbox"/>	Windows	Execution Control	Allow EXEMaker e
<input checked="" type="checkbox"/>	3	<input type="checkbox"/>	Windows	Advanced	Ignore ABC Suite
<input checked="" type="checkbox"/>	4	<input type="checkbox"/>	Windows	Execution Control	Allow ABC Execut

The Export Rules dialog appears. It shows the number of rules to be exported, provides a field in which to name the file, and includes other export options.

4. Enter the file name (without extension) for the new Export File. This is the only mandatory field.

Export Rules

Number of rules to export: 4

Export File Name: ABC and EXEmaker Rules .rules

Password (optional): Enter password to encrypt export file

Confirm Password: Enter password to confirm

Export SIDs

Description: Enter description

Export Cancel

5. Exported rules files are not readable as text, but if you would like to further protect the file, enter and confirm a password. Be sure to have the password available for the users who will be importing the file.
6. Check the **Export SIDs** box if all of the following is true:
 - One or more of the rules you are exporting specify that they should be applied only for specific users or groups.
 - These users or groups are *not* one of the well-known security identifiers (SIDs) on Windows systems.
 - You are planning to import these rules to a server on which your non-well-known SIDs will be present. This is more likely to be the case if you are transferring rules within the same organization.
7. If you choose, add a Description that will help anyone importing rules from this file better understand what their purpose is.
8. When you are ready to save the Export File, click the **Export** button. The dialog closes and the rules file is created using the standard download mechanism of the browser running the console. For example, if you entered “New Custom Rules” in the Export File Name field, a file named “New Custom Rules.rules” might be written to the Downloads folder.

Once you have exported rules to a file, you can copy it to the host of another App Control Server or make it available via a network connection for import.

Importing Rules

When you want to import rules from another server, you need access to a rules file. In addition, if the file was passworded, you need the password to open it in the import dialog and choose the rules to import.

The steps for importing rules are shown in [“To import rules from a file:”](#) on page 437. Before doing an import, it is highly recommended that you read the following sections.

Selecting Rules to Import

When you enter the name of a rules file in the Import Rules dialog, the file is checked to determine whether it is properly formed and the rule type matches the page on which the import is being attempted. If it is passworded, you are prompted to enter the password. Assuming it passes these checks, the rules it contains are listed in the dialog box.

Import Rules

Import File Name ABC and EXE...Rules.rules

Description Rules to allow execution and ignore temp file writes for two new company apps.

Select Rules To Import

Overwrite Existing Rules Import SIDs

<input type="checkbox"/>	Exists	Name	Type	Platform	Action
<input type="checkbox"/>	Yes	Allow ABC Executions	Custom	Windows	Allow executes
<input type="checkbox"/>	Yes	Ignore ABC Suite Temp Files	Custom	Windows	Ignore writes
<input type="checkbox"/>	Yes	Allow EXEMaker executions	Custom	Windows	Allow executes
<input type="checkbox"/>	Yes	Ignore EXEMaker Temp Files	Custom	Windows	Ignore writes

Table 59 describes the fields on the Import Rules dialog, most of which are described in more detail later in this section.

Table 59: Import Rules Dialog Fields

Field	Description
Import File Name	The name of the file from which rules will be imported to this server. Enter file names using the Choose File button and file chooser dialog.
Description	The description provided when the rules were exported, if any.
Overwrite Existing Rules	If not checked (the default), there is no checkbox next to rules that already exist on the target server. If checked, all rules in the table have checkboxes, and you may choose to overwrite an existing rule.
Import SIDs	If not checked, user and group specifications in rules are not imported if those users or groups are not well-known Security IDs (SIDs), such as Local Administrator. If checked, all user and group specifications in rules are imported. Note that there is a matching option for exporting rules, and so some rules in an exported file might not include user and group specifications that are in the original rule.
Enter Password	Appears only if a password was specified during rule export. If present, shows a field in which to enter the password to open this file and an Open Import File with Password button.

Field	Description
Rules Table	<p>All rules included in the import file are listed in a table. The row for each rule includes the following columns:</p> <ul style="list-style-type: none"> • (Checkbox) – A checkbox appears next to each rule that can be selected for import. • Exists – Indicates whether the rule already exists on the target server. • Name – The name of the rule as it appears on the rules page. • Type – The type of rule as indicated by the tab on which it appears (Custom, Memory, or Registry). • Platform – The operating system/platform to which the rule applies (Windows, Mac, Linux). • Action – The action type taken by the rule.

Each rule on a App Control Server has a globally unique identifier (GUID), and that ID is included when it is exported to a file. When a rules file is chosen for import, the GUIDs of the incoming rules are compared to the GUIDs of existing rules, and if a rule already exists on the server, that fact is shown on the Import Rules dialog.

Depending upon the source of the rules (internal to your organization, the Carbon Black community, Carbon Black Support), you might make different decisions about which rules to import. You do not have to import all rules in a file. A checkbox next to each available rule allows you to choose which rules to import.

By default, any rules in the import file that already exist on the server do not have a checkbox next to them. However, there is a checkbox named Overwrite Existing Rules that activates checkboxes for these rules, allowing you to import any rule (including existing rules) listed on the page.

Import Rules

Import File Name: Choose File ABC and EXE...Rules.rules

Description: Rules to allow execution and ignore temp file writes for two new company apps.

Select Rules To Import

Overwrite Existing Rules Import SIDs

Exists	Name	Type	Platform	Action
<input type="checkbox"/> Yes	Allow ABC Executions	Custom	Windows	Allow executes
<input type="checkbox"/> Yes	Ignore ABC Suite Temp Files	Custom	Windows	Ignore writes
<input type="checkbox"/> Yes	Allow EXEMaker executions	Custom	Windows	Allow executes
<input type="checkbox"/> Yes	Ignore EXEMaker Temp Files	Custom	Windows	Ignore writes

Import Cancel

Differences in Settings for Imported Rules

Rules contain a variety of field types, including processes and paths, actions to take, and notifiers to use if a block is involved. Most of the settings for an imported rule remain the same as they were on the server from which they were exported, but there are some variations depending on the following factors:

- Whether an imported rule is *new or updates an existing rule* on the target server
- Whether the rule specifies that it applies only to *certain policies*
- Whether the rule specifies that it applies to *certain users or groups*


The following setting differences depend upon whether a rule is new or already existed on the server:

- **Enabled or Disabled** – New rules are disabled when imported and must be enabled to take effect. This gives you the ability to customize a rule, including providing any site-specific policy or user fields, before enabling it. When existing rules are overwritten by an import, the enable/disable settings on the target server are kept.
- **Rank** – New rules are ranked at the highest level when imported. Existing rules that are overwritten by an import maintain their previous relative rank on the target server (moving down in rank accordingly if new rules were also part of the import)
- **Notifier** – If a new imported rule requires a notifier (i.e., if it blocks an action), the default notifier is used. If an imported rule overwrites an existing rule, the notifier specified in the existing rule will be kept.

Some rules are specified to apply only to computers in certain policies. However, policies on one server might not exist on another. If an imported rule is new, any previous policy specification is removed and the rule applies to all policies. If an imported rule overwrites an existing rule, the policy setting in the existing rule on the target server is maintained – any policy specification in the rule from the exporting server is not applied.

Some rules are specified to apply only if certain users or members of certain groups are taking an action. There are user and group names that are well known Security Identifiers (SIDs) that can be expected to be available on all Windows computers. However, users and groups that are not well known might not exist on computers to which rules are imported. If an exported rule specifies users or groups, the results of an import depend on whether the user or group is well known and on whether several things:


- All well-know SIDs will always be exported and imported in a rule specification.
- If the Export SIDs checkbox was checked on the Export Rules dialog when the rules were exported, specifications for users and group that are not well-known will also be exported with their rules.
- If the Import SIDs checkbox is checked on the Import Rules dialog, specifications for users and group that are not well-known will also be imported, *if they were exported with the rules.*



- If a both well-known and non-well-known SIDs are specified in an exported rule and the Import SIDs checkbox is not checked, the rule is exported with the well-known users or groups only. If the rule only specifies users or groups that are not well known, the user or group specification is removed from the rule and it applies to all users.

To import rules from a file:

1. In the console menu, choose **Rules > Software Rules** and click on the tab for the type of rules (Custom, Registry, or Memory) you want to import.
2. Click the **Import Rules** button.
The Import Rules dialog appears.
3. Click the **Choose File** button to open a standard Windows file chooser dialog and choose the file whose rules you want to import.
If no password is required, the Import Rules dialog show a table of the available rules, as shown in [“Selecting Rules to Import”](#) on page 433.
4. If the file requires a password, the dialog shows fields for password entry before displaying any rule names. In this case, enter the password and click Open Import File With Password.



When the password is verified the dialog shows the table of available rules, as shown in [“Selecting Rules to Import”](#) on page 433.

5. If you want to include any user or group fields that do are not well-known SIDs, check the **Import SIDs** box. See [“Selecting Rules to Import”](#) on page 433 for an explanation of these feature.
6. By default, there is no checkbox for any rule that is already present on the target server. If you want the option of choosing to overwrite one or more existing rules, check the **Overwrite Existing Rules** box. See [“Differences in Settings for Imported Rules”](#) on page 436 for what happens when you overwrite an existing rule.

- In the dialog, examine the information about each rule, check the box next to each rule you want to import, and then click the **Import** button. The dialog closes and the rules are imported to the server. Rules that have been imported appear in bold italic on the rules page for the duration of the current session.

Software Rules

[Updaters](#)
[Rapid Configs](#)
[Publishers](#)
[Users](#)
[Directories](#)
[Files](#)
[Custom](#)
[Memory](#)
[Registry](#)
[Scripts](#)
[Reputation](#)

Saved Views: (none)
Group By: (none)

[Show Filters](#) | [Show Columns](#) | [Export to CSV](#) | [Refresh Table](#)

Automatically app

<input type="checkbox"/> Select 57	Rank	Status	Platform	Rule Type	Name
<input type="checkbox"/>	1	<input checked="" type="checkbox"/>	Windows	Advanced	<i>Ignore EXEMaker Temp Files</i>
<input type="checkbox"/>	2	<input checked="" type="checkbox"/>	Windows	Execution Control	<i>Allow EXEMaker executions</i>
<input type="checkbox"/>	3	<input checked="" type="checkbox"/>	Windows	Advanced	<i>Ignore ABC Suite Temp Files</i>
<input type="checkbox"/>	4	<input checked="" type="checkbox"/>	Windows	Execution Control	<i>Allow ABC Executions</i>
<input type="checkbox"/>	5	<input type="checkbox"/>	Windows	Expert	Report execution of potential risk files identified by yar
<input type="checkbox"/>	6	<input type="checkbox"/>	Windows	Expert	Block execution of malicious files identified by yara
<input type="checkbox"/>	7	<input checked="" type="checkbox"/>	Windows	Expert	Examine powershell script contents

Custom Rule Types and Examples

The Rule Type menu on the Add/Edit Custom Rule page provides the following options:

- **File Integrity Control** – Protects specified folders or files from being modified.
- **Trusted Path** – Defines folders or files for which file execution is always allowed.
- **Execution Control** – Controls behavior when an attempt is made to execute a file matching the rule.
- **File Creation Control** – Controls behavior when an attempt is made to write a file matching the rule.
- **Performance Optimization** – Specifies folders or files to avoid tracking (execution will still be monitored).
- **Advanced** – Provides a larger selection of menu options for controlling file execution, creation, and/or tracking.
- **Expert** – Provides the greatest selection of rule options via a checkbox-based interface. You can select one or more of the internal actions underlying the other rules types. Expert Rules do not have the error-checking that other Custom Rule types have, and it is possible to create unexpected (and unwanted) outcomes without feedback during rule creation. These rules are intended for use by Carbon Black Support or Services representatives, or customers working with them. See [Chapter 17, “Expert Rules,”](#) for more details.

The Custom Rules table includes several rules marked as *[Sample]* – these rules are disabled by default. For example, *[Sample] Developer - Visual Studio Ignore Intermediate Files* is a Performance Optimization rule that instructs App Control to ignore certain intermediate files typical of many build environments. In the Custom Rules table, you can click the View Details button next to any of these samples to examine the types of field choices that might be applied to accomplish similar results.

The sections below provide general examples of some of the different rule types.

File Integrity Control

Write Action Options: Block, Report

Execute Action: Does not apply to this rule type (not shown)

Users: Applies to all users (fixed value for this rule type, not shown)

File Integrity Control rules allow you to control modifications to a specific folder (or file) or folders (or files) matching your specification. You can write-protect the folder(s) by choosing Block as the Write Action, or you can monitor (but not block) changes by choosing Report as the Write Action.

Definition

Platform: Windows

Rule Type: File Integrity Control

Write Action: Block Use Policy Specific Notifier

Path or File: Specific Path...

Process Exclusion:

Rule Applies To

Policies: All Current and Future policies
 Selected policies

For example, perhaps you use an application called ScheduleCreator to generate schedules for everyone at your company and put the results in a **Schedule** folder in the **My Documents** folder on each user's computer. Assume that the ScheduleCreator executable is called **makesched.exe**. You want to be able to generate the schedule for each user, but you want to make sure nobody can change the schedules in the designated location once generated. You could choose **File Integrity Control** as the rule type and leave **Block** as the Write Action. Then you could enter **<MyDocuments>\Schedule** as your Path or File. Note that **<MyDocuments>** is a macro that maps to the My Documents folder for each user on computers running the agent. Finally, in the Process Exclusion box, you could enter ***\makesched.exe** so that this process will be allowed to write to the path in the rule. Use of a macro in the Process Exclusion box could further restrict the allowable process to one run from a specific location, such as **<ProgramFiles>\Schedule Maker\makesched.exe**.

Add Custom Rule

General

Rule Name: Block Unauthorized Writes to Schedule Folder

Description: Block any attempts to modify schedules in the Schedule folder unless by Schedule Maker application

Status: Enabled Disabled

Definition

Platform: Windows

Rule Type: File Integrity Control

Write Action: Block Use Policy Specific Notifier

Path or File: Specific Path...
<MyDocuments>\Schedule\

Process Exclusion: <ProgramFiles>\Schedule Maker\makesched.exe

Rule Applies To

Policies: All Current and Future policies
 Selected policies

Save & Exit Save Cancel

Trusted Paths

Execute Action: Allow, Allow and Promote, Promote

Users: Applies to all users (fixed value for this rule type, not shown)

One use of custom rules is designation of a trusted path. You can designate a network location as a trusted path and place installers there so that computers in certain policies or all policies can execute them.

A trusted path is an access method, *not* a global approval method. It allows execution of files in a specific location without globally approving files generated by the executable.

Any files in a trusted path must be executed in the specified location; the destination of the files *resulting* from an execution can be another computer (i.e., the computer accessing the executable via a trusted path). Computers with access to files on the trusted path cannot execute an installation package by copying it to their own machine and executing it there.

The local state of any files written by a file in a trusted path depends upon the Execute Action command used:

- If the Execute Action is Allow, an installer is allowed to write files but those files are not locally approved by the action. In this case, if the new files have not been seen by the App Control Server before, they are added to the File Catalog tab of the Files page with a status of Unapproved.
- If the Execute Action is Allow and Promote, the installer can write files and those files will be locally approved (unless already banned).

Important

- Any user who is able to write executables or scripts into the trusted path can make any application available to any computer that (a) has access to that location and (b) permits executions from remote drives. Before you enable a trusted path, check the platform's security settings for that location to ensure that it is properly protected.
- In the console, one way to help protect a Trusted Path is to create a user-specific File Integrity Control or File Creation Control rule for the same path. If you rank that new rule higher than the Trusted Path rule, you can control writes to the path while still allowing its use as a software distribution location.
- Whether or not you create a companion rule as suggested above, be careful about the rank of Trusted Path rules. Because new rules are ranked first by default, you will need to move the rule down in the ranking if you want the internal rules for blocking undesirable activity take precedence.

To create a trusted path for installers, follow the instructions in [“Creating a Custom Rule”](#) on page 396, choosing **Trusted Path** as the Rule Type. Note that when you choose Trusted Path, other fields on the page change to reflect your choice. The Execute Action menu shows **Allow**, meaning that files matching this rule will be allowed to execute.

Definition

Platform:

Rule Type:

Execute Action: ⓘ

Path or File: ⓘ

Process: ⓘ

Rule Applies To

Policies: All Current and Future policies
 Selected policies

For example, you might use an application called FileDistributor to distribute your company software via some distribution server. Assume that the FileDistributor application is actually an executable called **filedist.exe**, and that your company's software is deployed from a distribution server located at **\\FILE2DEPLOY\Apps**. You could choose **Trusted Path** as the rule type and enter **\\FILES2DEPLOY\Apps*** as your Path or File.

If you leave the Process field for this rule set to **Any Process**, any process on a client affected by the rule can run applications and installers from that location. To reduce the security gaps in your custom rule, you might want to limit the right to execute files in this directory to FileDistributor itself, such that *only* FileDistributor can install applications from the named directory. By making the Process ***filedist.exe**, you create just such a restriction. You can be even more specific by using a macro to identify the file location; for

example, **<ProgramFiles>\FileDistributor\filedist.exe**. A user *manually* trying to run those same files will be blocked.

Add Custom Rule

General

Rule Name:

Description:

Status: Enabled Disabled

Definition

Platform:

Rule Type:

Execute Action: ⓘ

Path or File: ⓘ
 ✓

Process: ⓘ
 ✓

Rule Applies To

Policies: All Current and Future policies
 Selected policies

Save & Exit
 Save
 Cancel

You can further limit trusted paths and any other custom rules to computers in one or more specific policies, using the “Rule applies to” buttons. By combining all of these fields, you have the opportunity to define a rule that allows you to accomplish necessary operations while exposing your systems to as little security risk as you can.

Execution Control

Execute Action Options – Allow, Block, Allow and Promote, Promote, Prompt, Report

Write Action – Does not apply to this type (not shown)

Execution Control rules are exactly what they sound like. They allow you to create a rule that responds in the way you choose when a file matching the rule attempts to execute. They do not have any effect on attempts to write (create, modify, or delete) matching files.

Execution Control rules are similar to Trusted Path rules, except that Execution Control rules allow you to specify a user or group and they offer more Execute Action options.

Definition

Platform: Windows

Rule Type: Execution Control

Execute Action: Allow

Path or File: Specific Path... i

Process: Any Process i

User or Group: Any User i

Rule Applies To

Policies: All Current and Future policies
 Selected policies

For example, perhaps your developers use a tool called MyDevTool to develop and compile DLLs. The MyDevTool application is set up to run the DLLs it creates. You might create a rule that prevents this execution from being blocked.

Since the files created by MyDevTool are all DLLs, you can use `*.dll` as your Path or File. If you were certain of the location of these files, you could further specify the path, but for this example we will leave the location open.

If you leave the Process field for this rule set to **Any Process**, any process on a client affected by the rule can run any DLL. To make this rule more secure, you might want to limit the right to execute files in this directory to MyDevTool application itself. To do this, you could use a macro to help specify the exact location of the tool, for example `<ProgramFiles>\ToolCo\MyDevTool\runtool.exe`.

If you have defined Active Directory groups, you might choose to further restrict the ability to run these DLLs to the group known to have permission to use this tool. To do this, you could choose **Specify User or Group...** on the User or Group menu and then enter the AD Group name for the permitted group, **Developers**, for example.

Now you have a rule that will allow execution of DLL files in any location as long as they are being executed by user in the Developers group using MyDevTool.

Add Custom Rule

General

Rule Name:

Description:

Status: Enabled Disabled

Definition

Platform:

Rule Type:

Execute Action: ⓘ

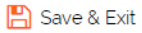

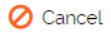
Path or File: ⓘ
 ⓘ

Process: ⓘ
 ⓘ

User or Group: ⓘ
 ⓘ

Rule Applies To

Policies: All Current and Future policies
 Selected policies

 Save & Exit
 Save
 Cancel

File Creation Control

Write Action Options – Ignore, Block, Approve, Approve as installer, Prompt, Allow

Execute Action – Does not apply to this rule type (not shown)

File Creation Control rules allow you to control what happens when there as an attempt to write (create) a file that matches the rule. They do not have affect file execution attempts.

Like File Integrity Control rules, File Creation Control rules allow you to Block writes. However, File Creation Control rules allow you to specify a user or group and they offer more Write Action options for cases in which you are not blocking file writes.

Definition

Platform: Windows ▼

Rule Type: File Creation Control ▼

Write Action: Block ▼ i Use Policy Specific Notifier

Path or File: Specific Path... ▼ i
 ▼

Process: Any Process ▼ i

User or Group: Any User ▼ i

Rule Applies To

Policies: All Current and Future policies
 Selected policies

Performance Optimization

Write Action – Ignore (value fixed for this rule type, not shown)

Execute Action – Does not apply to this rule type (not shown)

Users – Any User (value fixed for this rule type, not shown)

Unless instructed otherwise, the App Control Server keeps track of any files written to a computer running its agent. Normally, this is useful for monitoring purposes. However, there are cases in which a particular process writes many files to the same directory as part of its normal operation, and monitoring these write operations uses system and network resources unnecessarily while providing no important information. In cases such as these, you might choose to create a Performance Optimization custom rule for the uninteresting directory.

To create a rule that eliminates tracking for certain files, follow the instructions in [“Creating a Custom Rule”](#) on page 396 and choose **Performance Optimization** as the Rule Type. When you choose Performance Optimization, some other fields on the page change to reflect your choice. Note that although not shown, the Write Action for this rule is **Ignore**, meaning that writing of files matching this rule will not be tracked by the App Control Server.

Definition

Platform: Windows ▼

Rule Type: Performance Optimization ▼

Path or File: Specific Path... ▼ i
 ▼

Process: Any Process ▼ i

Rule Applies To

Policies: All Current and Future policies
 Selected policies

For example, perhaps an application called MyVirusGuard is writing a lot of temporary files to `c:\temp2\`.

You could create a Performance Optimization rule that specifies **c:\temp2*** in the Path or File field. The App Control Server would not track any files written to, modified in, or deleted from that location by anyone. This reduces processing and information collection, but it also means that you are not tracking *any* files being written to that directory.

If MyVirusGuard uses the executable MVGuard.exe for its operations, including writing files, you could add ***\MVGuard.exe** to the rule as the Process, which lets MyVirusGuard write to the directory without tracking. The server continues to track files written to c:\temp2\ by any other process. Specifying the process allows you to accomplish the task you wanted while maintaining as much protection as possible. Note also that because you used the asterisk wildcard and a slash before the process name in the Process field, it does not matter where you installed MVGuard.exe – it is always allowed to write to the designated directory without tracking.

Add Custom Rule

General

Rule Name:

Description:

Status: Enabled Disabled

Definition

Platform:

Rule Type:

Path or File: i
 x

Process: i
 x

Rule Applies To

Policies: All Current and Future policies
 Selected policies

Since the (hidden) Execute Action for a Performance Optimization rule is **Default**, any *executions* in c:\temp2 are still tracked and executions are still blocked if other rules would block them – only file *writing* has been allowed and not tracked, and only if attempted by the process you specified.

Pairing Ignore and Block Rules



In one previous example, a Process Exclusion was used to allow a specific process to write to the location normally blocked by the rule. You also can create an exclusion to a rule by pairing it with a second rule and ranking the exclusion rule before the main rule.

Perhaps a program called Super App has a log file called **superapp.log** in a **logs** subfolder. You might not want to create an exception for a process but instead only allow

the files to be written in the particular subfolder while protecting the rest of the application folder. To do this, you could create two rules with the following characteristics:

- **Ignore Rule** – Create a Performance Optimization rule to ignore and allow writes (the automatic action choice) to **<ProgramFiles>\superapp\logs***
- **Block Rule** – Create a File Integrity Control rule with a Write Action of Block for the path **<ProgramFiles>\superapp***, and rank that rule lower than the Allow rule.

With the Performance Optimization rule ranked before the Block rule, App Control first checks to see whether a modification attempt matches the exception, and if it does, the Block rule is not evaluated.

Rank ▲	Status	Rule Type	Name	Action	Path
↑ ↓ 1		Performance Optimization	Allow Super App Log Writes	Ignore writes	c:\temp\superapp\logs*
↑ ↓ 2		File Integrity Control	Protect Super App Folder	Block writes	c:\temp\superapp*

Chapter 15

Registry Rules

This chapter describes Registry Rules, which control what happens when there is an attempt to make changes in the Windows Registry at locations that match paths you specify. If you choose, you can limit enforcement of the rules to specified users and/or processes.

Unified Management

Registry Rules can be centrally managed for multiple servers through the Unified Management feature. This is described in [“Unified Management of Rules”](#) on page 791.

Platform Note: Registry rules affect only computers running Windows operating systems.

Sections

Topic	Page
Overview	450
Specifying the Notifier for Registry Rules	451
Creating Registry Rules	452
Registry Rule Fields	455
Specifying Registry Paths	459
Specifying Processes in Registry Rules	460
Rule Ranking	465
Disabling or Deleting Registry Rules	466
Sample Registry Rules	467
Autostart Rules	469

Overview

Registry rules enable you to block, report, allow, or prompt the user for a choice when there are attempts to write to Windows Registry locations matching paths you specify. Creation, modification and deletion of keys or values all count as “writes”.

You can view a list of related events, including any blocks caused by registry rules, by going to the Events page and choosing Registry on the Saved Views menu

Note

For computers in Visibility mode policies, registry rules that would block writing or prompt users for a decision are treated as report-only rules, and therefore will not block or prompt.

Rule Scope

You can create registry rules that apply to all users and all processes that try to make a registry change on any Windows computer. You also can create a more focused rule scope by specifying one or more of the following criteria:

- **Process-specific** – You can make a rule apply only when *certain processes* attempt to write to the specified location.
- **User- or group-specific** – You can make the rule apply only to a particular *user or group of users*.
- **Policy-specific** – You can choose to limit a rule to *computers in specified policies*.
- **Server-specific** – If you have Unified Management enabled, you can choose to limit a rule to *computers reporting to specified servers* in the management group. This is described in [“Unified Management of Rules”](#) on page 791.
- **Rule order** – Registry rules are evaluated in order of *Rank*, a column that is displayed by default on the Registry Rules table. The rule ranked ‘1’ has the highest rank, ‘2’ is next, and so on. You can change the order of rules. For example, you can create a rule that applies when *a particular user* attempts to access a specified Registry path, and put that before a rule that applies when *any other user* attempts to access that path.
- **Conditional Macros** – You can use certain macros to restrict the conditions under which specific parameters in rules are applied. Only agents meeting the “test” described in the macro will attempt to match the parameter prefixed with the macro. Most of these macros are *OnlyIf* macros with different arguments, such as `<OnlyIf:OSVersionIs:10.6.8>` and `<OnlyIf:HostName:*SMITH-1*>`.

Important

Registry rules generally should be as narrowly targeted as possible to avoid unintended effects.

Sample Rules

A new installation of a App Control Server is pre-configured with built-in registry rules, disabled by default, which you can view by clicking the Registry tab on the Software Rules page. Some of these are samples that you may either enable as is or use as a guide to creating your own rules. The Autostart rule, which also is disabled by default, protects a long list of registry locations potentially affected on startup. See the section “[Sample Registry Rules](#)” on page 467 for an example of how a rule can be configured.

Exporting and Importing Registry Rules

You can export registry rules from one server and import them to another. There are buttons for this purpose on the Registry Rules page. See “[Exporting and Importing Rules](#)” on page 431 in the Custom Rules chapter for more information.

Specifying the Notifier for Registry Rules

The App Control Agent provides *notifiers* that can be displayed when a rule blocks an action or prompts the user for a decision to allow or block an action. For each registry rule, you can choose one of two sources for the notifier:

- **Use Policy Specific Notifier** – Each Policy includes an Advanced Setting, “Enforce registry rules”, which is always on. This setting has a Notifier field in which you can specify the notifier that appears on agent computers when a registry rule blocks an action.

If you choose Use Policy Specific Notifier for a rule, it is possible that the policy specifies <none> as the Notifier for *Enforce registry rules*. In this case, a notifier will not be shown, even for a Prompt rule. Unless you are certain that you never want to prompt the user for a response to a rule, choosing <none> for the custom rule notifier in a policy is not recommended. See “[Advanced Settings](#)” on page 187 for more information.

- **Custom Notifier** – If you do not choose the policy-specific notifier, you can choose (or create) a different notifier for a registry rule. The choices appear on a menu on the Add/Edit Registry Rule page.

When you choose Block as the rule action, you can choose <none> on the Custom Write Notifier menu since it is possible you want the rule to block actions without notification. A Prompt rule requires a user choice, so when you choose Prompt as the rule action, the Custom Notifier menu does not include <none>.

Note

If you are using Unified Management and create a Registry Rule that applies to more than one server, client servers will use default notifiers, even if a custom notifier is specified on the management server.

See [Table 60](#) below for the registry rule notifier settings. See [Chapter 20, “Endpoint Notifiers and Approval Requests,”](#) for more on notifiers.

Creating Registry Rules

In addition to providing a name, to create a registry rule, you need to provide the information shown in bold in the left column of the table below and enter it in the Add Registry Rule page in the locations on the right:

General Description	Field on Add/Edit Registry Rule Page
If this/these source process(es) ...	Process
...and/or this/these user(s) ...	User or Group
... attempt to modify the Windows Registry at this/these location(s) ...	Registry Path
... on computers in this/these policy(ies) ...	Rule applies to/Policies:
... on computers reporting to this/these App Control server(s) ...	Rule applies to/Servers (if Unified Management is enabled)
.. then this action should be taken.	Write Action

* Additional actions and other options are available in Expert Mode. See [Chapter 17, "Expert Rules,"](#) for more details.

For each of these fields, there could be multiple matching items, or the rule could specify all items in that class (for example, the rule applies to all users, or all policies, or all source processes).

To add (create) a registry rule:

1. In the console menu, choose **Rules > Software Rules**. The Software Rules page appears.
2. Click **Registry**, either on its tab or in the menu to the left of the page. The Registry Rules page appears:

Software Rules

Automatically apply

<input type="checkbox"/> Select 6	Rank	Status	Platform	Name	Action
<input type="checkbox"/>	1	<input type="radio"/>	Windows	[Sample] Block Suspicious System Changes	Block, Finish Rule Group
<input type="checkbox"/>	2	<input type="radio"/>	Windows	[Sample] Prompt on Suspicious changes to Startup Files	Prompt, Finish Rule Group
<input type="checkbox"/>	3	<input type="radio"/>	Windows	[Sample] Report Typical Changes to Startup Files	Report
<input type="checkbox"/>	4	<input type="radio"/>	Windows	[Sample] Report Changes to Trusted Zones	Report
<input type="checkbox"/>	5	<input type="radio"/>	Windows	[Sample] Report Changes to Home Page or Search Settings	Report
<input type="checkbox"/>	6	<input type="radio"/>	Windows	Autostart Rules	Report

- To create a new rule, click the **Add Registry Rule** button. The Add Registry Rule page appears.

Add Registry Rule

General

Name:

Description:

Status: Enabled Disabled

Definition

Expert Mode: On Off

Platform: Windows

Write Action: Block Use Policy Specific Notifier

Registry Path: i

Source Process: Any Process i

User or Group: Any User i

Rule Applies To

Policies: All Current and Future policies
 Selected policies

- In the Name field, enter the name you want to appear on the list of rules.
- If you want to add other comments about the rule, such as its purpose or its relationship to other rules, you may provide an optional Description.
- Enter the remaining information you want for this rule (see [Table 60, “Registry Rule Fields,”](#) on page 455) and then click the **Save** button if you need to remain on the page or **Save & Exit** if to go to the Registry Rules table. By default, a new registry rule is **Disabled** and ranked #1, listed at the *top* of the Registry Rules table.
- Before you enable a rule, change its rank unless you want it to take precedence over (and perhaps preempt) all other rules. You can change rank using the arrows in the Rank column or drag-and-drop (if the table is sorted by rank), or you can click on the rank number and enter a new rank in the dialog box. See [“Rule Ranking”](#) on page 465 for more details.
- When you are satisfied with the rank and want to enable the rule, click the toggle switch in the Status column of the Registry Rules table. The button in the switch moves to the right and the background turns from white to green.

Editing a Registry Rule

Editing a Registry Rule is very similar to creating one. If you have permission to edit the rule, you can edit any field, including the rule name.

To edit a registry rule:

1. In the console menu, choose **Rules > Software Rules**. The Software Rules page appears.
2. Click **Registry**, either on its tab or in the menu to the left of the page. The Registry Rules page appears.
3. On the Software Rules page, click the View Details button for the rule you want to edit.
4. On the Edit Registry Rule page for that rule, make your changes. [Table 60](#) describes the custom rule fields that you might choose to edit.
5. Click either **Save** (to remain on the edit page) or **Save and Exit** (to return to the table page).
6. If you clicked Save, a confirmation message appears on the page. Click the message to clear it from the page. If an error occurs, review the error message and correct the conditions that caused the error before saving again.

Unified Management

If you are using Unified Management, and you edit a unified rule shared with other servers, a “wizard” shows the progress as the edited rule is saved on each server.

See [“Unified Management of Rules”](#) on page 791 for more details.

Copying a Registry Rule

There is a *Copy this rule* command on the right menu on the Edit Rules page for Custom, Registry, and Memory Rules. This is for making copies of the rule on the same server. You might do this so that you can customize a sample rule while preserving the original settings as a template. It also allows you to make slightly different rules for different policies without having to manually provide all of the settings for each one.

To copy a registry rule:

1. On the Registry Rules table page, click the View Details button to open the details page for the rule you want to copy.
2. On the details page, click **Copy this rule...** in the Actions menu on the right side of the page. This opens a dialog box. By default, the copy is named using the original rule name plus “(copy)”.
3. In the dialog box, change the rule name if you want something more descriptive than what is there.
4. If you want the new rule enabled immediately, check the **Enable copied rule** box. Otherwise, uncheck the box.
5. Click **OK**. The copied rule is created and its details page replaces the details page for the original rule.
6. Make any changes you would like to make in the new (copied) rule and **Save** or **Save & Exit**.

Registry Rule Fields

Table 60 shows the fields available on the Add/Edit Registry Rule page. Column headings on the rule table page are shown when they differ from the Add/Edit page.

Table 60: Registry Rule Fields

Field	Description
Name	Name by which this rule is identified in the Registry Rules table. (Required)
Description	Optional information about the registry rule. This can be any text you choose to enter.
Rank (Table only)	The rank of this rule in order of evaluation. The rule ranked '1' in the table is evaluated before the rule ranked '2', and so on.
Status	Radio buttons that make this rule Enabled or Disabled. This allows you to create a rule that you use only at certain times, or to temporarily disable the rule without losing the information used to create it.
Expert Mode	Radio buttons that make turn Expert Mode on and off (the default). Expert mode provides more options than standard mode but does not have all of the error-checking that other rule types have, so it is possible to create unexpected (and unwanted) outcomes without being warned during rule creation. These rules are intended for use by Carbon Black Support or Services representatives, or customers working with them. See Chapter 17, "Expert Rules," for more details.
Platform	Platform for which this rule is effective. This is a read-only field and the value is always Windows. Registry rules do not have any impact on non-Windows platforms.
Write Action (Add/Edit page only)	The action to take when there is a write attempt matching this rule. See Table 61 for the action options. For all Windows platforms except Windows Server 2003 64-bit, write rules also control changes to registry permissions.
Action (Rule table only)	The type of action the rule takes. The possible values include all of those shown for Write Action plus other actions made available in Advanced and Expert rules.
Operation (Rule table only)	The type of operation the Write Action rule affects. Standard Write Rules for the registry affect the following operations: Create Key, Rename Key, Delete Key, Change value, Delete value, Set Security, and Open Key with Write Access. Rules created or edited in Expert Mode might not include all of these operations and also might include other operations.

Field	Description
Action (Legacy) (Rule table only)	<p>This column shows actions and operations for the rule as shown in the Action column in pre-8.1.6 versions, or it shows “Expert Action(s)” in cases where expert rule information was not previously shown.</p> <p>Note: This field is present strictly for continuity with older versions – you should use the separate Action and Operation columns for the most accurate description of the rule.</p>
Use Policy Specific Notifier	<p>If you choose Block or Prompt as the Write Action, this checkbox appears to the right of the Write Action choice. If you check the box, the notifier that appears when a registry rule blocks an action is the notifier specified for the Enable Registry Rules setting in the policy for the computer on which the action was blocked. If not checked, you can choose a custom notifier from the Custom Write Notifier menu.</p>
Custom Write Notifier	<p>If you choose Block or Prompt as the write action, and you do not check the Use Policy Specific Notifier box, this menu appears.</p> <p>If you choose Block as the write action, you can choose any notifier from the menu. The menu also includes a <none> option so that you can disable the notifier for this rule.</p> <p>If you choose Prompt as the write action, you can choose any notifier on the menu. Prompt rules must display a notifier, so there is no <none> choice in this case.</p> <p>Note: If you use Unified Management to create a rule that applies to more than one server, client servers will use default notifiers, even if a custom notifier is specified on the management server.</p>
Registry Path (Path in table)	<p>Registry path to which this rule applies. See “Specifying Registry Paths” on page 459 for details on your options for specifying the path.</p>
Source Process (Process in table)	<p>This field allows you to limit the rule so that it is applied only when certain processes attempt to execute or write files matching the path specification. See “Specifying Processes in Registry Rules” on page 460 for details on specifying a process and Table 62 for process menu options.</p>
User or Group	<p>This field allows you to specify users or groups to which this rule applies. See “Specifying Users or Groups” on page 464 for details on specifying users or groups.</p>
Rule Applies To: Servers (Add/Edit page only)	<p>Radio buttons allow you to apply the rule to the current server, All Servers or Selected Servers. If you choose Selected Servers, a list that includes the current server and of all App Control servers managed by this server appears, each with a checkbox. In addition, policies for the servers you include appear in the Selected policies list.</p> <p>Unified Management: This field appears only if Unified Management is configured on the server you are logged into.</p>

Field	Description
Unified Server Source (Table only)	If this is a unified rule, the name of the unified management server that created or copied the rule.
Rule Applies To: Policies (Policy in the table)	Radio buttons allow you to apply the rule to All policies or Selected policies . If you choose Selected policies , a list of all policies available on your App Control Server appears, each with a checkbox. You can check as many policies as you choose. Unified Management: If Unified Management is configured on the server you are logged into, and if you have chosen to apply the rule to additional servers, policies for all selected servers appear in this list.
Is Global (Table only)	Indicates whether the rule applies to all policies (Yes) or only selected policies (No).
Rule Applies To: Override Permissions (Add/Edit page only)	Radio buttons allow you to specify whether administrators on other servers can modify rules sent via Unified Management on their own server. The options are No Override , Partial Override (allows changing rank) and Full Override (allows editing and changing rank). Unified Management: This field appears only if Unified Management is configured on the server you are logged into and this rule is applied to more than the current server in the Rule Applies To:Servers field.
History	For existing rules, a History panel on the Edit Rule page appears showing some or all of the following fields. In addition, these fields can be added as columns on the rules table page. <ul style="list-style-type: none"> • Created By – If the rule was created on this server, the user who created it. Rules created during server installation or upgrades show “System” in this field. • Date Created – If the rule was created on this server, when it was created. • Last Modified By – If the rule has been modified since creation or import, the user who modified it. • Date Modified --If the rule has been modified since creation or import, when it was modified. • CL Version – Rules created after server installation also show the CL (config list) number that first contained the rule so that you can compare an agent CL number to determine whether the agent has received the rule. • Imported – (In the table only) indicates whether the rule was imported (Yes/No). • Imported By – If the rule was imported to this server, the user who imported it. • Imported Date – If the rule was imported to this server, when it was imported.

Specifying a Write Action

The Write Action in a registry rule is the action to take when there is a registry write attempt matching this rule. [Table 61](#) shows the options. Write action includes creation, deletion and modification of registry keys on all platforms. It also includes changes to registry permissions on all Windows platforms except Windows Server 2003 64-bit.

Table 61: Write Action Menu Options

Option	Description
Block	<p>Prevent creation, deletion and modification of registry keys and values at locations matching this rule.</p> <p>When Block is chosen, the Use Policy Specific Notifier checkbox and a Custom Write Notifier menu appear. These allow you to specify the notifier, if any, that appears when the rule blocks an action. See Table 60 for more details.</p>
Prompt	<p>Present a notifier dialog to the computer user when an attempt to modify the registry is made at this location. The dialog choices are Block or Allow. Once the user responds to the dialog, the choice applies anytime the same process matches the same rule on the same computer with the same user – the user will not be prompted again in this case.</p> <p>When Prompt is chosen, the Use Policy Specific Notifier checkbox and a Custom Notifier menu appear. These allow you to specify the notifier that appears to prompt the user. See Table 60 for more details.</p>
Report	Do not block modifications at this registry path but report them as events.
Allow	<p>Allow creation, deletion and modification of registry keys and values at locations matching this rule. This is the default behavior if there is no rule for a path.</p> <p>Use of Allow gives you a way to create an exception to a more general rule that blocks at a particular location. For example, if you create a rule that blocks all writes to</p> <pre>*\Software\MyApp*</pre> <p>you could create an exception by creating a higher ranking rule that allows writes to</p> <pre>*\Software\MyApp\SpecialKey</pre>

Specifying Registry Paths

The Registry Path specifies the locations in the Windows Registry to which a rule applies.

Definition

Expert Mode: On Off

Platform: Windows

Write Action: Block Use Policy Specific Notifier

Registry Path:

Source Process: Any Process

User or Group: Any User

All registry paths must begin with one of the following strings:

- HKLM\
- HKCU\
- HKLM-SoftwareX86\
- HKLM-SoftwareX64\
- HKCU-SoftwareX86\
- HKCU-SoftwareX64\
- *\

Notes

- You cannot use macros in the Registry Path.
- If you enter a path that uses a key that is actually a link to other keys, the rule will not work. For example, a rule that uses a path containing *CurrentControlSet* will fail to work. You might consider using wildcards in place of the linked item (for example, *ControlSet** in the previous case).

Using Wildcards

You can use wildcards (“*” for zero or more characters, “?” for one character) in the Registry Path. You can use wildcards to specify partial paths or multiple paths in the registry. The number of wildcards in a path is not restricted.

You can use wildcards to skip a level and make a rule apply to values (or sub-keys) of a sub-key, even if you don’t know their names. For example:

```
*\myapp\*\*
```

applies the rule only to keys or values *below a sub-key of myapp*, such as

```
HKLM\myapp\apprunner\4.0
```

but it does not apply to sub-keys or values in *myapp* itself, such as

```
HKLM\myapp\sharedfiles
```

Caution

When you use wildcards, do not to create a rule that is so broad that it will interfere with activity that is required for legitimate use by an application or the operating system. Do not use the asterisk wildcard by itself in the Registry Path field, especially with rules that block all writes, unless you are certain it will not interfere with necessary operations on the agent computer. Registry rules may seriously impact the performance of an application or system.

Specifying Keys or Values

If a path ends with a "\", it matches only the key at that path. If a path ends in "*", the rule applies to all keys, sub-keys and values underneath that path.

If a path ends without a slash or wildcard, it applies only to a value (not a key) matching the path. For example:

```
HKLM\SOFTWARE\FileReader\9.0\ViewOutput
```

would match a *value* named "ViewOutput" but not a *key* named "ViewOutput"

You can add more than one path to a Registry Rule. See [“Entering Multiple Paths or Processes”](#) on page 463 for details. In the Registry Rule table, rules with more than one path show the first path in the Registry Path field followed by **(multiple)**.

Specifying Processes in Registry Rules

The Process field on the Add/Edit Registry Rule page allows you to fine-tune the rule according to the process – that is, the running file – attempting to modify the registry.

Definition

Expert Mode: On Off

Platform: Windows

Write Action: Block Use Policy Specific Notifier

Registry Path:

Source Process:

User or Group: Any User

You can make the rule effective for all processes, certain types of processes, specific processes, or all processes except the one(s) you name. [Table 62](#) shows the Process options.

Table 62: Process Menu Options

Menu Option	Description
Any Process	Applies the rule to any process that attempts to write to the registry.
Any Promoted Process	Applies the rule to any process that is promoted at the time the rule is evaluated. A promoted process is any approved process that is marked as an installer, or has been promoted as a consequence of a custom rule, or is an approved process launched by a promoted process.
Any System Process	Applies the rule to every process that is running under the security context of the Local System user. This choice has the same effect as choosing Local System in the User or Group menu, but may be more efficient.
Specific Process...	Opens a text box below the menu; you can enter the names of processes you want controlled by this rule. See “Specifying Processes in Registry Rules” on page 460 for the guidelines and requirements for specifying a process.
Any Process Except...	Opens a text box below the menu, in which you can enter processes you <i>do not</i> want controlled by this rule. See “Specifying Processes in Registry Rules” on page 460 for the guidelines and requirements for specifying a process. Note: If you specify a User or Group and also choose Any Process Except from the process menu, the rule is enforced <i>unless the exception process is being executed by the user or group</i> .

When you choose a Process option that requires entry of a path (either *Specific Process...* or *Any Process Except ...*), you have several options for defining paths:

- **Specify a specific process or a directory** – You can enter a process specification that exactly identifies a process by path and name so that only that file matches the rule. You also can enter a specification that identifies a directory, which matches all processes in that directory and its subdirectories.
- **Specify a local drive or UNC path** – You can identify a local process by using a local drive name, such as *C:\folder1\subfolder\application.exe*. You also can enter a remote process by using a UNC path, such as *\\computername\dir\application.exe*. Mapped drives in a path or process specification are not recognized.
- **Use wildcards** – You can use wildcards ('?' for any one character and '*' for zero or more characters) to expand the scope of a process specification or help you match a file or folder whose exact location you don't know. Wildcards may be used at the beginning, end or middle of a path.
- **Use macros** – You can use special App Control macros to identify certain well-known folders in the Microsoft Windows environment, even if you don't know their exact location on all agent computers.
- **Specify multiple process paths** – You can add more than one process definition per rule.

Specifying Processes or Directories

You can choose to enter a directory or a specific file as your process path. When you specify a directory, you are instructing the rule to apply when any process in that directory or in any of its subdirectories attempts to write to the registry location specified (unless there are higher-ranked rules that match the current process).

To indicate that a Process definition is a directory, you must end it with a backslash (\) or a backslash and asterisk (*). If you do not include the backslash, the rule will attempt to match a *file* by the name you provided, not a directory. For example, either of the following correctly identifies “subfolder2” as a directory in a process definition:

```
c:\folder1\subfolder2\  
c:\folder1\subfolder2\*
```

However, the following is *not* recognized as a directory:

```
c:\folder1\subfolder2
```

If you use path macros in a process definition, the expanded macro is treated as a directory, even if you don't follow it with a backslash. See [Using Macros](#).

Using Wildcards

You can use wildcard characters in the Process field. Asterisk (*) indicates zero or more characters and question mark (?) indicates one character. You can also use them to specify processes that appear in different locations on different computers (although macros might be a more effective way to accomplish this – see [Using Macros](#)).

The number of wildcards in a process specification is not restricted. For example, you could define a path as:

```
*\Win*\folder?\
```

Automatic Process Path Conversions

The Process field undergoes automatic path conversions if it contains certain symbols:

- A process path that ends with a slash has the “*” wildcard added at the end of the path.
- A process path with no slash or drive letter has “*\” added at the beginning of the path.
- Drive letters may be used in a path as long as they are for local fixed volumes. Mapped drive letters should not be used because there is no guarantee that the mapping exists on all computers.
- The string “*:\” applies to all attached storage volumes except for floppy disks and CD-ROMs.

Specifying Devices in Process Path

You can specify that a rule applies when writes are attempted by processes running from some or all devices on the agent computer by including `\device\` in the path. For example:

- `\device*\` specifies all devices.
- `\device\harddisk*\` specifies attached storage volumes except for floppy disks and CD-ROMs.
- `\device\cdrom*\` specifies CD-ROM devices.

Using Macros

You can use certain macros in the Process field of a Registry Rule. You can see a menu of macros by typing the left angle bracket (<) character in the Process field. There are two types of macro supported in Registry Rule processes:

- **Path macros** – These are a subset of the well-known folders in the Microsoft Windows environment, and they always identify a location rather than a specific file. A path macro can be used only at the *beginning* of a Path or File specification in a rule (i.e., with no other text before it in the string).
- **Registry macros** – These are macros that specify strings in the Windows Registry. A registry macro can be used anywhere in the Path or File specification.

Macros can be an effective way to define a rule that works on all Windows computers even when the files you want to affect are in different locations on different computers.

See “[Using Macros in Rules](#)” on page 410 of the Custom Rules chapter for a description of path and registry macros. These macros may be used in the Process field of a registry rule.

Notes

Macros may be used in the Process field of a Registry Rule but not in the Registry Path field.

Entering Multiple Paths or Processes



For both the Registry Path and the Process field in a rule, you can enter more than one string. For example, when you have entered the first Registry Path for this rule, click the Expand button to the right of the box.


Definition

Expert Mode: On Off

Platform: Windows

Write Action: Block Use Policy Specific Notifier

Registry Path: HKLM\software\microsoft\windows\currentversion  

Source Process: Specific Process... 

You can then add additional paths by typing them in the box and clicking **Add** after each one.

Definition

Expert Mode: On Off

Platform: Windows

Write Action: Block Use Policy Specific Notifier

Registry Path: HKLM\software\microsoft\windows\currentversion\exp
 HKLM\software\microsoft\windows\currentversion

Source Process: Specific Process...

User or Group: Any User

You can remove any path by clicking the Expand button, selecting the file or path in the list below the Registry Path box, and clicking the **Remove** button. Adding or removing items in the Process field works in a similar way.

If you enter multiple paths or processes for a rule, the Registry Rules page shows the first path and then **(multiple)** in the relevant column for this rule. Moving the mouse over the value shows a tooltip with the complete list of paths or processes for the rule.

Select	Rank	Status	Platform	Name	Path
<input type="checkbox"/>	1	<input type="radio"/>	Windows	[Sample] Block Suspicious System Changes	HKLM\software\microsoft\windows nt\currentversion\image file exec
<input checked="" type="checkbox"/>	2	<input type="radio"/>	Windows	[Sample] Prompt on Suspicious changes to Startup Files	*software\microsoft\windows\currentversion\policies\explorer\run\
<input checked="" type="checkbox"/>	3	<input type="radio"/>	Windows	[Sample] Report Typical Changes to Startup Files	*software\microsoft\windows\currentversion\policies\explorer\run\
<input checked="" type="checkbox"/>	4	<input type="radio"/>	Windows	[Sample] Report Changes to Trusted Zones	*s
<input checked="" type="checkbox"/>	5	<input type="radio"/>	Windows	[Sample] Report Changes to Home Page or Search Settings	*s
<input checked="" type="checkbox"/>	6	<input type="radio"/>	Windows	Autostart Rules	<A

Showing 6 out of 6 item(s)

Specifying Users or Groups

You can create a rule that applies only when specific users or users in specific groups attempt an action. The choices for User or Group on the Add/Edit Custom Rule page are:

- **Any Users** – applies the rule to all users.
- **Specific User or Group...** – opens a text box below the menu, into which you can enter AD users or groups in the format *userorgroupname@domain* or *domain\userorgroupname*

- The other menu choices are built-in Windows groups, such as **Authenticated Users** and **Local Administrators**.

Notes

- When running on Windows Vista and later, membership in pre-defined security groups like Administrators requires that the application run as an administrator. If a group definition is necessary for a rule, consider using security groups you have defined rather than the pre-defined groups
- There is a brief delay after a user logs in before group membership is established and group-based rules become effective. This delay may be longer if you have a large number of rules. If you need a rule to be effective as soon as possible after a user logs on, do not specify a user *group* in the rule. Rules that specify a *username* or *SID* are always active and won't be affected by this delay.

Rule Ranking

Registry rules have a “Rank” number and are evaluated from lowest number to highest number, beginning with the rule ranked ‘1’. By default, rules appear in their rank order, but you can re-sort the table by other columns if you choose. If the path of an action matches two different rules, the highest ranking rule (that is, the one with the lowest number), takes precedence and the lower-ranked (higher number) rule has no effect. There is one exception to this behavior – rules whose action is Report do not stop processing of lower ranked rules.

You can change the ranking of rules.

To change the rank of a registry rule:

1. On the Registry Rules page, make sure the Rank column is displayed.
2. If you want to be certain you can see the rules ranked immediately before and after the rule you are moving, sort the table by rank and remove any filters.
3. To change the rank of a rule, you have three options:
 - In any table that displays the rank column, you can click on the rank number and enter a new rank number in the dialog box.
 - If the table is sorted by rank and not filtered, arrows appear next to the rank, and you can click the up or down arrow button next to the rule to change its rank.
 - If the table is sorted by rank and not filtered, you can hold down the left mouse button with the cursor over the rule and drag the rule to a new location.

<input type="checkbox"/> Select 6		Rank	Status	Platform	Name	Action	Path
<input type="checkbox"/>		↑ ↓ 1	<input type="radio"/>	Windows	[Sample] Block Suspicious System Changes	Block, Finish Rule Group	HKLM
<input type="checkbox"/>		↑ ↓ 2	<input type="radio"/>	Windows	[Sample] Prompt on Suspicious changes to Startup Files	Prompt, Finish Rule Group	*\sof
<input type="checkbox"/>		↑ ↓ 3	<input type="radio"/>	Windows	[Sample] Report Typical Changes to Startup Files	Report	*\sof
<input type="checkbox"/>		↑ ↓ 4	<input type="radio"/>	Windows	[Sample] Report Changes to Trusted Zones	Report	*\sof
<input type="checkbox"/>		↑ ↓ 5	<input type="radio"/>	Windows	[Sample] Report Changes to Home Page or Search Settings	Report	*\sof
<input type="checkbox"/>		↑ ↓ 6	<input type="radio"/>	Windows	Autostart Rules	Report	<Aut

Note

When using drag-and-drop, your target location must be visible in the current view (including rows you can scroll to but not rows that have not been loaded). If you need to move a rule to a ranking not currently shown, you can use the Click to Show More bar at the bottom of the rules table to add rows to the current view. You also can use the dialog box described in the next procedure.

Disabling or Deleting Registry Rules

If you do not want a registry rule to be effective anymore, you can either disable it, which leaves it in the registry rules table, or delete it from the table. In either case, the rule stops affecting newly discovered files. However, files that were affected by the rule before it was disabled retain any file state assigned to them by the rule.

If you think you might use the rule again, disabling it temporarily is the best choice.

To disable a registry rule:

1. In the console menu, choose **Rules > Software Rules**, and when the Software Rules page appears, click the **Registry** tab. The Registry Rules table appears.
2. Click the View Details button next to the rule you want to disable. The Edit Registry Rule page appears.
3. In the Status line, click the **Disabled** radio button, and then click the **Save** button at the bottom of the page. The rule is now disabled.

Deleting a rule eliminates it permanently – there is no undo or retrieval for a deleted rule. Because of that, be sure you actually want to delete the rule.

To delete a registry rule:

1. In the console menu, choose **Rules > Software Rules**, and when the Software Rules page appears, click the **Registry** tab. The Registry Rules table appears.
2. Click the Delete button (red circle with X) next to the rule you want to delete, and click **OK** on the confirmation dialog. The rule is now deleted.

Viewing Rule Status on Computers

Depending upon the number of agents managed by your App Control Server and the number that are disconnected, not all agents might receive new or updated rules in a short amount of time. The Related Views menu on the Edit page for an enabled rule provides links to two different filtered views of the Computers page to help determine the status of the rule on agent-managed computers. The choices are:

- **All Computers that have received this rule**
- **All Computers that have not yet received this rule**

This menu does not appear for rules that have never been enabled.

Sample Registry Rules

App Control is shipped with a series of disabled sample registry rules. You can examine the rules to see whether you might want to enable them, or to consider using them as templates that you modify to accomplish exactly what you want for your own registry protection.

Important

Do not enable any of the sample registry rules without examining their fields, including which registry paths they apply to and what action (Block, Prompt, Report, Allow) they involve. You also can configure the Action for a rule to Report for a period of time before you make it fully active (i.e., blocking, prompting or allowing actions).

Example: Report Changes to Internet Explorer Trusted Zone

The example here starts with fields from the sample rule “[Sample] Report Changes to Trusted Zones,” which is included in the console but disabled by default. This rule reports changes to the sites or IP addresses in the Internet Explorer Trusted Zone on machines running the App Control Agent. Because you may give higher privileges to sites in the trusted zone, any changes to that zone could be a security risk.

To begin the process, go to the **Registry** tab and then click on the View Details button next to the “[Sample] Report Changes to Trusted Zones” rule.

Edit Registry Rule

General

Name: [Sample] Report Changes to Trusted Zones

Description: Generate an event whenever a change is made to the sites or IP addresses in the Internet Explorer Trusted Zone

Status: Enabled Disabled

Definition

Expert Mode: On Off

Platform: Windows

Write Action: Report

Registry Path:

- *\software\microsoft\windows\currentversion\internet s
- *\software\microsoft\windows\currentversion\internet s

Source Process: Any Process

User or Group: Any User

Rule Applies To

Policies: All Current and Future policies
 Selected policies

History

Created By: System
Date Created: Jun 29 2020 07:00:29 PM
Last Modified By: System
Date Modified: Jun 29 2020 07:00:29 PM

Save & Exit Save Cancel

As the description says, this rule generates a App Control event whenever a registry change is made that will change the sites or IP addresses in the Internet Explorer Trusted Zone. The fields are:

- **Write Action: Report** – This indicates that the rule only reports changes matching the rule – it does not block an action or allow an action that would otherwise be blocked. If you wanted to create a more restrictive rule, you could change this to **Prompt**, in which case each user on a computer running the App Control Agent would have the opportunity to block or allow Registry changes matching the rule. Or you could **Block** any changes matching the rule.
- **Registry Path:**
 - *\software\microsoft\windows\currentversion\internet settings\zonemap\domains*
 - *\software\microsoft\windows\currentversion\internet settings\zonemap\ranges*

– This rule includes two paths. Because the paths starts with *, any attempt to write to them, whether it starts with HKCU, HKLM, or another allowed prefix, will match the rule. Because the paths end with a slash and asterisk, keys and values at and below **domains** and **ranges** (respectively) will match the rule.

- **Process: Any Process** – Any process attempting registry writes that match the other fields activates the rule.
- **User or Group: Any User** – Any user attempting registry writes that match the other fields activates the rule.
- **Rule applies to: All policies** – All policies, and therefore all Windows computers running the App Control Agent, are subject to this rule.

If you enable this rule, registry write attempts matching the rule appear on the Events page. You can search for them by clicking the Show Filters button on the Events page and creating a filter for “Subtype is Report write (registry rule)”. When you find an event report matching this rule, you might respond in one of several different ways:

- If the change is undesirable, undo the change (outside of App Control) and create a new rule preventing that change from happening again (rather than just reporting it). Use wildcards or multiple paths to make the rule as narrow or broad as necessary.
- Allow the change if you consider it benign or desirable.
- Use the file information on the App Control Server to obtain information about the process that has attempted the modification.

Autostart Rules

The table of Registry Rules for this release includes an Autostart Rules rule that is actually a collection of rules. It is disabled by default. When activated, this rule set reports and optionally blocks attempts to modify registry locations that control what happens when you startup a computer. For example, one of the many paths covered by the Autostart Rules is:

HKLM\Software\Microsoft\Windows NT\CurrentVersion\Winlogon

If you want to test the impact of this set of rules before making it active, you can choose Report on the Write Action menu for the rule. Then, after some time has elapsed, you can go to the Events page and filter for “Rule name contains Autostart” to see what events have been triggered by this rule set. If you determine that activating the rule will not interfere with your operations, you can change the Write Action value to Block (or Prompt).

On the Edit Registry Rule page for Autostart Rules, the Registry Path is shown as <AutostartRules>. This macro refers to the current list of locations controlled by this rule. The list is maintained within the App Control Server and not enumerated in the rule definition. It is expected to be updated and expanded with future releases. If you need more detail about specific locations affected by this rule in your version, please contact Carbon Black Support.

Note

Pre-7.0 releases of Bit9 Parity had Registry Rules that affected a small subset of the locations included in the new Autostart Rules set. If you used any of these startup rules, you might want to use the Autostart Rules instead for greater protection on startup.

Chapter 16

Memory Rules

This chapter describes Memory Rules, which can protect a process from being accessed or altered by other processes.

Unified Management

Memory Rules can be centrally managed for multiple servers through the Unified Management feature. This is described in [“Unified Management of Rules”](#) on page 791.

Platform Note: Memory rules affect only computers running Windows operating systems.

Sections

Topic	Page
Overview	471
Specifying the Notifier for Memory Rules	472
Creating Memory Rules	473
Memory Rule Fields	476
Specifying Target and Source Processes	481
Rule Ranking	485

Overview

Memory Rules allow you to monitor attempts to access a process on a Windows computer, and if you choose, protect the process from being accessed or altered by any other process(es) or user(s). When a rule matches your criteria, you can *block* read, write or execution access to a matching process, *report* on access, or *prompt* the user on the agent system to block or allow access. There also are advanced options for special cases.

If an in-memory malicious attack occurs on a system protected by the App Control Agent, a properly configured memory rule can prevent that attack from spreading to other processes, or even from accessing information in other processes. Memory rules limit the vulnerability of a protected computer. They can also protect specific applications or processes from termination or other manipulation by users or malicious code.

You can view a list of memory-rule-related events, including blocked actions caused by memory rules, on the Events page by choosing **Memory** on the Saved Views menu.

Important

There are two built-in rules named *Tamper Protection*, ranked 1 and 2 by default, that help protect agent computers. Do not edit these rules, and do not disable or reorder them unless instructed to do so by Carbon Black Support. Check the description field for any rule before you consider modifying it.

Rule Scope

You can create memory rules that apply to all Windows computers, regardless of which user and what process attempts to access the process you specify. You also can create a more focused scope for a rule by specifying one or more of the following criteria:

- **Source-process-specific** – You can make a rule apply only when a particular source process attempts to access the target process you are monitoring or protecting.
- **User- or group-specific** – You can make the rule apply only to a particular user or group of users.
- **Policy-specific** – You can choose to limit a rule to computers in specified policies.
- **Server-specific** – If you have Unified Management enabled, you can choose to limit a rule to *computers reporting to specified servers* in the management group. This is described in [“Unified Management of Rules”](#) on page 791.
- **Rule order** – Memory rules are evaluated in order of *Rank*, a column that is displayed by default on the Memory Rules table. The rule ranked ‘1’ has the highest rank, ‘2’ is next, and so on. You can change the order of rules to have a more specific rule evaluated before a more general one. For example, you can create a rule that applies when *a particular user* attempts to access a process, and put that before a rule that applies when *any other user* attempts to access the process. See [“Rule Ranking”](#) on page 485 for more details.
- **Conditional Macros** – You can use certain macros to restrict the conditions under which specific parameters in rules are applied. Only agents meeting the “test” described in the macro will attempt to match the parameter prefixed with the macro. Most of these macros are *OnlyIf* macros with different arguments, such as `<OnlyIf:OSVersionIs:10.6.8>` and `<OnlyIf:HostName:*SMITH-1*>`.

There are certain restrictions on where memory rules are effective:

- A memory rule cannot be used to protect a process from itself. For example, you cannot create a rule that prevents a process from terminating itself, or from modifying its own memory.
- Memory Rules are not supported on Mac or Linux computers, or computers running Windows Server 2003 64-bit.
- Kernel Memory Access rules are supported only on computers running Windows XP.
- Dynamic Code Execution rules are supported only on 32-bit versions of Windows XP, Windows 2003, Windows Vista and Windows 7. They are not supported on any 64-bit Windows operating systems, nor are they supported on any Windows 8 or 10 versions.
- For computers in Visibility mode policies, memory rules that would block writing or prompt users for a decision act as report-only rules, and do not block or prompt.

Exporting and Importing Memory Rules

You can export memory rules from one server and import them to another. There are buttons for this purpose on the Memory Rules page. See [“Exporting and Importing Rules”](#) on page 431 in the Custom Rules chapter for more information.

Specifying the Notifier for Memory Rules

The App Control Agent provides *notifiers* that can be displayed when a rule blocks an action or prompts the user for a decision to allow or block an action. For each memory rule, you can choose from two sources for the notifier:

- **Use Policy Specific Notifier** – Each Policy includes an Advanced Setting, “Enforce memory rules”, which is always on. This policy setting has a Notifier field in which you can specify the notifier that appears on agent computers when memory rules block an action.

If you choose Use Policy Specific Notifier for a rule, it is possible that the policy specifies <none> as the Notifier for registry rules. In this case, a notifier will not be shown, even for a Prompt rule. Unless you are certain that you never want to prompt the user for a response to a rule, choosing <none> for the custom rule notifier in a policy is not recommended. See [“Advanced Settings”](#) on page 187 for more information.

- **Custom Notifier** – Instead of the policy-specific notifier, you can choose or create a custom notifier for a memory rule. The choices appear on a menu on the Add/Edit Memory Rule page. Custom Notifiers for Prompt rules must have a notifier. Custom Notifiers for Block rules allow you to choose <none> so that no notifier appears.

When you choose Block as the rule action, you can choose <none> on the Custom Notifier menu since it is possible you want the rule to block actions without notification. A Prompt rule requires a user choice, so when you choose Prompt as the rule action, the Custom Notifier menu does not include <none>.

Note

If you are using Unified Management and create a Memory Rule that applies to more than one server, client servers will use default notifiers, even if a custom notifier is specified on the management server.

See [Table 63](#) below for the memory rule notifier settings. See [Chapter 20, “Endpoint Notifiers and Approval Requests,”](#) for more on notifiers.

Creating Memory Rules

In addition to providing a name, to create a memory rule, you need to provide the information shown in bold in the left column of the table below and enter it in the Add Memory Rule page in the locations on the right:

General Description	Field on Add/Edit Memory Rule Page
If this/these source process(es) ...	Source Process
...and/or this/these user(s) ...	User or Group
... attempt the following memory access type ...	Permissions
... to process(es) at this/these location(s) ...	Target Process
... on computers in this/these policy(ies) ...	Rule applies to/Policies:
... on computers reporting to this/these App Control server(s) ...	Rule applies to/Servers (if Unified Management is enabled)
... then this action should be taken.	Action

* Additional actions and other options are available in Expert Mode. See [Chapter 17, “Expert Rules,”](#) for more details.

Each of these fields could have multiple matching items, or the rule could specify *all* items in that class (e.g., the rule applies to all users, or all policies, or all source processes).

To add (create) a memory rule:

1. In the console menu, choose **Rules > Software Rules**. The Software Rules page appears.
2. Click **Memory** on the left menu or the tab on the page. The Memory Rules table appears, showing several built-in rules and any other memory rules that have been created on your server.

Software Rules

Updaters Rapid Configs Publishers Users Directories Files Custom **Memory** Registry Scripts Reputation

Saved Views: (none) Add Group By: (none) Ascending

Show Filters Show Columns Export to CSV Refresh Table

+ Add Memory Rule Export Rules Import Rules Search: Autom:

Select	Rank	Status	Platform	Name	Action	Operation	Per
<input type="checkbox"/>	1	●	Windows	Tamper Protection	Block, Silent	Create Handle, Duplicate Handle, Al...	Ad
<input type="checkbox"/>	2	●	Windows	Tamper Protection	Block, Silent	Create Handle, Duplicate Handle, Al...	Dyr

- Click the **Add Memory Rule** button. The Add Memory Rule page appears.

Add Memory Rule

General

Name:

Description:

Status: Enabled Disabled

Definition

Expert Mode: On Off

Platform:

Action: Use Policy Specific Notifier

Permissions:

Target Process:

Source Process:

User or Group:

Rule Applies To

Policies: All Current and Future policies
 Selected policies

Save & Exit
 Save
 Cancel

- In the Name field, enter the name you want to appear on the list of rules. You may also provide a longer, optional Description below the Name field.
- Enter the remaining information you want for this rule (see [Table 63, “Memory Rule Fields”](#) on page 476) and then click the **Save** button if you need to remain on the page or **Save & Exit** if to go to the Memory Rules table. By default, a new memory rule is **Disabled** and ranked #1, listed at the *top* of the Memory Rules table.
- Before you enable a rule, change its rank unless you want it to take precedence over (and perhaps preempt) all other rules. You can change rank using the arrows in the Rank column or drag-and-drop (if the table is sorted by rank), or you can click on the rank number and enter a new rank in the dialog box. See [“Rule Ranking”](#) on page 485 for more details.
- When you are satisfied with the rank and want to enable the rule, click the toggle switch in the Status column of the Memory Rules table. The button in the switch moves to the right and the background turns from white to green.

Editing a Memory Rule

Editing a Memory Rule is very similar to creating one. If you have permission to edit the rule, you can edit any field, including the rule name.

To edit a memory rule:

1. In the console menu, choose **Rules > Software Rules**. The Software Rules page appears.
2. Click **Memory**, either on its tab or in the menu to the left of the page. The Memory Rules page appears.
3. On the rules table page, click the View Details button for the rule you want to edit.
4. On the Edit Memory Rule page for that rule, make your changes. [Table 63](#) describes the custom rule fields that you might choose to edit.
5. Click either **Save** (to remain on the edit page) or **Save and Exit** (to return to the table page).
6. If you clicked Save, a confirmation message appears on the page. Click the message to clear it from the page. If an error occurs, review the error message and correct the conditions that caused the error before saving again.

Unified Management

If you are using Unified Management, and you edit a unified rule shared with other servers, a “wizard” shows the progress as the edited rule is saved on each server.

See [“Unified Management of Rules”](#) on page 791 for more details.

Copying a Memory Rule

There is a *Copy this rule* command on the right menu on the Edit Rules page for Custom, Registry, and Memory Rules. This is for making copies of the rule on the same server. You might do this so that you can customize a sample rule while preserving the original settings as a template. It also allows you to make slightly different rules for different policies without having to manually provide all of the settings for each one.

To copy a memory rule:

1. On the Memory Rules table page, click the View Details button to open the details page for the rule you want to copy.
2. On the details page, click **Copy this rule...** in the Actions menu on the right side of the page. This opens a dialog box. By default, the copy is named using the original rule name plus “(copy)”.
3. In the dialog box, change the rule name if you want something more descriptive that what is there.
4. If you want the new rule enabled immediately, check the **Enable copied rule** box. Otherwise, uncheck the box.
5. Click **OK**. The copied rule is created and its details page replaces the details page for the original rule.
6. Make any changes you would like to make in the new (copied) rule and **Save** or **Save & Exit**.

Memory Rule Fields

Table 63 shows the fields available on the Add/Edit Memory Rule page. Column headings on the rule table page are shown when they differ from the Add/Edit page.

Table 63: Memory Rule Fields

Field	Description
Name	Name by which this rule is identified in the Memory Rules table. (Required)
Description	Optional information about the memory rule. This can be any text you choose to enter.
Rank (Table only)	The rank of this rule in order of evaluation. The rule ranked '1' in the table is evaluated before the rule ranked '2', and so on.
Status	Radio buttons that make this rule Enabled or Disabled. This allows you to create a rule that you use only at certain times, or to temporarily disable the rule without losing the information used to create it.
Expert Mode	Radio buttons that make turn Expert Mode on and off (the default). Expert mode provides more options than standard mode but does not have all of the error-checking that other rule types have, so it is possible to create unexpected (and unwanted) outcomes without being warned during rule creation. These rules are intended for use by Carbon Black Support or Services representatives, or customers working with them. See Chapter 17, "Expert Rules," for more details.
Platform	Platform for which this rule is effective. This is a read-only field and the value is always Windows. Memory rules do not have any impact on non-Windows platforms.
Action	The action you want the App Control Agent to take when there is an attempt to access or alter a process matching this rule. Table 64 shows the options for this field when creating or editing a non-expert rule. Other actions configured in Expert rules or underlying the standard menu commands may appear in the table column for this field.
Operation (Table only)	The type of operation the rule affects.
Action (Legacy) (Table only)	This column shows actions and operations for the rule as shown in the Action column in pre-8.1.6 versions, or it shows "Expert Action(s)" in cases where expert rule information was not previously shown. Note: This field is present strictly for continuity with older versions – you should use the separate Action and Operation columns for the most accurate description of the rule.

Field	Description
Use Policy Specific Notifier	If you choose Block or Prompt as the Action, this checkbox appears to the right of the Action choice. If you check the box, the notifier that appears when a memory rule blocks an action is the notifier specified for the Enforce Memory Rules setting in the policy for the computer on which the action was blocked. If not checked, you can choose a custom notifier from the Custom Write Notifier menu.
Custom Write Notifier	If you choose Block or Prompt as the Action, and you do not check the Use Policy Specific Notifier box, this menu appears. When Block is the Action, you can choose any notifier from the menu. The menu also includes a <none> option so that you can disable the notifier for this rule. When Prompt is the Action, you can choose any notifier on the menu. However, Prompt rules <i>must</i> display a notifier, so there is no <none> choice in this case. Note: If you use Unified Management to create a rule that applies to more than one server, client servers will use default notifiers, even if a custom notifier is specified on the management server.
Permissions	The type of access you want to affect with this rule. Table 65 shows the permissions options.
Target Process (Path in table)	The process(es) you want this rule to restrict, monitor, or allow access to. See “Specifying Target and Source Processes” on page 481 for a description of the ways you can define a target process.
Source Process (Process in table)	This field allows you to apply the rule only when a specified Source Process requests access to the Target Process. Table 66, “Source Process Menu Options,” on page 484 shows the menu choices. “Specifying Target and Source Processes” on page 481 describes options for entering a path. Note: No Target Process specification is needed for Kernel Memory Access or Dynamic Code Execution rules because the Source Process applies the rule to itself.
User or Group	This field allows you to specify users or groups to which this rule applies. See “Specifying Users or Groups” on page 484 for detail on specifying users or groups.
Rule Applies To: Servers (Add/Edit page only)	Radio buttons allow you to apply the rule to the current server, All Servers or Selected Servers . If you choose Selected Servers , a list that includes the current server and of all App Control servers managed by this server appears, each with a checkbox. In addition, policies for the servers you include appear in the Selected policies list. Unified Management: This field appears only if Unified Management is configured on the server you are logged into.
Unified Server Source (Table only)	If this is a unified rule, the name of the unified management server that created or copied the rule.

Field	Description
Rule Applies To: Policies (Policy in the table)	<p>Radio buttons allow you to apply the rule to All policies or Selected policies. If you choose Selected policies, a list of all policies available on your App Control Server appears, each with a checkbox. You can check as many policies as you choose.</p> <p>Unified Management: If Unified Management is configured on the server you are logged into, and if you have chosen to apply the rule to additional servers, policies for all selected servers appear in this list.</p>
Is Global (Table only)	<p>Indicates whether the rule applies to all policies (Yes) or only selected policies (No).</p>
Rule Applies To: Override Permissions (Add/Edit page only)	<p>Radio buttons allow you to specify whether administrators on other servers can modify rules sent via Unified Management on their own server. The options are No Override, Partial Override (allows changing rank) and Full Override (allows editing and changing rank).</p> <p>Unified Management: This field appears only if Unified Management is configured on the server you are logged into and this rule is applied to more than the current server in the Rule Applies To:Servers field.</p>
History	<p>For existing rules, a History panel on the Edit Rule page appears showing some or all of the following fields. In addition, these fields can be added as columns on the rules table page.</p> <ul style="list-style-type: none"> • Created By – If the rule was created on this server, the user who created it. Rules created during server installation or upgrades show “System” in this field. • Date Created – If the rule was created on this server, when it was created. • Last Modified By – If the rule has been modified since creation or import, the user who modified it. • Date Modified --If the rule has been modified since creation or import, when it was modified. • CL Version – Rules created after server installation also show the CL (config list) number that first contained the rule so that you can compare an agent CL number to determine whether the agent has received the rule. • Imported – (In the table only) indicates whether the rule was imported (Yes/No). • Imported By – If the rule was imported to this server, the user who imported it. • Imported Date – If the rule was imported to this server, when it was imported.

Specifying the Rule Action

The Action for a Memory Rule defines what you want App Control to do if there is a memory access attempt matching the rule. [Table 64](#) shows the options.

Table 64: Action Menu Options

Field	Description
Block	<p>Prevent access to, termination of, or modification of processes matching this rule.</p> <p>When Block is chosen, the Use Policy Specific Notifier checkbox and a Custom Write Notifier menu appear. These allow you to specify the notifier, if any, that appears when the rule blocks an action. See Table 63 for more details.</p>
Block Silently	<p>Prevent access to, termination of, or modification of processes matching this rule. Do not display a notifier, and do not generate an event.</p>
Prompt	<p>Present a notifier dialog to the endpoint user when there is an attempt to access, terminate, or modify processes matching this rule. The dialog choices are Block or Allow. Once the user responds to the dialog, the choice applies anytime the same process matches the same rule on that computer – the user will not be prompted again in this case.</p> <p>When Prompt is chosen, the Use Policy Specific Notifier checkbox and a Custom Write Notifier menu appear. These allow you to specify the notifier that appears to prompt the user. See Table 63 for more details.</p> <p>Note: Use of Prompt as the action for Dynamic Code Execution rules is <i>not recommended</i> because the combination can have destabilizing effects on computers running the App Control Agent.</p>
Report	<p>Do not block access, termination, or modification of matching processes but report the actions as events.</p>
Allow	<p>Allow all memory/process operations that match this rule. This is the default behavior if there is no rule for a particular target or source process.</p> <p>Use of Allow gives you a way to create an exception to a more general rule that blocks at a particular location. For example, if you create a rule that blocks all memory operations at</p> <pre>c:\Program Files\InterestingApp*</pre> <p>you could use Allow to create a higher ranking rule that allows operations at</p> <pre>c:\Program Files\InterestingApp\Subfolder\</pre>

Specifying the Rule Permissions

Permissions define the type of access you want to affect with this rule, such as read, write or execution. Some options allow you to control multiple types of access. [Table 65](#) shows the options available on the permissions menu.

Table 65: Permissions Menu Options

Field	Description
Control Process	Access required to control the execution of a process or thread, including the ability to terminate the process.
Read Access	Access required to retrieve, copy or duplicate certain information about a process or thread. If all you are concerned about is data loss or theft, you might use this choice with the Block Action.
Write Access	Access required to modify a process or thread and its attributes.
Dynamic Code Execution	Affects whether an application can execute code not associated with an executable image. This protection prevents arbitrary or floating code execution of the sort used by many forms of malware. Protects against attempts to disable Dynamic Execution Protection (DEP). Applies only to 32-bit versions of Windows XP, Windows 2003, Windows Vista and Windows 7. Important: Do not create a Dynamic Code Execution rule with Prompt as the action choice – this could cause undesirable results on agent computers.
Kernel Memory Access	Affects whether a user-mode process can access kernel memory. You can create rules allowing access by a legitimate application while denying access for all other applications. Applies only to Windows XP.
Write + Control	Both write and control permissions. You can use this Permission choice and choose Block as the Action to prevent an attack on a process, such as a malicious code injection, termination, or other alterations.
Read + Write + Control	Read, write, and control permissions. This is the option you would use, along with the Block Action, to prevent data loss or theft as well as attacks. This does not include Dynamic Code Execution or Kernel Memory Access.
Advanced...	This option allows for very detailed control of memory access. Contact Carbon Black Support before using the Advanced option.

Specifying Target and Source Processes

You usually specify two processes in a memory rule:

- **Target Process** – The process(es) you want the rule to restrict, monitor, or allow access to.
- **Source Process** – The process(es) requesting access to the Target Process.

Definition

Expert Mode: On Off

Platform: Windows

Action: Block Use Policy Specific Notifier

Permissions: Read Access

Target Process:

Source Process: Any Process

User or Group: Any User

When you specify Target Process in a Memory Rule, you have several options for defining the string for that field. These same options can be used when you choose one of the two Source Process options that require entry of a path (*Specific Process...* or *Any Process Except ...*). These options are:

- **Specify a directory or a process** – You can enter a process specification that exactly identifies a file by path and name so that only that file matches the rule. You also can enter a specification that identifies a directory, and so affects processes running from files in that directory and its subdirectories.
- **Specify a local drive or UNC path** – You can identify a process by using a local drive name, such as *C:\folder1\subfolder\application.exe*. You also can enter a remote process by using a UNC path, such as *\\computer\dir\application.exe*. Mapped drives in a path or process specification are not recognized.
- **Use wildcards** – You can use wildcards ('?' for any one character and '*' for zero or more characters) to expand the scope of a process specification or help you match a file or folder whose exact location you don't know. Wildcards may be used at the beginning, end or middle of a path.
- **Use macros** – You can use special App Control macros to identify certain well known folders in the Microsoft Windows environment, even if you don't know their exact location on all agent computers.
- **Specify multiple paths or processes** – You can add more than one process path definition per rule.

Specifying a File or Directory

You can specify a directory or a file as the Target or Source Process path. Using a directory applies the rule to processes in that directory and any of its subdirectories (unless higher-ranked rules apply to processes or subdirectories in it).

To identify a Process definition as a directory, you must end it with a backslash (\) or a backslash and asterisk (*). Without the backslash, the rule will attempt to match a *file* by

the name you provided, not a directory. For example, either of the following correctly identifies a directory in a process definition:

```
c:\folder1\subfolder2\  
c:\folder1\subfolder2\*
```

However, the following is not recognized as a directory:

```
c:\folder1\subfolder2
```

If you use path macros in a process definition, the macro is treated as a directory, even if you don't follow the it with a backslash. See [Using Macros](#).

Using Wildcards

You can use wildcard characters in the Process fields. Asterisk (*) indicates zero or more characters and question mark (?) indicates one character. You can use wildcards to specify partial paths or multiple paths for directories that appear in different locations on different computers (although macros might be a more effective way to accomplish this – see [Using Macros](#)). Wildcards are not allowed inside of macros.

The number of wildcards in a process specification is not restricted. For example, you could define a path as:

```
*\Win*\folder?\
```

Caution

When you use wildcards, be careful not to create a rule that is so broad that it will interfere with activity that is required for legitimate use by an application or the operating system. Don't use the asterisk wildcard by itself in Target Process field, especially with rules that block multiple types of access, unless you are absolutely certain it will not interfere with necessary operations on the agent computer.

Automatic Path Conversions

When a rule is processed, file paths in a process field undergo several automatic path conversions if they contain certain symbols:

- Any path that ends with a slash has the "*" wildcard added at the end of the path.
- Any path that has no slash or drive letter has "*" added at the beginning of the path
- Drive letters may be used in a path as long as they are for local fixed volumes. Mapped drive letters should not be used because there is no guarantee that the mapping exists on all computers.
- The string ".*:" applies to all attached storage volumes except for floppy disks and CD-ROMs.

Specifying Devices in Paths

You can create rules that apply to processes on some or all devices on the agent computer by including `\device\` in the path. For example:

- `\device*\` specifies all devices.
- `\device\harddisk*\` specifies attached storage volumes except for floppy disks and CD-ROMs.

- `\device\cdrom*` specifies CD-ROM devices.

Using Macros

You can use certain macros in the Process fields. You can see a menu of macros by typing the left angle bracket (<) character in either of the Process fields. There are two types of macros supported in Memory Rule processes:



- **Path macros** – These are a subset of the well known folders in the Microsoft Windows environment, and they always identify a location rather than a specific file. A *path* macro can be used only at the beginning of a Path or File specification in a rule (i.e., with no other text before it in the string).
- **Registry macros** – These are macros that specify strings in the Windows Registry. A *registry* macro can be used anywhere in the Path or File specification.

Macros can be an effective way to define a rule that works on all agent computers even when the processes you want to specify are in different locations on different computers.





See “Using Macros in Rules” on page 410 of the Custom Rules chapter for a description of path and registry macros. The macros described there may be used in the Process fields of a memory rule.

Entering Multiple Target or Source Processes

For each Process field in a Memory Rule, you can enter more than one string. For example, when you have entered the first Memory Path for a rule, click the Expand button to the right of the box.

Target Process:  

You can then add process paths by typing them in the box and clicking **Add** after each one.

Target Process:    

You can remove any process path by clicking the Expand button, selecting the path in the list below the box, and clicking the **Remove** button.

If you enter multiple paths in either process field in a rule, the Memory Rules table shows the first path and then “(multiple)” in the relevant column for this rule. Moving the mouse over the value shows a tooltip with the complete list of processes for the rule.

Path	Process	Date Modified
<ProgramFilesx86>\Adobe\Acrobat 11.0* (multiple)	*	
<div style="border: 1px solid black; padding: 5px;"> <ProgramFilesx86>\Adobe\Acrobat 11.0*, <ProgramFilesx86>\Adobe\Creative Cloud*, <ProgramFilesx86>\Adobe Photoshop CS5.1*, <ProgramFilesx86>\Adobe\Flash Player* </div>		

The Source Process Menu

The Source Process field in a Memory Rule specifies the process that is requesting access to the Target Process. The Source Process menu includes options that are completely defined by your menu choice, such as **Any Process**, and options that require entry of a path to the process(es):

Table 66: Source Process Menu Options

Field	Description
Any Process	Applies the rule to any process that attempts to access the target process.
Any Promoted Process	Applies the rule to any source process that is promoted at the time the rule is evaluated. A promoted process is any approved process that is marked as an installer, or has been promoted as a consequence of a custom rule, or is an approved process launched by a promoted process.
Any System Process	Applies the rule to every source process that is running under the security context of the Local System user. This has the same effect as choosing Local System in the User or Group menu.
Specific Process...	Opens a text box below the menu, into which you can enter source process(es) you want controlled by this rule.
Any Process Except...	Opens a text box below the menu, into which you can enter the source process(es) you <i>do not</i> want controlled by this rule. Note: If you specify a User or Group and also choose Any Process Except from the process menu, the rule is enforced <i>unless the exception process is being executed by the user or group</i> .

Specifying Users or Groups

You can create a rule that applies only when specific users or users in specific groups attempt an action. The choices for User or Group on the Add/Edit Memory Rule page are:

- **Any Users** – applies the rule to all users.
- **Specific User or Group...** – opens a text box below the menu, into which you can enter AD users or groups in the format *userorgroupname@domain* or *domain\userorgroupname*
- The other menu choices are built-in Windows groups, such as **Authenticated Users** and **Local Administrators**.

Notes

- On Windows Vista and later, membership in pre-defined security groups like Administrators requires that the application is run as an administrator. If a group definition is necessary for a rule, consider using security groups you define rather than the pre-defined groups.
- There is a brief delay after a user logs in before group membership is established and group-based rules become effective. This delay may be longer if you have a large number of rules. If you need a rule to be effective as soon as possible after a user logs on, do not specify a user *group* in the rule. Rules that specify a *username* or *SID* are always active and won't be affected by this delay.

Rule Ranking

Memory rules have a “Rank” number and are evaluated from lowest number to highest number, beginning with the rule ranked ‘1’. By default, rules appear on the Memory Rules page in their rank order, but you can sort the table by other columns if you choose.

If a memory-related action matches a rule’s definition, that rule is evaluated. Rule processing continues down the rank order to see whether any other rules match the current memory-related action. If there is another match, what happens next depends on the *Permissions* setting for the rules:

- If the action matches two rules, but these rules have different permissions settings – for example, one is applied to *Read Access* and the other is applied to *Write Access* – both rules are evaluated. In this case, if there is a third matching rule that is applied to *Control Process*, that rule is also evaluated.
- If the action matches two (or more) rules and all have *the same* permissions settings – for example, both are applied to *Write Access* – only the first rule is evaluated. There is one exception to this behavior – a rule whose action is Report does not stop processing of lower ranked rules with the same permissions setting.

You can change the ranking of rules if you decide that you want one of your rules to be considered before its current rank position.

Important

There are two built-in rules named *Tamper Protection*, ranked 1 and 2 by default, that help protect the server. Do not rank other rules higher than these unless instructed to do so by Carbon Black Support.

To change the rank of a memory rule:

1. On the Memory Rules page, make sure the Rank column is displayed.
2. If you want to be certain you can see the rules ranked immediately before and after the rule you are moving, sort the table by rank and remove any filters.
3. To change the rank of a rule, you have three options:
 - In any table that displays the rank column, you can click on the rank number and enter a new rank number in the dialog box.

- If the table is sorted by rank and not filtered, arrows appear next to the rank, and you can click the up or down arrow button next the to rule to change its rank.
- If the table is sorted by rank and not filtered, you can hold down the left mouse button with the cursor over the rule and drag the rule to a new location.

<input type="checkbox"/> Select 2	Rank ▲	Status	Platform	Name	Action	Operation
<input type="checkbox"/>	↑ ↓ 1	<input checked="" type="checkbox"/>	Windows	Tamper Protection	Block, Silent	Create Handle, Duplicate Handle, AI...
<input type="checkbox"/>	↑ ↓ 2	<input checked="" type="checkbox"/>	Windows	Tamper Protection	Block, Silent	Create Handle, Duplicate Handle, AI...

Note

When using drag-and-drop, your target location must be visible in the current view (including rows you can scroll to but not rows that have not been loaded). If you need to move a rule to a ranking not currently shown, you can use the Click to Show More bar at the bottom of the rules table to add rows to the current view. You also can use the dialog box described in the next procedure.

Disabling or Deleting Memory Rules

If you do not want a memory rule to be active anymore, you can either disable it, which leaves it in the rules table, or delete it from the table.

If you think you might use the rule again, disabling it temporarily is the best choice.

To disable a memory rule:

1. In the console menu, choose **Rules > Software Rules**, and when the Software Rules page appears, click the **Memory** tab. The Memory Rules table appears.
2. Click the View Details button next to the rule you want to disable. The Edit Memory Rule page appears.
3. In the Status line, click the **Disabled** radio button, and then click the **Save** button at the bottom of the page. The rule is now disabled.

Deleting a rule eliminates it permanently – there is no undo or retrieval for a deleted rule. Because of that, be sure you actually want to delete the rule.

To delete a memory rule:

1. In the console menu, choose **Rules > Software Rules**, and when the Software Rules page appears, click the **Memory** tab. The Memory Rules table appears.
2. Click the Delete button (red circle with X) next to the rule you want to delete, and click **OK** on the confirmation dialog. The rule is now deleted.

Viewing Rule Status on Computers

Depending upon the number of agents managed by your App Control Server and whether any are disconnected, not all agents might receive new or updated rules in a short amount of time. The Related Views menu on the Edit page for an enabled rule provides links to two different filtered views of the Computers page to help determine the status of the rule on agent-managed computers. The choices are:

- **All Computers that have received this rule**
- **All Computers that have not yet received this rule**

This menu does not appear for rules that have never been enabled.

Chapter 17

Expert Rules

Expert Rules are Custom, Memory and Registry Rules created with a special interface that provides many more options than the standard interface for these rules. They are intended for expert users working with Carbon Black Support or Technical Services.

Sections

Topic	Page
Overview	489
Enabling the Expert Rules Interface	489
Expert Rule Definitions	492
Expert Rule Operations	492
Expert Rule Actions	495
Tags and Tagging Actions in Expert Rules	499
Expert Rule Examples	502

Overview

Custom, Memory and Registry Rules in App Control provide built-in rule types to help you create rules suitable for many situations. Beginning with v8.0.0, Custom, Memory, and Registry Rules include a new "Expert" option. Expert rules, as the name implies, are intended for expert users. This new rule type exposes more operations you can use to trigger a rule and more actions you can take when the rule is triggered. It also allows you to combine multiple operations and actions into one rule.

Important

Expert Rules do not have all of the error-checking that other rule types have, and it is possible to create unexpected (and unwanted) outcomes without being warned during rule creation. These rules are intended for use by Carbon Black Support or Services representatives, or customers working with them.

Enabling the Expert Rules Interface

Because Expert Rules offer a many more options, their interface differs from that of standard rules. Instead of menus, Expert Rules are configured through checkboxes.

To open the Expert Rules interface for Memory or Registry Rules:

1. On the Memory or Registry Rules table page, click the **Add Memory Rule** or **Add Registry Rule** button.
2. On the Add Memory Rule or Add Registry Rule page, click the **On** radio button in the Expert Mode field in the Definition panel.

The screenshot shows the 'Definition' panel for configuring an Expert Rule. The 'Expert Mode' field has two radio buttons: 'On' (selected) and 'Off'. Below it, the 'Platform' is set to 'Windows'. The 'Write Action' is set to 'Block', and the 'Use Policy Specific Notifier' checkbox is checked. The 'Registry Path' field is empty. The 'Source Process' is set to 'Any Process', and the 'User or Group' is set to 'Any User'. Information icons are present next to the 'Registry Path', 'Source Process', and 'User or Group' fields.

To open the Expert Rules interface for Custom Rules:

1. On the Custom Rules table page, click the **Add Custom Rule** button.
2. On the Add Custom Rule page, choose **Expert** on the Rule Type menu in the Definition panel.

Definition

Platform: Windows

Rule Type: File Integrity Control

Write Action: File Integrity Control

Path or File: Trusted Path

Process Exclusion: Execution Control

Rule Applies To: File Creation Control

Policies: Performance Optimization

Policies: Advanced

Policies: Expert

Select a policy

Add Custom Rule

General

Rule Name:

Description:

Status: Enabled Disabled

Definition

Platform:

Rule Type:

Operations

<p>Basic Operations</p> <p><input type="checkbox"/> Open</p> <p><input type="checkbox"/> Open Execute Intent</p> <p><input type="checkbox"/> Read</p> <p><input type="checkbox"/> Cleanup</p> <p><input type="checkbox"/> Lock File</p> <p><input type="checkbox"/> Mmap Read</p>	<p>Modifying Operations</p> <p><input type="checkbox"/> Write Intent</p> <p><input type="checkbox"/> Write</p> <p><input type="checkbox"/> Write Delayed</p> <p><input type="checkbox"/> Delete</p> <p><input type="checkbox"/> Delete on close</p> <p><input type="checkbox"/> Rename</p> <p><input type="checkbox"/> Create New</p> <p><input type="checkbox"/> Permission Change</p> <p><input type="checkbox"/> Owner Change</p> <p><input type="checkbox"/> Mmap Write</p>	<p>Execute Operations</p> <p><input type="checkbox"/> Execute</p> <p><input type="checkbox"/> Script Execute</p> <p><input type="checkbox"/> Process Create</p> <p><input type="checkbox"/> Process Terminate</p> <p><input type="checkbox"/> Image Load</p>
--	--	---

Actions

<p>Authorization Actions</p> <p><input type="checkbox"/> Allow</p> <p><input type="checkbox"/> Block</p> <p><input type="checkbox"/> Report</p> <p><input type="checkbox"/> Prompt</p> <p><input type="checkbox"/> Terminate Source Process</p> <p><input type="checkbox"/> Suspend Source Process</p>	<p>Approval Actions</p> <p><input type="checkbox"/> Promote process</p> <p><input type="checkbox"/> Demote process</p> <p><input type="checkbox"/> Don't Promote Children</p> <p><input type="checkbox"/> Approve as installer</p> <p><input type="checkbox"/> Approve</p> <p><input type="checkbox"/> Promote Target Process</p> <p><input type="checkbox"/> Demote Target Process</p>	<p>Other Actions</p> <p><input type="checkbox"/> Trigger Action</p> <p><input type="checkbox"/> Finish Rule Group</p> <p><input type="checkbox"/> Stop Rule Processing</p> <p><input type="checkbox"/> Silent</p>	<p>Tagging Actions</p> <p><input type="checkbox"/> Tag Target</p> <p><input type="checkbox"/> Remove Target Tags</p> <p><input type="checkbox"/> Tag Process</p> <p><input type="checkbox"/> Remove Process Tags</p> <p><input type="checkbox"/> Add Global Tags</p> <p><input type="checkbox"/> Remove Global Tags</p>	<p>File Tracking Actions</p> <p><input type="checkbox"/> Ignore</p> <p><input type="checkbox"/> Dirty</p> <p><input type="checkbox"/> Never report</p> <p><input type="checkbox"/> Track</p>
---	--	--	--	---

Process Tag(s): i

Target Tag(s): i

Global Tag(s): i

Global Tag Exception(s): i

Path or File: i

Process: i

User or Group: i

Rule Applies To

Policies: All Current and Future policies
 Selected policies

Expert Rule Definitions

The basic components of an Expert Rule are the same as those of a non-expert rule: an *operation* that is being monitored, some combination of other conditions that must be met to match the rule, and an *action* to take when the rule is triggered. At least one operation and at least one action are required in an Expert Rule definition.

When multiple operations are defined in a rule, the rule triggers if *any* of them is true, as long as the action defined in the rule is possible for that operation.

Expert Rule Definition

If all other rule criteria are met (**source process, target file/path/process, user, policies**) ...

... and if any of the **Operations** defined in the rule are attempted ...

... then take the **Actions** defined in the rule (if available for the operation).

Expert Rule Operations

Expert Versions of Custom, Memory, and Registry Rules each have their own set of Operations choices:

- Custom Rule Operations – see [Table 67](#).
- Memory Rule Operations – see [Table 68](#).
- Registry Rule Operations – see [Table 69](#).

Table 67: Expert Custom Rules: Operation Settings

Column Name	Attempted Operation	Description
Execute Operations	Execute	Execution of a file
Execute Operations	Image Load	Loading of a file (dll, ocx, etc.) into memory
Execute Operations	Process Create	Creation of a new process
Execute Operations	Process Terminate	Termination of a process
Execute Operations	Script Execute	Execution of a script. For App Control agents to see a script execute, the appropriate script rules should be defined in the console.
Modifying Operations	Delete On Close	This operation corresponds to someone opening/creating a file with the FILE_FLAG_DELETE_ON_CLOSE flag set (meaning someone intends to delete this file). Typically used for short-lived files, but also can be used by malware as an alternative way of deleting files.
Modifying Operations	Create New	Creation of a new file or directory

Column Name	Attempted Operation	Description
Modifying Operations	Delete	Deletion of a file
Modifying Operations	Mmap Write	Write to a memory mapped file
Modifying Operations	Owner Change	Change the owner of a file or directory
Modifying Operations	Permission Change	Change the permissions on a file or directory
Modifying Operations	Rename	Rename a file or directory
Modifying Operations	Write	<p>Write the contents of a file; unlike other rules, “write” in the Expert Rules interface does not mean any modification at all. However, operations such as modifying the length of a file are also considered writes.</p> <p>Note: If you specify a rule that allows creation of new files but blocks writes to existing files, the agent will allow the process that created a new file to make modifications to that same file for a short period of time. Without this, the process that created the file could not write the initial content, and you would be left with a zero-byte file.</p>
Modifying Operations	Write Delayed	Memory mapped writes in which an application “maps” a file into memory (RAM) and writes to that memory. This content is later flushed to disk by the operating system’s paging mechanism.
Modifying Operations	Write Intent	This option is no longer functional in expert rules. It is scheduled to be removed from the interface in a future release.
Basic Operations	Cleanup	<p>Cleanup is the file system report that a process is done using a file; it means the file has been “closed”. This corresponds to IRP_MJ_CLEANUP; see https://msdn.microsoft.com/en-us/library/windows/hardware/ff548608(v=vs.85).aspx. In App Control, cleanup signals that a file is ready to be analyzed, and it also triggers file deletion in a “delete on close” operation.</p> <p>Note: Cleanup cannot be blocked since doing so would result in a handle leak. You can choose a reporting action, however.</p>
Basic Operations	Lock File	Lock the file that matches the rule.

Column Name	Attempted Operation	Description
Basic Operations	Mmap Read	Read a memory mapped file.
Basic Operations	Open	File open action.
Basic Operations	Open Execute Intent	A file handle was acquired with the intent to execute, but execution has not happened yet.
Basic Operations	Read	Read the contents of a file.

Table 68: Expert Memory Rules: Operation Settings

Column Name	Attempted Operation	Description
Basic Operations	Access Kernel Memory	Rules can use this operation to close a bypass on XP and 2003 systems that prevents usermode processes from opening <code>\Device\PhysicalMemory</code> , which effectively allows them to read kernel memory. Windows versions from Vista forward prevent this action on their own.
Basic Operations	Allocate Memory	Corresponds to the <code>VirtualAlloc</code> system call, which is invoked when an application wants to obtain a block of memory with specific permissions.
Basic Operations	Debug Process	Corresponds to <code>OB_OPERATION_PROCESS_PTRACE</code> , which is invoked if there is an attempt to enable ptrace logging on another application.
Basic Operations	Kill Process	Corresponds to <code>OB_OPERATION_PROCESS_KILL</code> , which is used to signify that someone opened a handle to another process/thread and attempted to terminate it.
Process/Thread Operations	Create Handle	This operation occurs if there is an attempt to open a new handle to a process/thread. Rules can be used to strip or report the permissions on the handle to limit what the source process can do on the target object. The permissions available here are the same as those documented for non-expert memory rules in Table 65, "Permissions Menu Options" , on page 480.

Column Name	Attempted Operation	Description
Process/Thread Operations	Duplicate Handle	This operation occurs if there is an attempt to duplicate a handle that is already open. Rules can strip or report the permissions on the new handle to limit what the process can do with the duplicate handle.

Table 69: Expert Registry Rules: Operation Settings

Column Name	Attempted Operation	Description
Key Operations	Create Key	Creation of a registry key
Key Operations	Rename Key	Not implemented. Do not use. Rename is a delete operation plus a create operation. If you want to block or report renaming of keys, you can take the action on either of those operations.
Key Operations	Delete Key	Deletion of a registry key
Key Operations	Set Security	SetSecurity is invoked when someone tries to change the permissions on a given registry key/value.
Key Operations	Open Key with Write Access	Open a registry key with write access.
Value Operations	Change Value	Change value is invoked when there is an attempt to modify a registry value. The target name of the operation is the full path to the registry value (e.g. HKLM\key\value).
Value Operations	Delete Value	Delete value is invoked when there is an attempt to delete a registry value. The target name of the operation is the full path to the registry value (e.g. HKLM\key\value).

Expert Rule Actions

[Table 70](#) shows the actions an Expert Rule can take if an operation matching the rule is attempted. Unlike operations, most (but not all) actions are common to all three rule pages. The Where Available column in the table shows whether the action is limited to one page.

As with non-expert rules, Expert Rules are often most effective in pairs. For example, one rule might tag certain types of files and another one might take a specified action, such as allowing execution, when files with that tag appear later. [“Tags and Tagging Actions in Expert Rules”](#) on page 499 provides more information about this feature.

With Expert Rules, you can also combine actions that might otherwise require two rules. For example, you can configure rule to "promote" a process so that files it writes are locally approved, and in the same rule, demote children of the process so that files they

write are not locally approved. When you review the table, look for actions that form this kind of pairing.

The table includes brief descriptions of what these actions do. Many of the actions are described in more detail in the "Custom Rules," "Memory Rules," and "Registry Rules" chapters in Using App Control, which is available as online help in the v8.0.0 App Control console or as a PDF download on the Carbon Black User Exchange.

Note

The Actions column does not currently show a value for every Expert Rule.

Mutually Exclusive Actions

Some of Expert Rule actions are mutually exclusive. If you choose one of the options in the following list, the others will be grayed out on the page:

- Allow, Block, Report, Prompt
- Promote process, Demote process
- Tag Target, Remove Target Tags, Tag Process, Remove Process Tags, Add Global Tags, Remove Global Tags
- Ignore, Dirty, Never Report, Track

Table 70: Action Settings in Expert Rules

Column Name	Setting Name	Description	Where
Approval Actions	Approve	Locally approve the target (file). Note: Currently, you cannot disable sending approval events in an Expert Rule. If you do not want an Expert Rule to send approval events, first create it as a non-expert rule, turn off <i>Send Approval Event</i> , save the rule, and then change it to an Expert Rule to finish your configuration.	Custom Rules
Approval Actions	Approve As installer	Locally approve the target (file) and mark it as an installer.	Custom Rules
Approval Actions	Demote process	Demote the process that performed the operation.	All
Approval Actions	Demote Target Process	Demote the target process.	Custom Rules
Approval Actions	Don't Promote Children	Do not promote child processes of the process that performed the operation; used when the process itself is promoted (see below).	All

Column Name	Setting Name	Description	Where
Approval Actions	Promote process	Promote the process that performed the operation, locally approving <i>files</i> written by this process; promote new <i>processes</i> spawned by this process unless “Don’t Promote children” was also chosen (see above).	All
Approval Actions	Promote Target Process	Promote the target process when this operation happens. Only applicable with the “Create process” operation.	Custom Rules
Approval Actions	Query Reputation	Ask server for the global state of the target (file) when this operation happens. This setting is for built-in rules and cannot be activated in new rules or changed in existing rules.	All
Authorization Actions	Allow	Allow the corresponding operations to go through. Note: You can create an Expert Rule that allows creation of new files but blocks writes to existing files. However, the agent will allow the process that created a new file to make more modifications to the same file for a short time. This is necessary to allow the same process to both create the new file and write the initial content to the file.	All
Authorization Actions	Block	Block the corresponding operation.	All
Authorization Actions	Prompt	Prompt the user to decide whether to allow or block the operation. A notifier must be selected when this action is chosen.	All
Authorization Actions	Report	Report (as an event) that the operation would have been blocked, but do not block it. For example, generate an event for all new files written by Powershell.	All
Authorization Actions	Suspend Source Process	Suspend the process that performed this operation. This is typically used for malware research where a researcher might want to inspect the process and see what it did (or was about to do) before it is terminated.	All
Authorization Actions	Terminate Source Process	Terminate the process that performed this operation	All
Tagging Actions	Tag Process	Tag the process; if chosen, one or more tags must be specified in the “Tags to Add/Remove” field	All

Column Name	Setting Name	Description	Where
Tagging Actions	Tag Target	Tag the target object; if chosen, one or more tags must be specified in the “Tags to Add/Remove” field	All
Tagging Actions	Remove Process Tags	Remove tags from the process; if chosen, one or more tags must be specified in the “Tags to Add/Remove” field	All
Tagging Actions	Remove Target Tags	Remove tags from the target object; if chosen, one or more tags must be specified in the “Tags to Add/Remove” field	All
Tagging Actions	Remove Global Tags	Remove global tags that other rules can test; if chosen, one or more tags must be specified in the “Tags to Add/Remove” field	All
Tagging Actions	Add Global Tags	Add global tags that other rules can test; if chosen, one or more tags must be specified in the “Tags to Add/Remove” field	All
File Tracking Actions	Dirty	Trigger re-analysis of the file matching the Path or File definition to see whether its hash has changed	Custom Rules
File Tracking Actions	Ignore	Do not track modifications	Custom Rules
File Tracking Actions	Never report	Keep an agent record of these operations but do not them to the server	Custom Rules
File Tracking Actions	Track	Track the file regardless of ignore rules	Custom Rules
Other Actions	Finish Rule Group	Skip other user-created rules but continue evaluating all built-in rules.	All
Other Actions	Report Execution (Deprecated)	Trigger meter on first execution events; this field is deprecated and appears only in details of an internal rule. It is read-only.	Custom Rules
Other Actions	Silent	Perform all assigned rule actions, but don't generate notifiers or report events.	All
Other Actions	Stop Rule Processing	Stop processing other rules after this rule is processed; this may improve performance. Note that Allow also stops processing but allows the action to continue.	All
Other Actions	Trigger Action	Trigger agent action where usermode sends an event in response to a kernel operation. For use with internal rules only.	All
Other Actions	Unenforceable	Indicate that some other action could not be enforced due to platform limitations. This field appears only when details of an internal rule are displayed. It is read-only.	All

Tags and Tagging Actions in Expert Rules

Tags are labels that can be applied to different objects tracked in App Control for as long as those objects exist. An "object" in this case can be a running process, a registry key, a file, an image, or the entire global 'system' that those processes run on. The global system is the computer the process is running on.

Each operation has an "initiator" process, the process that initiated it. Each operation also has a target object that the operation is being carried out on. Target objects vary depending on the type of operation. For example:

- For the "file write" operation, the target object will be a file.
- For a "process start" operation, the target will be another process.
- For a "registry value creation" operation, the target is a registry value. The behavior and lifespan of the tag depends on the *type* of object being tagged.

On the Add or Edit Rule page, the Tagging Actions column provides options for adding and removing tags when the other conditions of the rule are met. There are separate 'add' and 'remove' options for initiator processes, target processes, and the global system.

Tags are primarily useful when several rules related to the same tag(s) are created. Once a rule applies tags to an object, other rules can use these tags as a factor in determining whether a process matches the rule conditions, taking an action when a match is found. In other words, to use tags:

- Create a rule that applies a tag to an object.
- Create a separate rule that uses the presence of that tag as a condition for matching the rule; if it is testing the same operation as the tagging rule, rank this rule lower.

To use an Expert Rule to apply tags to an object:

1. On the Add Custom/Memory/Registry Rule page, provide a name for the rule.
2. If the rule name does not include the tag(s) you intend to use, consider putting them into the Description field. Although you cannot add a "tags" column on the rules table pages, you can display the description.
3. Choose **Expert** as the Rule Type (Custom Rules) or click **On** in the Expert Mode radio button field (Memory and Registry Rules).
4. In the Operations list, choose the operation(s) that should trigger this rule.

- In the Tagging Actions list, choose the object that you want to tag (Tag Target, Tag Process, Add Global Tags). When you choose one of these actions, the *Tags to Add/Remove* field is added below the list.

Actions

Authorization Actions	Approval Actions	Other Actions	Tagging Actions	File Tracking Actions
<input type="checkbox"/> Allow	<input type="checkbox"/> Promote process	<input type="checkbox"/> Trigger Action	<input checked="" type="checkbox"/> Tag Target	<input type="checkbox"/> Ignore
<input type="checkbox"/> Block	<input type="checkbox"/> Demote process	<input type="checkbox"/> Finish Rule Group	<input type="checkbox"/> Remove Target Tags	<input type="checkbox"/> Dirty
<input type="checkbox"/> Report	<input type="checkbox"/> Don't Promote Children	<input type="checkbox"/> Stop Rule Processing	<input type="checkbox"/> Tag Process	<input type="checkbox"/> Never report
<input type="checkbox"/> Prompt	<input type="checkbox"/> Approve as installer	<input type="checkbox"/> Silent	<input type="checkbox"/> Remove Process Tags	<input type="checkbox"/> Track
<input type="checkbox"/> Terminate Source Process	<input type="checkbox"/> Approve		<input type="checkbox"/> Add Global Tags	
<input type="checkbox"/> Suspend Source Process	<input type="checkbox"/> Promote Target Process		<input type="checkbox"/> Remove Global Tags	
	<input type="checkbox"/> Demote Target Process			

Tags To Add/Remove: ⓘ

Process Tag(s): ⓘ

Target Tag(s): ⓘ

Global Tag(s): ⓘ

Global Tag Exception(s): ⓘ

- In the Tags to Add/Remove field, enter the name(s) of the tag(s) you want to apply when the conditions of this rule are met. To add more than one tag, separate the tag names with commas.
- Provide any additional conditions for matching this rule, such as paths or files, the processes and any restrictions by user or policy.
- When you have finished specifying the rule, click the **Save & Exit** button.

To create an Expert Rule applied only to operations with a specific tag:

- Navigate to the table page for the type of rule you want to create (Custom, Memory or Registry).
- Click the add rule button.
- On the Add Rule page, provide a name for the rule.
- If the rule name does not include the tag(s) you intend to use, consider putting them into the Description field. Although you cannot add a "tags" column on the rules table pages, you can display the description. This will be helpful in pairing the rule that creates a particular tag with a rule that uses that tag to identify matching operations.
- Choose **Expert** as the Rule Type (Custom Rules) or click **On** in the Expert Mode radio button field (Memory and Registry Rules).
- In the Operations list, choose the operation(s) that should trigger this rule.
- In the Actions list, choose the action to perform when an operation matches the rule.

Note: You do not need to use one of the Tagging Actions in this case unless you are using one tag to create another one.
- Enter the names of the tags you want to match in the appropriate field(s):

- **Process Tag(s):** Enter tags here if you want to apply this rule when the process that *initiates an operation* has a matching tag.
- **Target Tag(s):** Enter tags here if you want to apply this rule when the process, file, or registry key that is the *target of an operation* has a matching tag.
- **Global Tag(s):** Enter tags here if you want to apply this rule when the 'global system' on which the operation is *being performed* has a matching tag. This is equivalent to the computer on which the operation is performed.
- **Global Tag Exceptions(s):** Enter tags here if you want to *exclude* 'global systems' with any of the matching tags from being subject to this rule.

Actions

Authorization Actions	Approval Actions	Other Actions	Tagging Actions	File Tracking Actions
<input type="checkbox"/> Allow	<input type="checkbox"/> Promote process	<input type="checkbox"/> Trigger Action	<input type="checkbox"/> Tag Target	<input type="checkbox"/> Ignore
<input type="checkbox"/> Block	<input type="checkbox"/> Demote process	<input type="checkbox"/> Finish Rule Group	<input type="checkbox"/> Remove Target Tags	<input type="checkbox"/> Dirty
<input type="checkbox"/> Report	<input type="checkbox"/> Don't Promote Children	<input type="checkbox"/> Stop Rule Processing	<input type="checkbox"/> Tag Process	<input type="checkbox"/> Never report
<input type="checkbox"/> Prompt	<input type="checkbox"/> Approve as installer	<input type="checkbox"/> Silent	<input type="checkbox"/> Remove Process Tags	<input type="checkbox"/> Track
<input type="checkbox"/> Terminate Source Process	<input type="checkbox"/> Approve		<input type="checkbox"/> Add Global Tags	
<input type="checkbox"/> Suspend Source Process	<input type="checkbox"/> Promote Target Process		<input type="checkbox"/> Remove Global Tags	
	<input type="checkbox"/> Demote Target Process			

Process Tag(s): i

Target Tag(s): i

Global Tag(s): i

Global Tag Exception(s): i

Path or File: i

9. Provide any additional conditions for matching this rule, such as paths, file names, processes, and any restrictions by user or policy.

10. When you have finished specifying the rule, click the **Save & Exit** button.

Tag Syntax Requirements

Tags must meet the following requirements and restrictions:

- Commas are used to separate tags when multiple tags are used; do not use commas in the tag name itself.
- Tags must have at least one non-numeric character.
- The prefixes "<Bit9:", "YaraTags" or "<LegacyClassification:" are reserved for use by Carbon Black and should not be used in a tag unless advised by Carbon Black Support or Services. See Built-in Tags on page 14 for more information.
- A tag and the process pattern of a rule (i.e., the pattern in any of the process fields) should not be the same. This helps avoid conflicts during rule processing.
- Avoid extremely long tag names. All of the fields in a rule combined must not exceed 2048 characters.
- Do not use the pipe (|) character.

Built-in Tags

Special tags are used by App Control in some Custom Rules provided in this release. Do not use these tags unless advised otherwise by Carbon Black support or services.

- **<YaraTags:tagname>** – Yara tags come from Yara rule content. When used to match the process that initiated an operation or the target process, they refer to the *file* that the process's image was loaded from. Yara tag values persist as part of the tracked file state, including across reboots (unlike user-created tags).

Direct customization of YARA rules is not supported in this release, and these tags are currently for use by Carbon Black only.

- **Bit9:tagname** – The 'Bit9:' prefix is used on tags built-in to App Control for various purposes. Although usable in other rules, they are intended only for rules provided by Carbon Black, and their behavior could change in later releases without notice.
- **<LegacyClassification:tagname>** – This prefix is used for internal, Carbon Black rules to identify older, hexadecimal tags. It should not be used in other rules.

Tag Persistence

User-created tags, for processes and for the global system, do not persist across reboots of an agent. The rule that attaches the tag must detect the operation it describes and reattach the tag before a rule that uses the tag can discover the tagged process. A tag may also be explicitly removed by a rule that has a "remove tag" action defined. There are other conditions that affect tags on different objects:

- **Process/thread tag** – Process and thread tags persist until the process instance dies. If the agent process (parity.exe) is restarted, then the tags would still be set. If the full system is restarted or if the kernel filter driver (parity.sys) is unloaded and reloaded, then a process would lose its classifications.
- **File tag** – Currently, a file tag lives only during a single operation.
- **Yara Tag** – Yara tags persist for the life the hash they apply to in the agent cache.
- **Global Tag** – Global tags persist until the agent process (parity.exe) is restarted.

Expert Rule Examples

Several of the default Custom Rules included in v8.0.0 are Expert Rules. You can examine the following rules to get ideas about the kind of rules you might choose to create:

- Examine powershell script contents
- Block powershell scripts that execute memory
- Do not treat these processes as .NET applications
- Report read-only memory map operations on unapproved executables by .NET applications
- [Sample] Prompt for read-only memory map operations on unapproved executables by .NET applications in medium enforcement
- [Sample] Deny read-only memory map operations on unapproved executables by .NET applications in high enforcement

- Deny read-only memory map operations on banned executables by .NET applications

Note

Registry or Memory Rules included by default in this release do not use Expert Mode.

Example: Allow Execution in a Folder when Visual Studio is Running

Perhaps you want to restrict executions in a specific folder, called `projectfolder` in the example here, so that they are allowed only when Visual Studio is running. This can be done using a series of Custom Rules to create, use, and remove that tag. Global tags essentially tag the entire environment on a computer indicating that anything happening on it matches the tag.

If you create a suite of rules like this, be sure to name them in a way that makes their relationship clear, and consider providing more information about their interactions in the Description field for each one. You can also further refine the rules with the other standard options, such as specifying user and/or policy.

1. Create one Custom Rule that applies a global tag when it detects Visual Studio running. For example:
 - **Operations:** Process Create
 - **Actions:** Add Global Tags
 - **Tags to Add/Remove:** VWrite2projectfolder
 - **Target:** devenv.exe
2. Create a second Custom Rule that allows execution in a specific folder when the global tag is set. For example:
 - **Operations:** Execute
 - **Actions:** Allow
 - **Global Tag(s):** VWrite2projectfolder
 - **Path or File:** <ProgramData>\projectfolder\
3. Create a third Custom Rule that removes the global tag when the Visual Studio process terminates. For example:
 - **Operations:** Process Terminate
 - **Actions:** Remove Global Tags
 - **Tags to Add/Remove:** VWrite2projectfolder
 - **Target:** devenv.exe

Example: Tag a Process and Report its Children

Perhaps you want to tag all processes that are launched by `svchost.exe` so that you can report when the child process are running. You can create a pair of rules for this purpose. Name the rules in a way that makes their relationship clear, and consider providing more information in the Description field for each one.

1. Create one Custom Rule that applies a tag to a process if it is the launched by `svchost.exe`. For example:
 - **Operations:** Process Create

- **Actions:** Tag Target
 - **Tags to Add/Remove:** childofsvchost
 - **Process:** svchost.exe
2. Create a second Custom Rule that reports creation of processes identified with the tag created in the previous rule.
- **Operations:** Process Create
 - **Actions:** Report
 - **Process Tag(s):** childofsvchost

Example: Promote an Installer and Demote its Children

Perhaps you want to promote an installer so that it can successfully install an application, but you do not want the application to be able to create files that are automatically approved. For example, you might want to allow installation of notepad++, but not have scripts created by notepad++.exe be approved based on this promotion:

- **Operations:** Write
- **Actions:** Promote process, Don't Promote Children
- **Path or File:** *\notepad++.exe
- **Process:** Any Process

Switching to or from Expert Rule Mode

You can change an existing rule to an Expert Rule if you want to add operations or actions not available in its current rule type. When you make this change, the menu choices you began with should be converted to the correct checkboxes in Expert Mode.

Important

Avoid changing rules from expert to non-expert mode. Many of the operations and actions in Expert Rules are not available to other rule types, and you could lose your rule definition entirely or convert the rule to something that does not match the operations or take the actions you want.

Chapter 18

Rapid Configs

This chapter describes Rapid Configs, which are sets of rules that can be used to accomplish tasks such as application optimization, operating system and application hardening, and approval of files delivered by software distribution systems.

Sections

Topic	Page
Overview	506
Rapid Config Details	510
Configuring and Enabling Rapid Configs	511
Specifying Notifiers for Rapid Configs	516
Automatic Rapid Config Updates	518

Overview

App Control includes a variety of individual rules that you can use to monitor or protect your endpoints. In most cases, these rules perform a narrowly defined task, such as reporting or blocking the execution of files at a location you specify.

A Rapid Config (short for “rapid configuration”) is a *set of rules* that can be used to achieve more complex goals such as optimizing the interaction of App Control and a specific application, hardening of operating systems and applications, and approval of files created or delivered by certain tools or pathways. For example, a Rapid Config for a specific application might include rules that ignore writes to specified folders while approving certain files so that the application can be more easily used.

A single Rapid Config might combine actions that you could have taken by creating several individual rules. On the other hand, it might be entirely composed of internal rules not exposed in the user interface, or it might be a combination of internal and visible rules.

Rapid Configs are created by Carbon Black and are built in to App Control. Although you cannot create a Rapid Config, some of them allow (and require) configuration so that they work properly in your environment. Other Rapid Configs are completely configured and require only that you enable them. You can enable any of these configurations to take advantage of their efficiency or protection features, or disable them if they interfere with activities you need to perform on the server or an endpoint.

To keep your Rapid Configs current, you can allow automatic updating by Carbon Black File Reputation. This feature is enabled by default when the connection between your App Control Server and Carbon Black File Reputation is enabled.

Rule Scope

If you choose, you can enable Rapid Configs that apply on all computers on a platform (e.g., all Windows computers) under all conditions. You also have the option to more narrowly focus the scope of a Rapid Config by specifying one or more of the following criteria (not all of these options are available for all Rapid Configs):

- **Policy-specific** – For all Rapid Configs, you can choose to limit a rule to *computers in specified policies*, and where you enter configuration information, you can configure one Rapid Config with different settings for different policies. Within a Rapid Config, you can also specify different settings for different policies.
- **User- or group-specific** – For certain Rapid Configs (e.g., the Visual Studio config), you can make the rule apply only to a specific *user or group of users*.
- **Conditional Macros** – You can use certain macros to restrict the conditions under which certain parameter values in Rapid Configs are applied. Only agents meeting the “test” described in the macro will use the parameter as a condition within the Rapid Config. Note that these conditional macros must be applied for each parameter you want to conditionalize. See [“Specifying Paths and Processes”](#) on page 517 for details.

Rapid Configs are located on a tab view on the Software Rules page.

Note

Some of the items listed as Updaters in previous releases are now on the Rapid Configs tab.

Viewing Rapid Configs

To view the table of Rapid Configs:

- On the console menu, choose **Rules > Software Rules** and click the **Rapid Configs** tab on the Software Rules page.

Software Rules							
Updaters Rapid Configs Publishers Users Directories Files Custom Memory Registry Scripts Reputation							
Saved Views: (The Current View Has Unsaved Changes - Discard)				Group By:			
(none) Delete Save Add				(none) Ascending			
Show Filters Show Columns Export to CSV Refresh Table							
Action Showing 27 out of 27 item(s)							
Select	Name	Enabled	Configured	Auto Detection	Platform	Modified By	P
<input type="checkbox"/>	WMI Protection	No	No	No	Windows	System	A
<input type="checkbox"/>	Windows Installer Embedded File Protection	No	No	No	Windows	System	A

Initially, this page shows the list of Rapid Configs built in to the App Control release you installed. [Table 71](#) lists the Rapid Configs available when this publication was completed. If you have automatic cloud update enabled or install a later release, you might see more Rapid Configs, and some of the Rapid Configs listed here might be changed.

Table 71: Rapid Configs

Configuration	Platform	Description
Browser Protection	Windows	Reports or prevents browsers from performing potentially malicious operations.
Carbon Black App Control Server Tamper Protection	Windows	Protects the App Control Server from tampering. Disabled by default, but enabling it is recommended for extra protection. It may be disabled later if necessary for troubleshooting purposes. Note: There is tamper protection built into the App Control <i>agent</i> , which is on by default. The Rapid Config on this page is for tamper protection on the <i>server</i> .
Carbon Black EDR Tamper Protection	Windows	Protects the Carbon Black EDR <i>sensor</i> from tampering. If you have both the App Control Agent and the Carbon Black EDR sensor installed on endpoints, enabling this updater provides extra protection.

Configuration	Platform	Description
Cryptomining Protection	Windows	Reports or prevents potentially malicious behavior related to file based cryptomining attacks. Minimum App Control agent version to use this Rapid Config is 8.0.0.
Delivery Optimization	Windows	Approve files written by the Delivery Optimization Service (DoSvc). This Rapid Config is not needed for agents running version 8.1 and later because files written by the Delivery Optimization Service will automatically be approved in those versions. Minimum App Control agent version to use this Rapid Config is 7.2.0.
Domain Controller Logon Scripts	Windows	Allows and promotes all files under the Sysvol and NetLogon directories of specified domain controllers if the machine is a member of the specified domain.
Doppelganger Protection	Windows	Protect against the exploit known as Doppelganging on Windows systems. Reference: https://community.carbonblack.com/docs/DOC-11212 . Minimum App Control agent version to use this Rapid Config is 8.0 P7.
Linux Hardening	Linux	Improves the security of computers running Linux by reporting or blocking modification of critical Linux system files.
Microsoft Exchange Server	Windows	Improves the performance of Microsoft Exchange servers when running along side App Control. Minimum App Control agent version to use this Rapid Config is 7.2.0.
Microsoft Office Protection	Windows	Improves security by watching for suspicious behavior by Microsoft Office apps, such as spawning of other applications or creating executable file types.
Microsoft SCCM	Windows	Approves software delivered via Microsoft SCCM. Optionally allows and promotes files you specify that are executed directly from SCCM distribution points.
Microsoft SQL Server	Windows	Improves the performance of Microsoft SQL servers when running alongside App Control. Minimum App Control agent version to use this Rapid Config is 7.2.0.

Configuration	Platform	Description
Mimikatz Protection	Windows	Protect against Mimikatz based attacks on windows systems. Mimikatz is a credential abuse tool effective at retrieving cleartext passwords, NTLM hashes, Kerberos Ticket Granting Tickets (TGT) and more. Developed by Benjamin Delpy to illustrate flaws within the Windows Authentication subsystem, it is a tool frequently used by malicious actors due to its reliability and efficiency. Several successful attacks leverage or mimic Mimikatz to dump credentials from memory, enabling actors to move laterally across systems using legitimate credentials - undetected. Minimum App Control agent version to use this Rapid Config is 8.1.0.
Powershell Protection	Windows	Improve security by watching for suspicious executions of Powershell.exe. Minimum App Control agent version to use this Rapid Config is 8.0.0.
Ransomware Protection (Most Types)	Windows	Protect against ransomware by reporting or blocking modification to files typically targeted by ransomware. Specifically targets types of ransomware that create temporary files then overwrite or delete the original. It does this by blocking renames or deletions of typically targeted files. Minimum App Control agent version to use this Rapid Config is 7.2.0.
Script Processors	Windows	Improves the security of computers by ensuring that script processors only run from expected locations. Minimum App Control agent version to use this Rapid Config is 8.0.0. Note: See Chapter 13, "Script Rules," for more information on the definition and control of scripts.
Self-Service Approvals	Windows	Provides a folder from which normal end-users can approve the execution of unapproved files even when in high enforcement. For more details on the benefits of this Rapid Config see this document: https://community.carbonblack.com/docs/DOC-4162 . Minimum App Control agent version to use this Rapid Config is 7.2.0.
Visual Studio	Windows	Approves Visual Studio builds and ignores intermediate build files.
Windows App Store	Windows	Approves Windows App Store installs and updates to specified directories.
Windows Hardening	Windows	Improves security of machines running Microsoft Windows.
WMI Protection	Windows	Protects against Windows Management Instrumentation (WMI) exploitation on windows systems. Minimum App Control agent version to use this Rapid Config is 8.0.0.

For App Control and Carbon Black EDR tamper protection configurations, your options are to enable or disable them and choose the policies to which they will be applied; no other changes can be made. Other Rapid Configs allow or require you to provide other parameters, such as paths and processes, that will specify how they work.

Note

When you choose the parameters for a Rapid Config, consider the potential number of matches and the volume of events that could be generated by those parameters. This is especially true for the Windows Hardening and Browser Protection Rapid Configs. To test the event volume for a particular configuration, consider first enabling the Rapid Config for a policy with a small number of computers before applying it generally.

Rapid Config Details

The table on the Rapid Config tab on the Software Rules page shows information about the purpose and status of each configuration. [Table 72](#) shows the fields that can be displayed in the Rapid Configs table. Most of these fields also appear on the Rapid Config Details page, although in some cases they have a different name.

Table 72: Rapid Config Table Fields

Field	Description
Name (Rapid Config Name on the Details page)	The name of the Rapid Config
Description	The operational description of the Rapid Config (i.e., the actions it takes)
Purpose	The goal or outcome that the Rapid Config is intended to accomplish
Enabled (Status on the Details page)	Is the configuration enabled
Configured (Table only)	Is the Rapid Config configured. Some Rapid Configs require configuration and some do not. You cannot enable a Rapid Config until all required fields are provided.
Auto Detection	This is an unused field that will be removed in a future release.
Platform	The operating system platform to which this configuration applies
Modified By (Table only)	The user that last modified the configuration; if the rule has not been changed, this will be System.
Policy (Applies To on the Details page)	The policies for which the configuration has been configured and enabled

Field	Description
CL Version (Table only)	The configuration list version in which the current definition of this configuration is included
Created By (Table only)	The user who created the configuration; in most cases, this will be System
Date Created	The date and time when this configuration was created
Date Modified	The date and time when this configuration was last modified
Date Upgraded	The date and time when this configuration was last upgraded. Initially, this will be the same as the Date Created. If cloud-based automatic updates are enabled, the date and time will change when updates are provided
Version	The version of the configuration; this begins at 1 and increments as significant updates are provided (or made during development prior to shipment)

Configuring and Enabling Rapid Configs

A Rapid Config must be configured before it can be enabled. The actions required for configuration vary depending upon the config:

- In some cases, all configuration settings are built in, and the only changes a user can make are to enable or disable the Rapid Config and change the policies to which it applies.
- If a Rapid Config has editable fields but all of these are either optional fields or already have values, you can enable the config immediately. If a Rapid Config has a setting that could block or report an action, the default is usually “Report”.
- If a Rapid Config has required fields that do not have defaults or values you previously entered, values must be entered into those fields before the config can be enabled.

The Rapid Configs table includes a column showing the configuration status of each config. Any config whose Configuration column shows *Yes* can be enabled directly from the table page. Rapid Configs whose Configuration column shows *No* cannot be selected on that page (their check boxes are grayed out), and must be configured on the Rapid Config Details page before being enabled.

The screenshot shows the 'Rapid Configs' tab in the 'Software Rules' section. The table displays the following configurations:

Configuration	Description	Enabled	Configured	Aut
Browser Protection	Reports or prevents potentially malicious behavior related to browsers. This includes execution of files downloaded by browsers, modification of the hosts file, and modification of browser related registry entries. Minimum Cb Protection agent version to use this Rapid Config is 8.0.0.	No	No	No
<input checked="" type="checkbox"/> Cb Protection Server Tamper Protection	Provides protection against tampering with the Cb Protection Server. Minimum Cb Protection agent version to use this Rapid Config is 7.2.0.	No	Yes	No
<input checked="" type="checkbox"/> Cb Response Tamper Protection	Prevents tampering with Cb Response. Minimum Cb Protection agent version to use this Rapid Config is 7.2.0.	No	Yes	No
<input type="checkbox"/> Cryptomining Protection	Reports or prevents potentially malicious behavior related to file based cryptomining attacks. Minimum Cb Protection agent version to use this Rapid Config is 8.0.0.	No	No	No
<input type="checkbox"/> Delivery Optimization	Approve files written by the Delivery Optimization Service (DoSvc). This Rapid Config is not needed for agents running version 8.1 and later because files written by the Delivery Optimization Service will automatically be approved in those versions.	No	Yes	No

To enable a configured Rapid Config from Rapid Config table page:

1. On the console menu, choose **Rules > Software Rules**.
2. Click the **Rapid Configs** tab to view configurations.
3. Check the box next to any config(s) you want to enable. Only boxes for configured Rapid Configs can be checked.
4. On the Action menu, choose **Enable Rapid Config**. The config is enabled for all policies, or if you are re-enabling a previously configured Rapid Config, it uses the policy settings choices you made before.

If you want to enable a Rapid Config that is already configured but you want to choose the policies it applies to, use the Rapid Config Details page.

To enable and select policies for a configured Rapid Config:

1. On the console menu, choose **Rules > Software Rules**.
2. Click the **Rapid Configs** tab to view configurations.
3. Click the View Details button next to the configuration you want to view or edit. The details you see vary among the available Rapid Configs.

Edit Rapid Config

Rapid Config Name: Microsoft Teams

Version: 3

Description: Approve Updates to Microsoft Teams. Minimum agent version to use this Rapid Config is 7.2.0.

Purpose: To avoid blocking updates to Microsoft Teams

Status: Enabled Disabled

Platform: Windows

Applies To: All Current and Future Policies

Date Created: Jul 1 2020 09:43:41 AM

Date Modified: Jul 1 2020 09:43:41 AM

Date Upgraded: Jul 1 2020 09:43:41 AM

▼ Rapid Config settings for All Current and Future Policies Delete settings for these policies...

Approve updates to Microsoft Teams i

***Approve Files Here:** i

***Approve Files Written By This Process:** i

***Process Publisher:** i

Settings Apply To: All Current and Future Policies
 Selected Policies

+ Add settings for additional policies

Save & Exit
Save
Cancel


4. If the Rapid Config you are enabling appears like the one above (i.e., does not have a second panel with text boxes to fill in), it is pre-configured. Click the **Enabled** button.

Note: If the Rapid Config Details page includes additional user-configurable fields, all mandatory fields (with a red asterisk) must be filled out before the config can be enabled. See “[User-Configured Rapid Configs](#)” on page 513 for more information.
5. In the Applies To field, click the radio button for **All Current and Future Policies** or **Selected policies**.
6. If you chose *Selected policies*, check the box next to each policy for which you want the Rapid Config to be enabled.
7. When you have finished selecting policies, click the **Save** button (to stay on the page) or **Save & Exit** button (to return to the table page) to save the enabled configuration.

User-Configured Rapid Configs

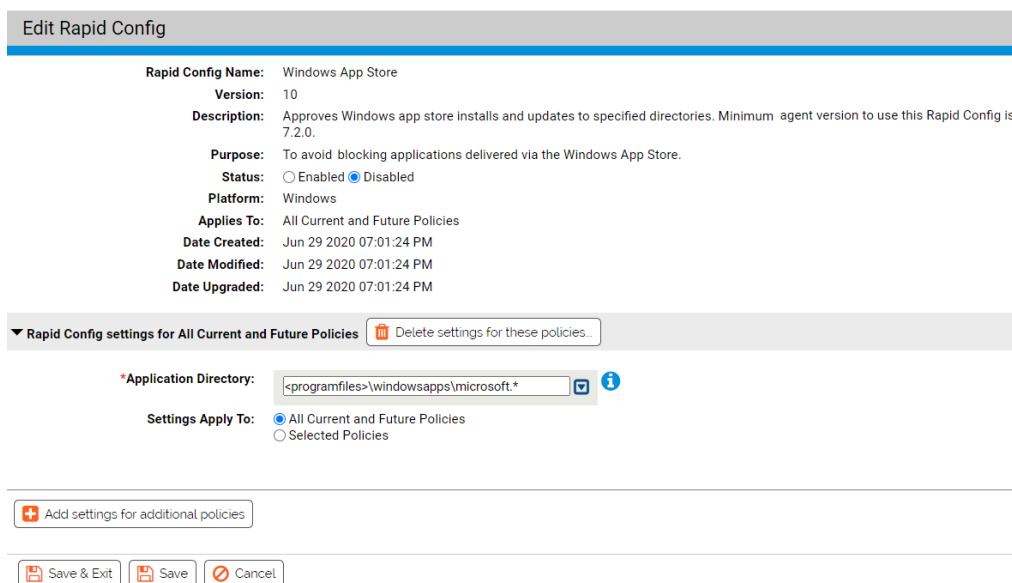
Some Rapid Configs require configuration beyond choosing the policies they apply to. They have additional panels that include a group of configuration settings, and the panels may allow you to choose which policies the *settings* apply to. If you choose to apply the settings to All policies, there is only one additional panel. If you choose specific policies, you can provide different setting values for different policies. You can also provide special settings for some policies and not apply the Rapid Config to other policies at all.

The specific parameters needed vary, but for each customizable configuration, any fields requiring user input have an information icon that displays popup help when you hover the mouse over it. Fields with a red asterisk are mandatory.



To configure and enable a policy-specific Rapid Config:

1. On the console menu, choose **Rules > Software Rules**.
2. Click the **Rapid Configs** tab to view configurations.
3. Click the View Details icon next to the configuration you want to view or edit. Rapid Config details vary.



If the Rapid Config you are enabling appears like the one above (it has a second panel with text boxes to fill in), it requires configuration before it can be enabled.

4. Click the **Enabled** button.
5. In the Rapid Config settings for All Current and Future Policies panel, hover over the information (i) icons next to each field to see the type of data to enter there. Any field with a red asterisk next to its name is mandatory.
6. Enter the values you want in each mandatory field and any other fields you choose. In some cases, a default value is provided, and you can leave that as the configuration if you choose.

7. In the Settings Apply To field, **All Current and Future Policies** is the default. If you leave that as the setting, you are finished configuring the Rapid Config and can **Save** or **Save & Exit**.
8. If you chose **Selected policies**, the panel name initially changes to New Rapid Config Settings Group. Check the box next to each policy to which you want this group of settings applied. The panel name changes to “Rapid Config settings for *<each policy you checked>*”.
9. If you want to configure a different group of settings for another group of policies, click **Add settings for additional policies** and repeat steps 5 through 8. Any policies covered by an existing settings group are not available when you are configuring a new settings group.
10. When you have finished selecting policies, click the **Save** button (to save the rule and stay on the page) or **Save & Exit** button (to save the rule and return to the table page).

Edit Rapid Config

Rapid Config Name: Windows App Store
Version: 10
Description: Approves Windows app store installs and updates to specified directories. Minimum agent version to use this Rapid Config is 7.2.0.
Purpose: To avoid blocking applications delivered via the Windows App Store.
Status: Enabled Disabled
Platform: Windows
Applies To: APC-79895-uorrtr, Default Policy, Help Desk
Date Created: Jun 29 2020 07:01:24 PM
Date Modified: Jun 29 2020 07:01:24 PM
Date Upgraded: Jun 29 2020 07:01:24 PM

▼ Rapid Config settings for APC-79895-uorrtr, Default Policy Delete settings for these policies...

***Application Directory:** 🔍 ⓘ

Settings Apply To: All Current and Future Policies
 Selected Policies

Policy

APC-79895-uorrtr

Default Policy

▼ Rapid Config settings for Help Desk Delete settings for these policies...

***Application Directory:** 🔍 ⓘ

Settings Apply To: All Current and Future Policies
 Selected Policies

Policy

Help Desk

+ Add settings for additional policies

Save & Exit

Save

Cancel

If you later decide to delete a group of settings, you can use click the View Details icon next to this Rapid Config on the Rapid Config table page, and in the Rapid Config Details page, click **Delete settings for these policies**, and then save the change. The policies affected by that group of settings are no longer affected by this Rapid Config.

Note

If a group of settings has the All Current and Future Policies button activated, clicking **Add settings for additional policies** displays an error dialog. You must deselect at least one policy before creating a new group of settings.

Specifying Notifiers for Rapid Configs

Some rules within Rapid Configs can block actions a user takes. For example, several of the rules that make up the Browser Protection Rapid Config can block:

- Execution of applications *by* browsers
- Execution of applications that were *downloaded from* browsers
- Registry modifications
- Host file modifications

For each one of these actions, you can specify files and paths affected, and you can choose to do nothing, to report the action, or to block the actions matching those settings. If you choose to block them, a Notifier field appears in the panel. That field provides a menu of existing notifiers from which you can choose the appropriate one for each action. When a Rapid Config contains more than one action that can be blocked, you can choose different notifiers for each action you block or use the same one for all. You also can choose Block for some actions and Report or Do Nothing for others.

Edit Rapid Config

Rapid Config Name: Browser Protection
Version: 18
Description: Reports or prevents potentially malicious behavior related to browsers. This includes execution of files downloaded by browsers, modification of the hosts file, and modification of browser related registry entries. Minimum agent version to use this Rapid Config is 8.0.0.
Purpose: To defend against malicious actors who exploit browser vulnerabilities to attack your enterprise.
Status: Enabled Disabled
Platform: Windows
Applies To: All Current and Future Policies
Date Created: Jun 29 2020 07:01:24 PM
Date Modified: Jun 29 2020 07:01:24 PM
Date Upgraded: Jun 29 2020 07:01:24 PM

▼ **Rapid Config settings for All Current and Future Policies** Delete settings for these policies...

Executables ⓘ

*Report Or Block Execution Of Applications By Browsers: Do Nothing Report Block ⓘ

Executable Files To Block:

+ Add ⓘ
 Java.exe
 Javaw.exe - Remove

Notifier: Enforce custom (file and path) rules ⓘ

Files That Should Not Be Blocked: ⓘ

Downloaded Executables ⓘ

*Report Or Block Execution Of Executables Created By Browsers: Do Nothing Report Block ⓘ

Executable Files To Block:

+ Add ⓘ
 *.bat
 *.cmd
 *.com
 *.dll - Remove

Notifier: New Notifier 1 ⓘ

Files That Should Not Be Blocked:

+ Add ⓘ
 ^\google\chrome\user data\pepperflash*\pepflas
 ^\google\chrome\user data\swreporter*\software - Remove

Specifying Paths and Processes

When you specify Path or File in a Rapid Config, you have some of the options that are available in Custom, Registry, and Memory rules. These include:

- **Specify a directory or a file/process** – You can enter a path or process specification that exactly identifies a file by path and name so that only that file matches the rule. You also can enter a specification that identifies a directory, and so affects all files or processes in that directory and its subdirectories.
- **Specify a local drive or UNC path (Windows only)** – You can use a local drive name, such as *C:\folder1\subfolder\application.exe*, to identify a local path or process. For a remote path or process, use a UNC path, such as *\\computer\dir\app.exe*. Mapped drives in a path or process specification are not recognized.
- **Use wildcards** – You can use wildcards ('?' for any one character and '*' for zero or more characters) to expand the scope of a path or process specification, or to help

- you match a file or folder whose exact location you don't know. Wildcards may be used at the beginning, end or middle of a path.
- **Specify multiple paths or processes** – For some paths and processes, you can add more than one path definition per rule.
 - **Use path macros** – You can use special macros to identify certain well known folders, even if you don't know their exact location on agent computers. Macros are platform-specific.
 - **Use conditional macros** – If you use conditional macros (such as *OnlyIf*), the condition you set applies only to the specific parameter with the macro. For example, if you set the Browser Protection Rapid Config to report executions of Java.exe, Javaw.exe, and FlashPlayerApp.exe, plus you added the following parameter: `<OnlyIf:HostName:Desktop-1>MyApp.exe`, the Rapid Config would report execution of MyApp.exe only on *Desktop-1* but it would report execution of Java.exe, Javaw.exe, and FlashPlayerApp.exe on any machine. To apply make all executables conditional, you would have to add the OnlyIf macro to each one.

See [“Specifying Paths and Processes”](#) on page 407 for more information about options and requirements for path and process specification. Note, however, that some of the options described in that section do not apply to Rapid Configs.

Note

If you specify a script file in a Rapid Config that controls execution, the config will not recognize the script and control its execution unless there is a corresponding Script Rule for the file extension and the process that executes the script. See [Chapter 13, “Script Rules,”](#) for more details.

Automatic Rapid Config Updates


New Rapid Configs may become available, or existing configs may be modified with enhancements or changed because of changes to the applications they apply to. When you install a new version of App Control, the Rapid Config table is updated to reflect these changes (if any).

By allowing Carbon Black File Reputation to maintain Rapid Configs, you get new and modified versions as soon as they become available, without waiting for a new release of App Control. Enabling Carbon Black File Reputation updates also means that obsolete configurations are deleted from the table. This feature is enabled by default if you have Carbon Black File Reputation enabled.

If you have not activated the integration between the App Control Server and the Carbon Black File Reputation, you can do that on the Licensing tab of the System Configuration page. See [“Activating Carbon Black File Reputation”](#) on page 763 for details.

To enable or disable cloud updates of Rapid Configs:

1. On the console menu, choose **System Configuration** on the Administration (Gear) menu.
2. On the System Configuration page, click the **Advanced Options** tab. The Advanced Options Configuration page appears, with the Software Rules Options panel at the bottom.
3. At the bottom of the page, click the **Edit** button.
4. In the Software Rule Options panel, the Carbon Black File Reputation updater option is enabled by default:
 - a. If you *do not want* Carbon Black File Reputation to keep your updaters current, *un-check* the box next to *Automatically update Rapid Configs and Updaters from Carbon Black File Reputation* and then click the **Update** button at the bottom of the page.
 - b. If you want to *re-enable* automatic updates from Carbon Black File Reputation after they have been disabled, check the box and click the **Update** button.



5. In the Confirm Server Setting Change dialog, click **Yes** to save your changes.

Alerts for Rapid Config Changes from the Cloud

You can enable an alert that will notify you each time a Rapid Config is created, modified, or deleted from the cloud. This is recommended if you have automatic updates enabled.

To enable alerts for Rapid Config updates delivered from the cloud:

1. On the console menu, choose **Tools > Alerts**.
2. Check the box next to the Rapid Config Alert.
3. On the Action menu, choose **Enable Alerts**.

Chapter 19

Event Rules

This chapter describes Event Rules, which allow you to specify an action to be performed when an event matches filters you define.

Sections

Topic	Page
Overview	521
Enabling, Disabling, and Deleting Event Rules	522
Disabling Processing of All Event Rules	523
Testing a Rule before Enabling	524
Creating and Editing Event Rules	525
Sample Event Rules	536

Overview

Event Rules allow you to specify an action to be performed when a file- or computer-related event occurs that matches filters you define. To use this feature, a console user must have *Manage event rules* permission. See [“User Role Permissions”](#) on page 106.

You can create an alert that reports when a specified event rule is triggered. See [“Creating Alerts”](#) on page 612.

Events That Can Trigger Rule Actions

Only events that relate to files or computers can be used to trigger an event rule. Each rule is required to have one event subtype specified; for example, the rule might specify that its action is triggered when an event with the subtype *New file on network* occurs. You may add more subtypes so that the rule takes action under several different event conditions. You also may add other specifications to the rule, such as that the event included a reference to a particular IP address.

You also may add specifications that the rule only runs when the target file identified in the event, or its parent process, has certain properties. For example, you might specify that a new, unapproved file is uploaded to an analysis service only if it does not have approval by reputation enabled.

Actions A Rule Can Take

The following actions can be taken using Event Rules:

- **Change global file state** – An Event Rule can create a global Approval, Ban or Report Ban, and can remove a global Approval or Ban for a file referenced in an event. This may be done for all computers or by policy. Rules that change global state may also be configured to resolve related approval requests from endpoint users.
- **Change global process state** – An Event Rule can create a global Approval, Ban or Report Ban, and can remove a global Approval or Ban for the file of the process referenced in an event. This may be done for all computers or by policy. Rules that change global state may also be configured to resolve related approval requests from endpoint users.
- **Change local file state** – An Event Rule can create or remove a Local Approval for a file referenced in an event. Rules that change local state may also be configured to resolve related approval requests from endpoint users.
- **Upload file** – An Event Rule can initiate upload of a file referenced in an event to the App Control Server.
- **Delete File** – An Event Rule can delete a file referenced in an event.
- **Analyze file** – An Event Rule can initiate upload of a file to any analysis service configured through the App Control Connector.
- **Move computer** – A computer referenced in a file-related event may be moved to a different policy and Enforcement Level.

Users will only see action options for which they have permission. For example, users without permission to submit files for analysis will not see the *Analyze file* option.

Simulating the Effect of a Rule

An important feature of Event Rules is the ability to simulate what would happen if you fully enabled a rule without actually taking the action specified. Event rules can have a significant impact on the App Control Server, and if not configured properly, they may have undesirable and unintended results. Because of this, it is strongly recommended that any new rule be run in *Simulate only* mode before it is fully enabled – this is one of the options on the Add and Edit Event Rule pages. See “[Testing a Rule before Enabling](#)” on page 524 for a recommended work-flow using Simulate only.

Re-Applying a Rule to Past Events

App Control also provides the ability to apply a new rule to past events. This can be useful in combination with Simulate only mode, allowing you to apply the rule to a larger set of past events to see the events that *would have been* processed by the rule. You can then review these results, and you may choose to fine tune the rule to reduce the conditions under which the rule is triggered. You might also re-apply a new rule to past events when it is fully enabled if, for example, you want to send all new, unapproved files that have appeared in the past week to an external service for analysis.

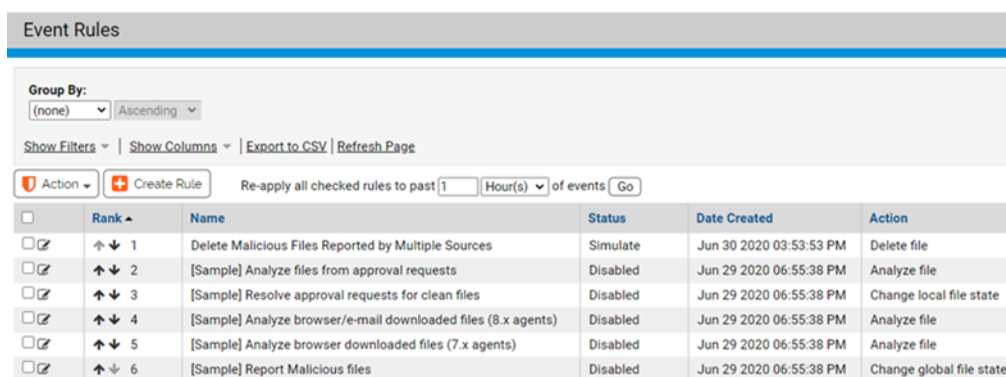
Enabling, Disabling, and Deleting Event Rules

You can enable, disable, or delete specific Event Rules on either the Event Rules (table) page or the Edit Event Rule page.

On the Event Rules page, you can select one or more rules and either enable or disable them using the Action menu. Note that you cannot choose the Simulate Only option on this menu. To enable a rule in Simulate Only mode, use the Edit Event Rule page.

To enable or disable rules in the Event Rules table:

1. On the console menu, choose **Rules > Event Rules**. The Event Rules page appears, showing the available rules and their status.



	Rank	Name	Status	Date Created	Action
<input type="checkbox"/>	1	Delete Malicious Files Reported by Multiple Sources	Simulate	Jun 30 2020 03:53:53 PM	Delete file
<input type="checkbox"/>	2	[Sample] Analyze files from approval requests	Disabled	Jun 29 2020 06:55:38 PM	Analyze file
<input type="checkbox"/>	3	[Sample] Resolve approval requests for clean files	Disabled	Jun 29 2020 06:55:38 PM	Change local file state
<input type="checkbox"/>	4	[Sample] Analyze browser/e-mail downloaded files (8.x agents)	Disabled	Jun 29 2020 06:55:38 PM	Analyze file
<input type="checkbox"/>	5	[Sample] Analyze browser downloaded files (7.x agents)	Disabled	Jun 29 2020 06:55:38 PM	Analyze file
<input type="checkbox"/>	6	[Sample] Report Malicious files	Disabled	Jun 29 2020 06:55:38 PM	Change global file state

2. Check the box next to one or more rules in the table.
3. On the Action menu, choose **Enable** or **Disable** and confirm your choice in the confirmation dialog. The checked rules are enabled or disabled according to your menu choice.

To enable or disable a single Event Rule, you can use the Edit Event Rule page. This is also where you can activate *Simulate only* mode, which allows you to see what the effect of an event rule *would be* without having it take the action specified in its definition.

To enable, disable or simulate the effect of an Event Rule:

1. On the console menu, choose **Rules > Event Rules**. The Event Rules page appears, showing the available rules and their status.
2. Click the View Details button next to the rule whose status you want to change. The Edit Event Rule page opens.

Edit Event Rule

General

Rule Name: [Sample] Analyze files from approval requests

Description: Automatically analyze all files related to approval requests with WildFire

Status: Enabled Simulate only Disabled

Select Event Properties

3. In the Status field, click the radio button for **Enable**, **Simulate Only** or **Disable**.
4. Make any other changes to the rule properties and click **Save** to stay on the page to monitor events processed by the rule or **Save & Exit** to leave the page. See [“Testing a Rule before Enabling”](#) on page 524 for a sample work flow for *Simulate only* rules.

Note


If an event rule depends on a particular configuration of an analysis tool, such as an analysis environment with a specific operating system, and if that environment becomes unavailable, the rule will be disabled automatically after waiting several minutes for the environment to become available again.

To delete event rules in the Event Rules table:

1. On the console menu, choose **Rules > Event Rules**. The Event Rules page appears, showing the available rules and their status.
2. Check the box next to one or more rules in the table.
3. On the Action menu, choose **Delete** and confirm your choice in the dialog. The checked rules are deleted.

Disabling Processing of All Event Rules

By default, the App Control Server is configured to process any enabled event rules. This means that the Enabled/Disabled/Simulate setting for each rule determines how and whether that rule functions. However, you can disable the event rule feature so that no event rules are processed. The checkbox for disabling and re-enabling event rule processing is on the **Advanced Options** tab of the **System Configuration** page. When you disable event rules, any alerts based on event rules will not be triggered.



To disable all event rule processing:

1. On the console menu, click the configuration (gear) icon and choose **System Configuration** and click on the **Advanced Options** tab.
2. Click the **Edit** button at the bottom of page.
3. In the Software Rule Options panel, *un-check* the Event Rules checkbox, then click the **Update** button.

To re-enable event rules, follow the same steps, except check the Event Rules box. Re-enabling event rules causes them to continue processing each rule from the point it was stopped. This can be useful when infrastructure activities would prevent event rules from completing their actions. For example, if you had many analysis rules that required access to a connected appliance and the appliance was down for maintenance, you could disable rule processing until the maintenance is completed.

Testing a Rule before Enabling

Event rule actions, such as bans, approvals and moving computers between policies, can cause serious security and operational issues if not properly configured. Because of this, it is strongly recommended that any new rule be run in *Simulate only* mode before it is fully enabled – this is one of the options on the Add and Edit Event Rule pages.

When you run an event rule in Simulate only mode, you can apply the rule to past notifications and view the events that *would have been* processed by the rule. You can then review these results, and you may choose to add or change filters for the rule to reduce the conditions under which the rule is triggered. If you open the Edit Event Rule pages for the sample rules, you can see some of the ways filters have been used to limit events processed by the rule. “[Sample Event Rules](#)” on page 536 describes these rules.

To test the effects of a rule with Simulate only mode:

1. On the console menu, choose **Rules > Event Rules**. The Event rules page appears, showing the available rules and their status.
2. Click on the View Details button next to the rule you want to test. The Edit Event Rule page for that rule opens.
3. Examine the configuration of the rule, changing it if necessary.
4. In the Status field, check the **Simulate only** radio button.
5. Make any other needed changes to the rule and click the **Save** button
Note: You must remain on the Event Rule details page to complete this process, so do not click the *Create & Exit* button.
6. On the Advanced menu to the right of the page, click on **Re-apply rule** choose a time period in the dialog box. This determines the window of past events the rule will be applied to. Depending upon the volume of matching events, you might want to limit the initial test to a short period, such as **1 day**. Choose the time period and click **Go**.

7. Continue to monitor the page, periodically clicking **Refresh Page** in the Processed Events panel until the Last Processed Event field in the History panel shows no more events to process. See [“Event Rule History and Processed Events List”](#) on page 534 for an example of the information shown in the Processed Events panel.
8. If you don't see events you expected to appear in the Processed Events panel, or if you see more or different events than expected, modify the rule accordingly, click **Save** again, and reapply the rule. If you followed the steps above, events related to the rule appear in the table of events with a *Simulated* in the Status field.
9. To simulate the rule for a longer period, change the *Re-apply rule* value and click **Go**.
10. Once you see events you expect and determine that the rule has no negative effects, change rule Status to **Enabled** and click **Save & Exit**. The rule is executed on new events – use the *Re-apply* menu if you want the rule to run actively on past events.

Creating and Editing Event Rules

You can create a new event rule by copying and modifying the settings of an existing rule or by creating the rule from scratch. In either case, you would need to provide at least the non-optional information shown in bold in the left column:

General Description	Section in the Add/Edit Event Rule Page
If a file- or computer-related event matches this/these criteria...	Select Event Properties
...and a file referenced in the event matches this/these criteria (optional)...	Select File Properties
...and the process referenced in the event matches this/these criteria (optional)...	Select Process Properties
... then take the following action...	Select Action
... on computers in this/these policy(ies) ...	Select Action/Create For:

The Select Event Properties, Select File Properties, and Select Process Properties sections can include multiple criteria for triggering the rule, and the Select Action section has different fields depending on the action you choose.

To add (create) an event rule:

1. On the console menu, choose **Rules > Event Rules**. The Event Rules page opens.
2. On the Event Rules page, choose **Create Rule**. The Create Event Rule page opens. [Table 73, “Event Rule Fields”](#), on page 528 describes the settings on this page.

Create Event Rule

General

Copy Settings From: (none) ▼

Rule Name:

Description:

Status: Enabled Simulate only Disabled

Select Event Properties

Add filter ▼

Subtype is ▼ ▼ +

Select File Properties

Add filter ▼

Select Process Properties

Add filter ▼

Select Action

Action: Change global file state ▼

Change Global State: Approve Ban Ban (Report Only) Remove Approval or Ban

Resolve Related Approval Request:

Create For: All Current and Future Policies Selected policies

Create & Exit
 Save
 Cancel

3. If there is an existing event rule similar to the one you want to create, choose that rule on the Copy Settings From menu. When you choose anything but (none) on this menu, the page is pre-populated with the fields from the rule you chose, and you need only change the fields that differ from the rule you copied from.
4. In the Rule Name field, provide a unique name for the rule. If you copied settings from an existing rule, the default name is that rule’s name followed by “(Copy)”.
5. In the Description field, you can provide a longer description of the rule if you choose. (Completing this field is optional).
6. In the Status field, you choose one of the following:
 - **Enabled** – Actions specified by the rule will be executed as specified.
 - **Simulate only** – Actions specified by the rule will be simulated. Events will be generated indicating what the rule would have done if enabled, but the actions specified will not actually be taken.

- **Disabled** – The rule and its settings will be saved but it will not execute or simulate the actions specified.

Important

Simulate only is strongly advised for a new event rule. See [“Testing a Rule before Enabling”](#) on page 524 for more about this Status choice.

7. In the Select Event Properties panel, use the Add filter menu to choose one or more event properties. For these filters:
 - At least one Subtype filter must be included.
 - Because only file- or computer-related events may be used to trigger an event rule, the selections on this menu are limited accordingly.
 - Some file-related properties that appear in events are not included here because they appear on the File Properties menu.
 - To use file names or path names in an event rule filter, specify them using the Event Properties filter rather than File Properties filter. The Event Property *File name* usually matches more of the relevant events than the File Property *First seen name*.
8. In the Select File Properties panel, use the Add filter menu to choose one or more file properties with which to further refine the conditions under which this rule will be triggered. Most of the choices here are the same as the fields in the App Control File Catalog, although there are some additional fields. See [“File and Process Properties in Event Rule Definitions”](#) for detailed information about certain choices in this panel.

Note

For both Select File Properties and Select Process Properties, if you choose an Extension filter, you must use the file extension *without* the initial dot (for example, *bat*, not *.bat*). Otherwise the rule will not function properly.

9. In the Select Process Properties panel, use the Add filter menu to choose one or more process properties with which to further refine the conditions under which this rule will be triggered. Most of the choices are the same as the fields in the App Control File Catalog, although there are some additional fields. See [“File and Process Properties in Event Rule Definitions”](#) for detailed information about certain choices in this panel.

Note

The process to which this configuration choice applies is the parent process of the file referenced in the event or event rule, not the process that appears in the operating system task manager when a file executes.

10. In the Select Action panel, use the Action menu to choose the action that will be taken when events and files match this rule. The options that appear on this menu depend upon the permissions of the console user creating or editing the rule – see [“User Role Permissions”](#) on page 106. The Action choices are listed in [Table 73, “Event Rule Fields”](#), on page 528.

11. If you choose an action that changes the state of a file, you can automatically resolve any approval request for the file. To do this, check the **Resolve Related Approval Request** box. If you do not check the box, any approval request for the related file will be left open until you manually close it. This box has no effect if there is not a related approval request. See [“Approval Requests and Justifications”](#) on page 566 for more on how approval requests are submitted and resolved.
12. When you have completed the rule definition, click **Save** to remain on the page, and follow the steps described in [“Enabling, Disabling, and Deleting Event Rules”](#).
-or-
To create the rule and leave the Create Event Rule page, click **Create & Exit**.

[Table 73](#) shows the fields available on the Create/Edit Event Rule page.

Table 73: Event Rule Fields

Panel:Field	Description
Copy Settings From:	Existing rule from which this rule copies its initial settings. If you do not want to copy any settings, leave the default value of (none).
Rule Name	Name by which this rule is identified. (Required)
Description	Additional information about the rule. This can be any text you choose to enter. (Optional)
Status	Radio buttons that determine whether and how this rule is activated: <ul style="list-style-type: none"> • Enabled – Actions specified by the rule will be executed as specified. • Simulate only – Actions specified by the rule will be simulated. Events will be generated indicating what the rule would have done if enabled, but the actions specified will not actually be taken. This is the default value for newly created rules. • Disabled – The rule and its settings will be saved but it will not execute or simulate the actions specified. This is the default value for the sample rules.
Select Event Properties: Add Filter	The properties of the event that triggers this rule: <ul style="list-style-type: none"> • Subtype – At least one event Subtype filter must be included in this filter (For example, <i>New file on network</i>). Additional Subtypes may be added so that, for example, a rule is triggered for either <i>New file on network</i> or <i>New unapproved file to computer</i> events. • Other Event properties – Other properties may be added to this filter. Some file-related properties that appear in events are not included here because they appear on the File Properties menu.

Panel:Field	Description
Select File Properties: Add Filter	<p>File properties to further refine the conditions for triggering this rule. Most of the choices here are the same as the fields in the App Control File Catalog. See “File and Process Properties in Event Rule Definitions” on page 532 for detailed information about certain choices in this panel. File properties are not required in an Event Rule.</p> <p>Note: If you specify a file property and that property is unavailable, the rule cannot be executed, and events matching the rule are placed in a Pending state until the property becomes available. For example, if you specify that a rule that requires that the Carbon Black File Reputation reputation for a file has a Trust level of 5 or less, if Carbon Black File Reputation is not configured and there is no trust information for the file, the rule will not be executed, even if all other rule specifications are met. This also applies to file prevalence and metadata.</p>
Select Process Properties: Add Filter	<p>Process properties to further refine the conditions for triggering this rule.</p> <p>Most of the choices here are the same as the fields in the App Control File Catalog. See “File and Process Properties in Event Rule Definitions” on page 532 for detailed information about certain choices in this panel. Process properties are not required in an Event Rule.</p> <p>Note: If you specify a process property and that property is unavailable, the rule cannot be executed, and events matching the rule are placed in a Pending state until the property becomes available. For example, if you specify that a rule that requires that the Carbon Black File Reputation data for a file shows a Trust level of 5 or less, if Carbon Black File Reputation is not configured and there is no trust information for the file, the rule will not be executed, even if all other rule specifications are met. This also applies to file prevalence and metadata.</p>

Panel:Field	Description
Select Action: Action	<p>The following options appear on the Action menu:</p> <ul style="list-style-type: none"> <p>Change global file state – This automatically changes the global state of matching files. You can approve, ban, or create a report-only ban for matching files. You can also remove approvals or bans. You also can apply the state change to All policies or selected policies.</p> <p>• Change global process state – This automatically changes the global file state of matching processes. You can approve, ban or create a report-only ban for matching processes. You can also remove approvals or bans. You also can apply the state change to All policies or selected policies.</p> <p>• Change local file state – This automatically changes the local state of matching files. You can locally Approve matching files or Remove local approval.</p> <p>• Upload file - This initiates an upload to the App Control Server of matching files from the agent-managed computer on which they appear. You can choose the default upload location or a custom location on the server or another accessible computer. For example, you can send all newly found files to a specific folder for manual examination or scanning by a tool on a different computer. Note: This option is available only for console users with one or both <i>Manage uploads of inventoried files</i> permission. See “User Role Permissions” on page 106.</p> <p>• Delete file – This initiates a request to delete the files referenced in the event that triggered to Event Rule. Deletion of the files on the endpoint is completed soon after the request is sent, the exact latency depending upon how many files are affected by the request and any other activities scheduled on the server. Important: When a file is deleted from an endpoint in this way, it is permanently deleted. It is not sent to a “recycle bin” or other location that allows it to be restored. See Chapter 9, “Deleting Files,” for more information about this feature.</p> <p>• Analyze file – This initiates upload of a file to a connected device or service for analysis when the rule conditions are met. You check the box for one or more enabled analysis services integrated with the App Control Server through the App Control Connector. If no services are configured, this option does not appear.</p> <p>• Move computer – This moves the computer referenced in the event to a different policy, with the following options:</p> <ul style="list-style-type: none"> <p>Specify policy – This displays a menu of the policies available on this App Control Server.</p> <p>Restore to normal enforcement level – This returns a computer that is in Local Approval mode to its previous policy. If the computer is not in Local Approval mode, this has no effect.</p> <p>Local approval – This moves a computer into Local Approval mode. See “Moving Computers to Local Approval Mode” on page 294 for details.</p> <p>Automatic policy – This moves a computer into the policy to which Active Directory mapping assigns it. If AD Mapping is not enabled, this setting has no effect.</p>
Resolve Related Approval Request	<p>When the Action choice for the rule is Change Global file state or Change local file state, this checkbox is displayed. If the box is checked, any approval request related to the file referenced in this file has its status changed to Resolved.</p>

Panel:Field	Description
Priority	When the Action choice for a rule is Upload file or Analyze file, you can set the priority for the upload or analysis to Low, Medium, or High, which determines the order in which the action is taken relative to other upload or analyze requests. Priority can be changed on the Requested Files page once a request is in progress.

Editing an Event Rule

You can edit existing event rules, modifying the fields described in [Table 73, “Event Rule Fields”](#) on page 528. However, you cannot change the Action setting for a rule once it is created; different actions may require different App Control Console user account permission, and also, rule history might not make sense if the rule recorded a mix of different actions. If you need to change the Action, create a new rule. You can use the *Copy Settings from* field to copy most of an existing rule’s definitions and then change the action before saving. Note also that you can change some of the options underneath an Action (such as changing which policies it applies to or changing Approval Request settings).

Edit Event Rule Page Menus

The Edit Event Rule page has two menus on the right side of the page. The Related Views menu has one or more of the following commands (which vary depending upon the Action chosen for the rule):

- **All file rules created by this rule** – Displays the Software Rules: Files Approvals and Bans page filtered to show file rules created by this event rule (does not include local file approvals, which are not tracked on this page)
- **All file uploads created by this rule** – Displays the Requested Files: Uploaded Files page filtered to show uploads initiated by this rule
- **All file analysis submissions created by this rule**-- Displays the Requested Files: Analyzed Files page filtered to show analysis submissions to analysis services configured through the App Control Connector
- **Related events** – Displays the Events page, filtered by this rule name

The Action menu includes one or more of the following commands:

- **Cancel all file analysis submissions created by this rule** – For file analysis rules, cancels all unprocessed file submissions made to analysis services configured through the App Control Connector. This has no effect if a file submitted because of this rule has already been sent to the analysis service.
- **Cancel all file uploads created by this rule** – For file upload rules, cancels all unprocessed file uploads initiated by the rule. This has no effect if a file has already been uploaded.
- **Create Alert** – This opens the Add Alert page and partially configures the alert with information from the event rule. If completed and saved, the alert reports each time this event rule is triggered.

The Advanced menu includes one or more of the following commands:

- **Re-apply rule** – This allows you to choose a starting point in the past and re-apply this rule to all events that occurred between that point and the current time. This is

useful for testing new or edited rules in *Simulate only* mode before switching to *Enabled* mode. It also can be used to re-apply rules to older events after switching to enabled mode.

- **Clear processed events** – This clears Simulated, Executed, and Skipped events in the Processed Events panel. Pending events are not cleared.

Event Rule Ranking

Event Rules are processed in the order of the rank, with the highest ranked (lowest numbered) rule processed first. Processing order does not depend on the current sorting order of the table, only on the rank number of the rule. All matching rules that are currently enabled are processed. You can sort by rank and then use the up and down arrows next to each rule on the Event Rule page to change the rank of the rules.

File and Process Properties in Event Rule Definitions

Certain choices in the Select File Properties and Select Process Properties panels of the Add/Edit Event Rule page have special behaviors affecting how they are evaluated. A common issue is what happens when an event occurs that is missing data specified in an event rule filter. Evaluation of that event is put into a Pending state until the data becomes available. The following sections describe this and other conditions affecting rule evaluation.

Carbon Black File Reputation Trust and Threat Data

If you choose one of the Carbon Black File Reputation-provided fields, *Trust* or *Threat*, in the File or Process Properties for a rule, only events that have a value for these fields will trigger the rule. Events whose files do not have a Trust or Threat value will go into Pending state (visible in the Processed Events list for the rule) until Carbon Black File Reputation data is available. Once data becomes available, the event will be evaluated against the rule.

Another behavior to be aware of is the treatment of Trust values that are unknown but not missing. If Carbon Black File Reputation and App Control Server have synchronized file information and there is no trust information for a file, no Trust value is shown in the console. However, the *stored* Trust value for a file whose trust is unknown is minus one (-1). Therefore, an event rule that specifies that an action is taken for files with less than a certain trust will be triggered for both low trust files *and files whose trust is unknown*. To limit the rule action to files for which the trust is known to be low (as opposed to unknown), add a second condition that specifies Trust must also be greater than or equal to zero.

File Prevalence

If you choose file *Prevalence* as a filter field, only events for which prevalence is calculated for a related file will trigger the rule. Events whose files have no prevalence value will go into Pending state until a Prevalence value is available. Also, keep in mind that certain settings will make it impossible to accurately report prevalence, including exclusion of Microsoft Support file tracking and exclusion of tracking in selected policies.

File Metadata

If any file metadata field (such as file type, file size, company, publisher, and product) is used as part of a file or process filter, an incoming event will be evaluated only after the specified metadata is reported for that particular file by the agent. The delay between

when an event is reported and when the related file message arrives is normally on the order of seconds. However, if an agent has a large backlog of files to report or goes offline just after sending an event, the delay could be long enough to place event rule evaluation into the Pending state.

File Extension

For both Select File Properties and Select Process Properties, if you choose file Extension as a filter, you must use the file extension *without* the initial dot. For example, to specify that a rule is triggered for batch files with the *bat* extension, you would use *bat* alone, not *.bat* (dot bat). Otherwise the rule would not function properly.

Analysis Results Options

The Select File Properties and Select Process Properties filter menus include file analysis options that are not available in the App Control File Catalog. These options can be used to take action based on the results of analysis by external devices. The current options are *Analysis Result: Check Point*, and *Analysis Result: Palo Alto Networks Wildfire*, – other options may be added after initial release. File analysis results can be one of the following values:

- **Unknown** – The file was not yet analyzed by this service.
- **Clean** – The file was analyzed with this provider and nothing suspicious was found.
- **Potential Risk** – The file was analyzed with this provider and a potential risk was detected.
- **Malicious** – The file was analyzed with this provider and is reported as malicious.
- **Analysis Pending** – The file is still being analyzed with this provider.
- **Analysis Error** – The file was analyzed but analysis returned an error.

As with the Carbon Black File Reputation and Prevalence filters, rules with analysis filters will go into the Pending state for an event that matches the rule but for which analysis results are not available.

Global Bans for Non-Cataloged Files

You can use an Event Rule to create a global ban for a file that has not yet been seen on a App Control Agent reporting to your server. This would happen if you specified a certain event subtype, such as Malicious file detected, in the Event Properties for the rule, and then an analysis service connected to the App Control Server reported a file that triggered an event matching the rule definition. If no other properties are defined for the rule, it immediately creates a “pre-ban” for the file so that if it does appear on any of the agent computers, it will already be banned.

However, if a File Properties filter is added to the rule definition, the rule goes into the Pending state until the reported file actually appears on an agent-managed computer and can be evaluated against the specified properties. If a Process Properties filter is defined

and an event has no process associated with it, the event will be silently skipped, leaving no record in the event view.

Notes

Event Rules that ban or approve MSI files should not rely on hashes reported by a third-party source. In addition, they should not use MD5 or SHA-1 hashes. See [“Approvals and Bans of MSI Files by Hash”](#) on page 303 for details.

How Event Rule Approvals Affect Endpoints

Local approvals initiated by an event rule happen immediately (or as soon as an agent connects to the server). However, unlike most other approvals, event rule global or by-policy approvals are not pushed to endpoints automatically. Like Reputation Rules, Event Rules have three conditions that cause a file approval to be sent to endpoints:

- If the App Control Server has a record of a file being blocked *on any endpoint* and that file is later approved by event rule, the server begins sending the approvals of the file to connected agents immediately.
- If a user attempts to execute an instance of an event-rule-approved file on a computer connected to the App Control Server, the server will allow the agent to run the file immediately, and also will begin sending the approval to other agents.
- If an event-rule-approved file is identified as an installer, the App Control Server begins sending the approval of the file to agents immediately.

Even if a file is approved by event rule and not blocked by another rule, until its approval is sent to agents because of one of the conditions above, instances of the file may be locally unapproved and may block if the agent computer is disconnected from the server before the approval is distributed.

If a file was approved globally or by-policy using an event rule and then an event rule removes that approval, the approval for that file is eliminated for connected computers, and the file state in the File Catalog reverts to unapproved. However, if an instance of this file was executed during the time it was approved by event rule, all instances on computers connected at that time remain *locally approved*.

Event Rule History and Processed Events List

A history of the events processed by each rule is included in the History panel on the Event Rule Details (Edit Event Rule) page. This history is automatically trimmed as events are trimmed from your App Control database.

History

Date Created: Jun 30 2020 03:53:53 PM
Created By: admin
Date Modified: Jun 30 2020 03:53:53 PM
Last Modified By: admin
Last Evaluation Time: Jul 10 2020 11:33:19 AM
Last Processed Event: Jul 10 2020 11:03:22 AM 0 events remaining to process

▼ Processed Events (3 items)

[Show Filters](#) | [Show Columns](#) | [Export to CSV](#) | [Refresh Page](#)

Date Executed	Status	Timestamp	Subtype	Source	Description	Process Name
May 05 2017 10:21:04 AM	Simulated	May 05 2017 07:27:57 AM	Approval request created	SRV2	Approval Request Id 1 was created by user 'SRV2\Admin'.	explorer.exe
May 05 2017 10:21:04 AM	Simulated	May 05 2017 07:29:00 AM	Approval request created	SRV2	Approval Request 2 was created by user 'SRV2\Admin'.	explorer.exe
May 05 2017 10:21:04 AM	Simulated	May 05 2017 07:38:50 AM	Approval request created	SRV2	Approval Request Id 3 was created by user 'SRV2\Admin'.	explorer.exe

The History includes the following information:

- **Date Created** – The time stamp for when this rule was created.
- **Created By** – The console login account of the user who created the rule.
- **Date Modified** – The time stamp for when the rule was last modified.
- **Last Modified By** – The console login account of the user who last modified the rule.
- **Last Evaluation Time** – The time stamp of the last time the rule was triggered by a matching event. In addition, this field shows statistics for any activations of the rule in the past hour, including the number of times it was triggered, the number of events processed, and the time elapsed for processing.
- **Last Processed Event** – The time stamp of the last event that was processed with this rule. This value can be useful in determining whether there is a significant backlog in processing events and also to determine events in the event log that might be processed next. Note that “processing” means the *rule* was processed, not that the resulting *action* has been completed.

Below the History panel, you can click on the Processed Events heading to show the table of events that have been processed by the current rule. This can help you monitor the impact of a rule. The Processed Events table shows the Status of each processed event, which is one of the following:

- **Pending** – The event matched the rule but the rule action has not been completed. If information is available about why the action is in this state, it is displayed as a tooltip when you hover over the Status.
- **Simulated** – The event was processed by the rule in Simulate only mode; the processing was recorded but the action was not executed. See [“Enabling, Disabling, and Deleting Event Rules”](#) for more information.
- **Executed** – The event was processed by the rule and the specified action was taken.
- **Skipped** – The rule was skipped because it would have taken an action that is prohibited or not relevant to the current conditions. For example, a rule cannot globally approve a banned file.

Sample Event Rules

The Event Rules page, which you access by choosing **Rules > Event Rules** on the console menu, includes several sample rules. You can click on the View Details button to open the Event Rule Details page for any of these rules to see how they were specified. You also can use them (or any other existing rule) as a template for a new rule. For example, you could modify a rule to ban or analyze files and processes referenced when an event with *Carbon Black EDR watchlist* subtype is reported.

In addition to the sample rules shown here, a sample rule for deleting files is shown in [“Automating File Deletion Requests”](#) on page 321.

Note

If an event rule depends on a particular configuration of an analysis tool, such as an analysis environment with a specific operating system, and if that environment becomes unavailable, the rule will be disabled automatically after waiting several minutes for the environment to become available again.

Sample Rule: Analyze files from approval requests

This rule sends any file for which an approval request is made to one or more analysis services. By default, it sends files to WildFire, but you can change the rule to send files to any of the analysis services you have configured through the App Control Connector tab on the System Administration page, and can require a result from more than one service. Files that have already been reported by the services you choose are not sent for analysis. See [“Approval Requests and Justifications”](#) on page 566 for more information about approval requests and [Appendix C, “App Control Connector,”](#) for more information about using App Control Connector to integrate analysis services with App Control.

The default properties of this rule are:

- **Event Properties:** Subtype is *Approval request created*
- **File Properties:** Analysis Result: Palo Alto Networks WildFire is *Unknown*
- **Process Properties:** None
- **Action:** Analyze file
 - **Priority:** Medium
 - (Analysis Service Choice): Unchecked

Sample Rule: Resolve approval requests for clean files

This rule performs two actions on files submitted in approval requests if they have been analyzed with WildFire and found to be Clean: it locally approves them, and it resolves the related approval request. If used, it should be enabled along with the *Analyze files from approval requests* rule and ranked after it, so that files are analyzed before their approval requests are resolved.

The default properties of this rule are:

- **Event Properties:** Subtype is *Approval request created*
- **File Properties:** Analysis Result: Palo Alto Networks WildFire is *Clean*
- **Process Properties:** None

- **Action:** Change local file state
 - **Change local state:** Approve
 - **Resolve Related Approval Request:** Unchecked

The rule can be modified to take action based on analysis results from multiple connected devices or services; it will be Pending until all of analysis requests have completed.

Sample Rule: Analyze downloaded files

This rule submits certain files downloaded to a App Control-managed computer from a web browser to Palo Alto Networks WildFire for analysis. It excludes files with properties that suggest they should be trusted or that have already been reported by or do not meet the requirements for WildFire analysis. It also excludes partially downloaded files.

The default properties of this rule are:

- **Event Properties:**
 - **Subtype** is *New file on network*
 - **Process** ends with *iexplore.exe* or *firefox.exe* or *chrome.exe*.
 - **File Name** doesn't contain *.crdownload* or *.part*
- **File Properties:**
 - **File Size** smaller than *10240000*
 - **File State** is not *Approved*
 - **File Type** is *Application*
 - **Analysis Result: Palo Alto Networks WildFire** is *Unknown*
- **Process Properties:** None
- **Action:** Analyze file
 - **Priority:** Medium
 - (Analysis Service Choice): Unchecked

Sample Rule: Report malicious files

This rule applies a global Report Only ban to all malicious files reported or detected by Carbon Black File Reputation or any of the appliances or services integrated with App Control through the Connector.

The default properties of this rule are:

- **Event Properties:** Subtype is *Malicious file detected*.
- **File Properties:** None
- **Process Properties:** None
- **Action:** Change global file state
 - **Change Global State:** Ban (Report only)
 - **Resolve Related Approval Request:** Unchecked
 - **Create for:** All policies

Because this is a Report only rule, it is not necessary to test this in Simulate only mode first.

Chapter 20

Endpoint Notifiers and Approval Requests

This chapter describes the App Control features that involve interactions with endpoint users, which are:

- notifiers that appear on agent-managed computers when a App Control rule blocks file access or related actions
- features that allow users on endpoints to request approval of a file or action, and console tools that manage these requests

Sections

Topic	Page
Notifiers: What Users See	539
The Notifiers Page	545
Assigning Notifiers to Settings and Rules	545
Customizing and Creating Notifiers	548
Notifiers in Windows Session Virtualization	564
Approval Requests and Justifications	566
Enabling Requests and Justifications	566
Submitting Requests and Justifications	568
Managing Requests and Justifications	569
Reviewing and Resolving Requests and Justifications	572
Customizing the Request/Justification Interface in Notifiers	587

Notifiers: What Users See

The App Control Agent runs silently in the background until it detects and blocks an action that matches a blocking rule. When the agent blocks an action, it can display a *notifier* on the computer where the action was attempted, informing the user why the action was blocked. Depending upon the attempted action and the configuration choices made on the App Control Server, notifiers can also give the user options for responding to the block.

Administrators can choose to block actions silently, but most of the descriptions in this section assume that notifiers are enabled for all rules and settings.

Notes

- Notifiers are supported for Windows 8 and Windows 8 Pro in version 7.2.2 and later, but only when these systems run in traditional desktop mode. Notifiers are not supported in the Metro interface.
- If you use pre-8.5.0 Mac and Linux agents, their notifiers will show previous product names. See the Carbon Black [User Exchange](#) for updates on the status of these agents.

Prompt Notifiers

Prompt notifiers tell the user what the attempted action was and why it was interrupted, but also give the user the option of allowing or blocking the action.

Security Notification - File and Path Custom Rule

Target: chrome_elf.dll
Path: c:\program files (x86)\google\chrome\application\84.0.4147.89\
Process: chrome.exe

Carbon Black App Control identified and paused access by chrome.exe to chrome_elf.dll. Choose Allow to grant access, or choose Block to deny access at this time. Scroll down for diagnostic data.

Allow Block

[Submit Justification>>](#)

	Process	Target	Path
1	chrome.exe	chrome_elf.dll	c:\program files (x86)\google\chrom
2	chrome.exe	version.dll	c:\windows\system32\

Justification

Enter your reason for access (512 characters max).

Your Email:

Priority: Medium

Submit

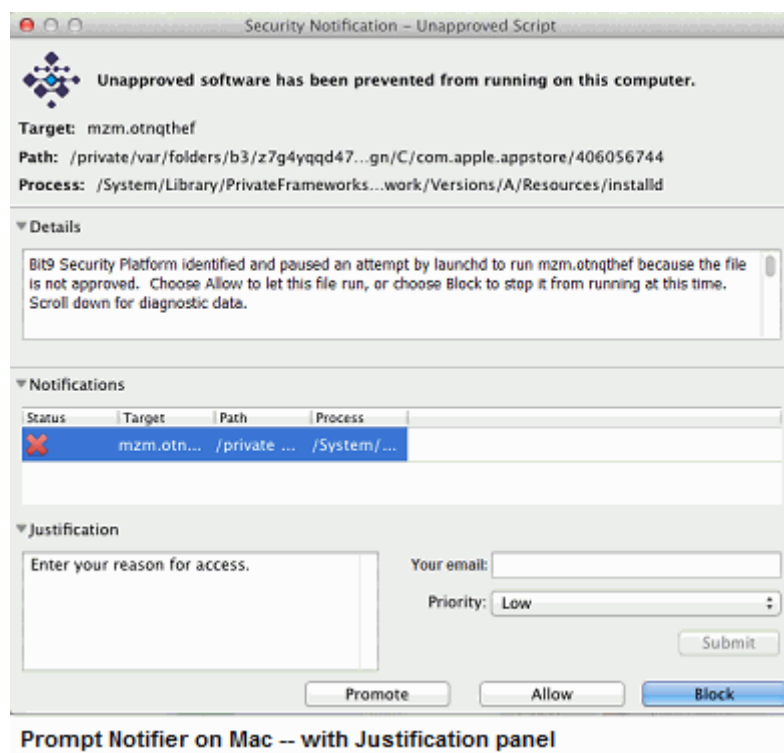
Users see Prompt notifiers under these conditions:

- When they attempt to execute an Unapproved file on a computer that is in Medium (Prompt Unapproved) Enforcement Level.

- When they attempt an action that is governed by a Custom (File and Path) Rule, Registry Rule, or Memory Rule, or a rule within a Rapid Config, and that rule is configured to prompt for a decision.

Because they require a response from the user, prompt notifiers cannot be disabled in rule definitions that have Prompt actions, and they *should not* be disabled for any policy setting that defines a rule that could prompt the user.

Prompt notifiers can include a Justification option, which allows users to send a *justification* of the choice to allow or block the action *before* making that choice. See [“Approval Requests and Justifications”](#) on page 566 for more information about this feature.



The choices on a prompt notifier depend upon the conditions that caused the block:

- **Block** leaves the block in effect, makes no changes in the state of files or devices, and dismisses the notifier.
- **Allow** lets the action take place. If it was a blocked execution of an Unapproved file because of Medium Enforcement on the computer, the file is locally approved and allowed to run. It is allowed to run indefinitely if it is a local file. If it is run remotely from a network share or removable device, it is temporarily approved to run for 14 days. If an allowed Unapproved file is recognized as an installer, files written by it are locally approved. If it is not recognized as an installer, files it writes are not locally approved.
- When an action is blocked by a file execution rule, holding down the Shift key activates the **Promote** button in Mac and Linux and changes **Allow** to **Promote** in Windows. Promote ensures that the file runs as a promoted process, meaning that files written by the process will be locally approved. This is useful if the notifier is displayed for an execution attempt on a file that installs other files but is not recognized by App Control as an installer.

- If the user takes no action on a *prompt* notifier after 10 minutes, the file is blocked, a block event is recorded on the App Control Server, and the notifier is dismissed. However, any interaction with the dialog (e.g., clicking on it or moving it) will prevent the timeout. For block-only notifiers, see the Notifier Timeout setting described in [Table 74](#).

Block-only Notifiers

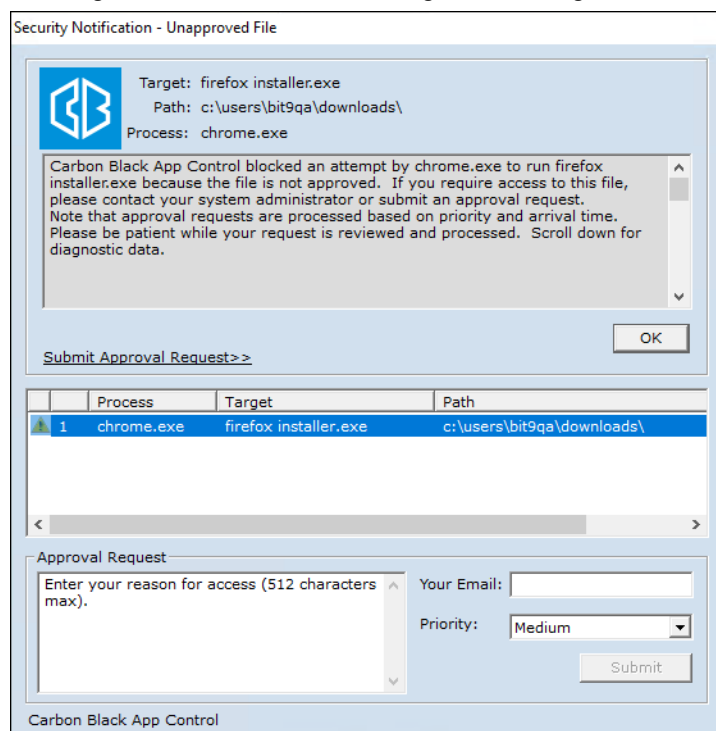
Block-only notifiers inform the user that their action was blocked and why, but do not give the user the option of allowing the action. Users see block-only notifiers, if enabled, under these conditions:

- When they attempt to execute a banned file on a computer that is in Control mode (High, Medium, or Low Enforcement Level).
- When they attempt to execute an unapproved file on a computer that is in High (Block Unapproved) Enforcement Level.
- When they attempt an action that is governed by a Custom Rule, Registry Rule, or Memory Rule, and that rule is configured to block the action.
- When they attempt a file action on a device that is governed by a Device Rule that blocks the action.
- When they attempt a file action that is governed by a Rapid Config that includes a rule to block the action.

The appearance and options of a block-only notifier depend on the platform of the endpoint.

Block Notifiers on Windows Computers

On Windows computers, block notifiers appear as full-sized dialogs. There is no option for allowing the blocked action. Clicking **OK** or using the **Esc** key dismisses the dialog.




If the Approval Request feature is enabled, users can send formal requests for access to files or devices that they can't currently access, and clicking the **Submit** button next to a request dismisses the dialog. Approval Requests have been enabled by default in new installations beginning with v7.0.0. See [“Approval Requests and Justifications”](#) on page 566 for more about this feature, including details about enabling approval requests if you are upgrading from a previous release.

Block-only notifiers can be disabled without disabling their underlying rules.

Block Notifiers on Mac and Linux Computers

On Mac (OS X and Mac OS) and Linux computers, a block notifier appears as a small, translucent notification panel with information about the operation and action that was blocked. Because the notification does not require action, this panel fades and disappears in five seconds unless the user clicks on it. If a new block happens while this notifier is displayed, the new block resets the timer to five seconds.



Clicking on the block notifier before it fades opens the Notifier history window, which provides a history of notifier events that have occurred on the computer. See [“Notifier Tray Icon and History Window”](#) on page 543 for details about the information and actions available there.

Notifier Components


Full-sized notifiers (all Windows notifiers and Prompt notifiers on Mac and Linux) can include the following components, some of which are always shown, some of which are optional, and some of which can be customized:

- The title appears at the top of the window. For example, “Security Notification – Unapproved File”.
- The notifier provides information about the Target of the action. For example, if an execution was attempted, target information includes the file the user attempted to execute, its path, and the process that attempted to execute it.
- A logo appears in the upper left of the notifier. By default, this is the App Control logo, but it can be changed or eliminated.
- On Mac and Linux computers, an additional subtitle appears, for example “Unapproved software has been prevented from running on this computer.”
- Notifier text, which appears in the top text box in the notifier, provides a description of what was blocked and why. For example, “App Control blocked an attempt by explorer.exe to run calc.exe because the file is not approved. If you require access to this file, please contact your system administrator.” On Mac and Linux computers, similar detail is available for each notifier event in the Notifier history window – see [“Notifier Tray Icon and History Window”](#) on page 543.
- On Windows computers, an optional notifier URL link can point to a site that explains security policy and/or provides an opportunity to request access to a blocked object. The link can also be configured to initiate a mail message to request access.
- On Windows computers, a history panel in the notifier shows files that have been blocked on that computer. A green check mark indicates that a file was allowed to run or write. A red ‘x’ indicates that the file or action was blocked, either by a App Control rule or by user choice. A yellow triangle indicates that the notifier timed out before the user took action (and so the action was blocked). A question mark indicates the current block event (i.e., the one that caused the current notifier to display). On Linux and Mac, a similar history is available in the Notifier history window – see [“Notifier Tray Icon and History Window”](#) on page 543.
- An Approval Request or Justification panel allows users to file formal approval requests for files or devices that they can’t currently access, or justifications for why they chose to allow an action if they were given a choice in the notifier. See [“Approval Requests and Justifications”](#) on page 566 for more about this feature.

Notifier Tray Icon and History Window


On Linux and Mac computers, installation of the App Control Agent adds a tray or panel icon that can be used to access a menu with the following options:

- **Show Notifications** – This opens the Notifier history window, which shows past blocks events and the notifier information associated with them. It also provides access to the interface for submitting approval requests for previously blocked files.
- **About** – This shows the agent version and copyright information.



App Control Notifier History Window

On Linux and Mac computers, the Notifier history window shows past blocks events. If the user selects a block event, they can get details about it and submit a request for the blocked file or action to be approved.



The list of block events includes the following information:

- **Status** – This is indicated by an icon: a red X for blocked files or actions; a green check for files or actions that were allowed because of user choice; a yellow triangle if the action was blocked because the notifier timed out before the user took action.
- **Path** – The full path to the file that was blocked.
- **Process** – The full path to the process that attempted the action.
- **Date** – The date and time the file or action was blocked.

Below the history list, the Requests panel allows the user to request approval of the blocked file selected in the list. This panel can be shown or hidden by clicking on the arrow next to its name.

Below the Requests panel, the Details panel provides a more detailed description of the file or action that was blocked. This panel can be shown or hidden by clicking on the arrow next to its name.

Note

On Windows computers, each notifier includes a history panel that functions much the same way as the history list in the Mac or Linux notifier. The key difference is that in Windows, the history is available only when a notifier is displayed.

The Notifiers Page

Available notifiers are shown in a table on the Notifiers page in the console. This page includes the default notifiers provided with the current App Control release and any notifiers you have added. In addition, if you upgraded from a previous version of Bit9 Platform or Parity and modified any of the notifiers, both the current default and the modified version are listed in the Notifiers table. The first modified version of a v6.0.2 notifier has “(custom 1)” appended to the name, the second “(custom 2)”, etc.

You can edit any notifier shown on the page, but you cannot delete the default notifiers.

Notifiers					
Group By: (none) Ascending Show Filters Show Columns Export to CSV Refresh Page Add Notifier Delete Selected Notifiers					
Notifier Name	Title	Link	Timeout	Use Count	
<input type="checkbox"/> Block banned file hashes	Security Notification - Banned File		0	3	
<input type="checkbox"/> Block banned file names	Security Notification - Banned File		0	3	
<input type="checkbox"/> Block executables run from a network drive	Security Notification - Unapproved Network Location		0	3	
<input type="checkbox"/> Block executions from banned removable devices	Security Notification - Banned Device		0	3	
<input type="checkbox"/> Block executions from unapproved removable devices	Security Notification - Unapproved Device		0	3	
<input type="checkbox"/> Block files with banned publishers or certificates	Security Notification - Banned Publisher or Certificate		0	3	

Assigning Notifiers to Settings and Rules

Notifiers can be assigned in two places in the console:

- On the Edit Policy page, for each policy setting
- On the Add/Edit Rule page for custom, registry, and memory rules; a rule can be configured to use the notifier assigned by a computer’s policy or to use a custom notifier specified in the rule details

Assigning Notifiers to Policy Settings

A default, setting-specific notifier is assigned to each policy setting, so notifier configuration is not required. However, you can choose a different notifier for each rule and setting in a policy. This section describes how you assign existing notifiers to settings. See [“Customizing and Creating Notifiers”](#) on page 548 for information about modifying notifiers or creating new ones.

To assign a notifier to a policy setting:

1. On the console menu, choose **Rules > Policies**. The Policies page appears.
2. On the Policies page, click the View Details button next to the name of the policy whose notifier assignments you want to change. The Edit Policy page appears.

- To change the notifier for an Advanced Setting, click the **Advanced** tab.]

The image shows two screenshots of the VMware Carbon Black App Control interface. The top screenshot shows the 'Advanced' tab selected, with a table of settings. The bottom screenshot shows the 'Device Control Settings' tab selected, also with a table of settings. Red boxes highlight the 'Advanced' and 'Device Control Settings' tabs in their respective screenshots.

Name	Status	Notifiers	Add	Edit
Block unanalyzed scripts and executables	Active	<default>: Block unanalyzed scripts and execut		
Block unapproved scripts	Active	<default>: Block unapproved scripts		
Block unapproved executables	Active	<default>: Block unapproved executables		
Block banned file names	Active	<default>: Block banned file names		
Block banned file hashes	Active	Block banned file hashes		
Block executables run from a network drive	Off	<default>: Block executables run from a netwoi		
Block files with banned publishers or certificates	Active	<default>: Block files with banned publishers or		
Enforce memory rules	Active	<default>: Enforce memory rules		
Enforce registry rules	Active	<default>: Enforce registry rules		
Enforce custom (file and path) rules	Active	<default>: Enforce custom (file and path) rules		
Enforce tamper protection	Active	<default>: Enforce tamper protection		
Terminate processes with banned images	Report Only	<default>: Terminate processes with banned im		

Locally approve unapproved files on transition from Visibility or Low Enforcement Level to Medium or High

Name	Status	Notifiers	Add	Edit
Block writes to unapproved removable devices	Off	<default>: Block writes to unapproved removab		
Block writes to banned removable devices	Active	<default>: Block writes to banned removable dr		
Report reads from unapproved removable devices	Off	<none>		
Report reads from banned removable devices	Off	<none>		
Block executions from unapproved removable devices	Off	<default>: Block executions from unapproved s		
Block executions from banned removable devices	Active	<default>: Block executions from banned remo		

- For each setting whose notifier you would like to change, make a new choice from the Notifiers menu.
You can choose <none> to display no notifier when a setting blocks an action. Consider all conditions for a setting, however, before changing its notifier to <none>. For example, if you choose <none> for *Block unapproved executables*, users in Medium Enforcement policies, who should be able to choose whether to block or allow execution of unapproved files, will not have the opportunity to make that decision. The file will be blocked without any notice from the agent.
- Click the **Save** button to preserve your Advanced settings notifier changes.
- If you want to change Device settings notifiers for this policy, click the Device Control Settings tab and repeat steps 4 and 5.
- When you are finished editing the notifiers for this policy, click the **Save & Exit** button to return to the Policies page. From there, you can select other policies and edit their notifiers, if you choose.

Policy Settings with Notifiers

Each of the following policy settings, which appear in the Device Control Settings and Advanced Settings lists on the Edit Policy page, has its own separately assigned notifier, except where noted:

Device Control Settings with Notifiers:

- Block writes to unapproved removable devices
- Block writes to banned removable devices
- Report reads from unapproved removable devices (will never display notifier)
- Report reads from banned removable devices (will never display notifier)

- Block executions from unapproved removable devices
- Block executions from banned removable devices

Advanced Settings with Notifiers:

- Block unanalyzed scripts and executables
- Block unapproved scripts
- Block unapproved executables
- Block banned file names
- Block banned file hashes
- Block executables run from a network drive
- Block files with banned publishers or certificates
- Enforce memory rules
- Enforce registry rules
- Enforce custom (file and path) rules
- Enforce tamper protection
- Terminate processes with banned images

Assigning Notifiers to Custom, Registry and Memory Rules

A notifier can be displayed when a custom, registry, or memory rule blocks an action or prompts the user for a decision to allow or block an action. For each rule, you can choose from two sources for the notifier:

- **Use Policy Specific Notifier** – Each Policy includes an Advanced Setting for each rule type. Each of these policy settings has a Notifier field in which you can specify the notifier that appears on agent computers when that type of rule blocks an action. You also can choose <none> to allow a rule to block an action without displaying any notifier. By default, rules that block or prompt use the policy-specific notifier.
- **Custom Notifier** – If you do not want to use the policy-specific notifier, you can assign any available notifier to any rule. The notifier choices appear on a menu on the Add/Edit page for the rule. You also can Add a new notifier or Edit an existing notifier. See [“Customizing and Creating Notifiers”](#) on page 548 for details.

Definition

Platform: Windows

Rule Type: File Integrity Control

Write Action: Block Use Policy Specific Notifier

Custom Write Notifier: Block unapproved executables Add Edit

Path or File: Specific Path...

Process Exclusion:

Rule Applies To

Policies: All Current and Future policies
 Selected policies

When you choose Prompt as the rule action, Custom Notifier menu does not include <none> as an option because a prompt rule requires a notifier to appear.

When you choose Block as the rule action, you can choose <none> on the Notifier menu for a rule since it is possible you want the rule to block actions without notification.

If you choose Use Policy Specific Notifier for a rule, it is possible that the policy specifies <none> as the Notifier for one of its rule types. In this case, a notifier will not be shown, even for a Prompt rule. Unless you are certain that you never want to prompt the user for a response to a rule, choosing <none> for the rule notifier in a policy is not recommended.

Assigning Notifiers to Rapid Configs

Rapid Configs are sets of rules, possibly including Custom, Memory, or Registry Rules. Some of the rules in a Rapid Config may block actions a user takes. If a Rapid Config can block an action, a Notifier field appears next to the setting that specifies the conditions under which that block takes place. When a Rapid Config contains more than one action that can be blocked, you can choose different notifiers for each action you block or use the same one for all. You also can choose Block for some actions in one Rapid Config and Report or Do Nothing for others. See [Chapter 18, "Rapid Configs,"](#) for more details.

Customizing and Creating Notifiers

You can edit existing notifiers, and you also can create new notifiers. If you edit one of the default notifiers, you can later reset that notifier to its original settings.

Note

The combination of notifier text, notifier link, notifier name, and custom logo path cannot exceed 1900 characters in length. You will see a warning if you exceed the limit.

To customize an existing notifier:

1. There are three ways to open the Edit Notifier page:
 - On the console menu, choose **Rules > Notifiers**, and in the Notifiers table, click the View Details button next to the name of the notifier you want to edit.
 - On Device Settings or Advanced Settings panel of the Edit Policy page, click **Edit** in the far right column next to the name of the notifier you want to edit.
 - On the Edit page for a Custom, Registry or Memory rule, if the Custom Notifier menu is showing, click **Edit** next to the name of the notifier.

Edit Notifier Block banned file hashes
?

Name:

Notifier Title:

Notifier Text:

<BlockText:Carbon Black App Control blocked an attempt by <ProcessName> to run <TargetName> because the file is banned. If you require access to this file, please contact your system administrator or submit an approval request. Note that approval requests are processed based on priority and arrival time. Please be patient while your request is reviewed and processed.> Scroll down for diagnostic data.

Notification Logo:

Notifier Link:

Notifier Timeout: seconds (0 = never timeout, -1 = never display)

Approval Request:

Notifier Applies To

Type	Name
Policy	Template Policy: Block banned file hashes
Policy	Local Approval Policy: Block banned file hashes
Policy	Default Policy: Block banned file hashes

Advanced

Remove Associations

Related Views

All Computers that have received this Notifier

All Computers that have not yet received this Notifier

2. Review and change the notifier settings you want to change (see [Table 74](#)).
3. Click the **Save** button to preserve your changes.

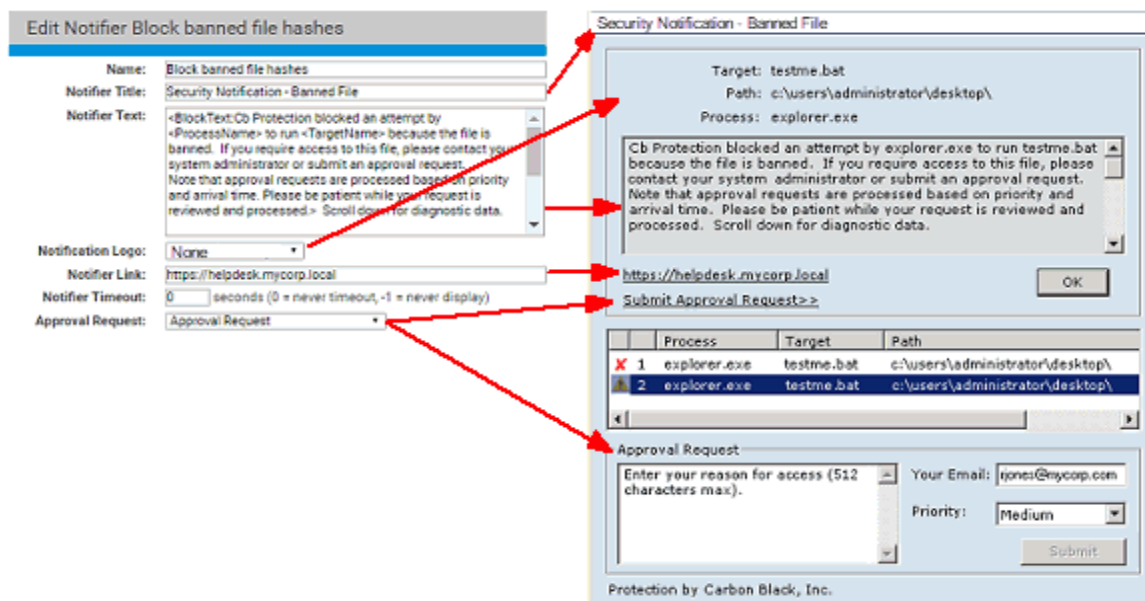
Table 74: Add/Edit Notifier Settings

Field	Description
Copy Settings From	(For Add Notifier page only) Existing notifier from which to copy the initial settings for the new notifier. You can use this to populate all of the new notifiers fields and then modify only those you want to change. Choose <i>(none)</i> if you want to fill in all notifier fields from scratch.
Name	The notifier name as it will appear in the Notifiers table and menus on the policy and rule pages. This name does not appear on notifier displayed to the computer user.
Notifier Title	Window title for the notifier message that the computer user sees when the agent blocks file execution as a result of this setting.

Field	Description
Notifier Text	<p>Explanatory message displayed in the notifier on Windows computers when the agent blocks file execution as a result of this setting. You can modify this message, tag different messages for block-only vs. block-and-prompt conditions, add tags that provide event-specific information, and add other conditional text. Tags here also can modify the Approval Request feature.</p> <p>See “Editing Notifier Text” on page 552 for a description of tags. See “Approval Requests and Justifications” on page 566 for a description of how to activate and configure Approval Requests.</p> <p>Platform Note: Notifier text appears on the Prompt notifier for all platforms, on the Block-only notifier for Windows, and on the App Control Notifier history dialog for a selected item in the history. Notifier messages also appear in the Windows event log.</p>
Notifier Logo	<p>By default, the Carbon Black logo appears in the notifier dialog box when a App Control setting blocks a file. The Notifier logo menu gives you these options:</p> <ul style="list-style-type: none"> • Leave Carbon Black Logo as the selection. • Choose None to display no logo or image in the notifier. • Choose Custom and provide a URL or file path to a different image. See “Specifying a Custom Notifier Logo” on page 560 for details about image format and file path requirements. <p>Platform Note: Custom logos are displayed on Windows agents only.</p>
Notifier Link	<p>Either:</p> <ul style="list-style-type: none"> • a link to an informational web page where the computer user can learn more about your security settings and procedures for responding to blocked files, or • a <i>mailto:</i> link to allow the user to send questions by mail <p>The URL or mailto link provided here can appear literally in the notifier or be represented by a “Friendly Text” description.</p> <p>Leave this field blank if you choose not to display a URL or mailto link at this time.</p> <p>Platform Note: For this release, Notifier Links appear only on Windows notifiers.</p>
Notifier Timeout	<p>The number of seconds that a <i>block-only</i> notifier stays on the screen on a Windows computer. After the specified period of time, the notifier is automatically closed.</p> <p>The default timeout value is zero (0), which leaves the notifier on screen so that the user must respond to it. A value of negative one (-1) instructs agents not to display the notifier at all. See “Disabling App Control Notifiers” on page 562 for additional information about enabling and disabling blocked action notifiers.</p> <p>This value does not affect <i>prompt</i> notifiers, which remain visible for 10 minutes and then automatically block the action if the user has not made a choice.</p> <p>Platform Note: This value affects Windows computers only. On Mac and Linux, a block-only notifier times out in 5 seconds by default.</p>

Field	Description
Approval Request	<p>Determines whether and how the Approval Request feature is enabled for this notifier. The choices are:</p> <ul style="list-style-type: none"> • None - No approval request panel is displayed. • Approval Request - The Approval Request panel appears when a rule completely blocks access to a file. • Justification - The Justification panel appears when a rule prompts a user to allow or block an action. • Approval Request and Justification - The Approval Request/Justification panel appears for both block and prompt conditions. <p>See “Approval Requests and Justifications” on page 566 for more details.</p>
Notifier Applies to	<p>(Appears only if the notifier is assigned to at least one setting or rule) This panel lists all of the rules and settings to which the notifier is assigned. You can remove all of these assignments by clicking Remove Associations in the Advanced menu. If you do this, the affected policy settings revert to their default notifier and the affected rules revert to the policy-specific notifier for their rule type.</p>

The illustration below shows where some of the changes in the Add/Edit Notifier dialog affect the notifier content.



Related Views on the Edit Notifier Page

Using the Related Views menu on the Edit Notifier page, you can get a list of all computers that have received the current version of the notifier, or a list of all computers that have *not* received the notifier.

Creating a New Notifier

Creating a new notifier is similar to editing an existing notifier, with the exception of the initial steps.

To add (create) a new notifier:

1. There are three ways to open the Add Notifier page:
 - a. On the console menu, choose **Rules > Notifiers**, and in the Notifiers table, click **Add Notifier** button.
 - b. On Device Settings or Advanced Settings panel of the Edit Policy page, click **Add** in the far right column next to the name of the notifier you want to edit.
 - c. On the Edit page for a Custom, Registry or Memory rule, if the Custom Notifier menu is showing, click **Add** next to the name of the notifier.
2. If you want to start with the settings of an existing notifier, choose a notifier from the Copy Settings From menu.
3. Enter or edit settings as necessary (see [Table 74](#)).
4. Click the **Save** button to preserve your changes.

Note: Once you click Save on the Add Notifier page, the notifier is saved and added to the Notifiers list. If you navigated to the Add Notifier page from a policy, the new notifier is saved even if you did not click Save on the Edit Policy page. To use the notifier, see [“Assigning Notifiers to Settings and Rules”](#) on page 545.

Editing Notifier Text

You can customize the notifier text a user sees when a App Control rule blocks an action. For example, you might want to add a description of the “Promote” option to the notifiers for your existing policies, unless you prefer not to highlight this option. The App Control Notifier supports conditional, meta and reporting tags that can be used to tailor the information reported to the end user.

Avoid using special characters in notifier text. In particular, the pipe character (|) is known to cause problems.

Platform Note: Notifier text appears on the Prompt notifier for all platforms, on the Block-only notifier for Windows, and on the App Control Notifier history dialog for a selected item in the history. Notifier messages also appear in the Windows event log.

Using Tags in Notifier Text

Notifier text and links can include tags that provide information specific to the event that caused the notification, such as the name of the computer the event occurred on and the policy in force at the time. [Table 75](#) shows the informational tags you can add to a notifier message. You might see other tags that are for Carbon Black Support purposes only.

Note

In addition to providing conditional information to the user, tags in the notifier text box can be used to customize the App Control Approval Request feature. See [“Customizing the Request/Justification Interface in Notifiers”](#) on page 587 for more information about these tags and how to use them.

Table 75: Informational Notifier Tags

Tag	Description	Example Values
<ComputerName>	The local name of the computer on which the block event occurred	"RJONES-LAPTOP"
<DebugInfo>	Technical information about the rule and policy that generated the event. This is a metatag (that is, it contains information represented by other tags)	
<DomainName>	The NetBIOS domain name of the computer on which the block event occurred	"MYCORP"
<EnforcementLevel>	The Enforcement Level of the agent at the time the block occurred	"High (Block Unapproved)", "Medium (Prompt Unapproved)", "Low (Monitor Unapproved)"
<Operation>	The type of operation that was blocked	"Execute", "Write", "Read", etc.
<OsVersion>	The version, build and release of Windows on the agent computer	"Microsoft Windows 7 x64 (build 7600)"
<Bit9AgentVersion>	The version of the agent running on the system on which the operation was blocked.	"8.0.0.256 (Patch 0)"
<Policy>	The policy the agent computer is in	"Research Team", "Sales Group", "Guests", etc.
<ProcessName>	The name (without the path) of the process that was blocked	"explorer.exe"
<ProcessPath>	The path (without the name) of the process that was blocked	"c:\windows\system32\"
<ProcessPathName>	The full path, including name, of the process that was blocked	"c:\windows\system32\explorer.exe"
<ProcessPublisher>	The publisher name for the source process, if signed	"Carbon Black, Inc.", "Google Inc.", "Microsoft Corporation", etc.
<ProcessSha256>	The SHA256 hash (hexadecimal) of the source process	
<RuleType>	The type of rule that was triggered	"File and Path", "Registry", "Memory", "Process", etc.
<TargetName>	The name (without the path) of the target file, registry key or process name to which access was attempted	"foo.bat"

Tag	Description	Example Values
<TargetPath>	The path of the target file, key or process (without the name)	"c:\test\"
<TargetPathName>	The full path and name of the target	"c:\test\foo.bat"
<TargetPublisher>	The publisher name for the target file, if signed	"Carbon Black, Inc.", "Google Inc.", "Microsoft Corporation", etc.
<TargetDevice>	The drive letter of the device on which an action was blocked. Unmapped devices are shown as \\device\<name>.	
<TargetShare>	The network path (without the filename) to the remote drive on which access to a file was blocked.	"\\SERVER3\temp\mydir"
<TargetSha256>	The SHA256 hash (hexadecimal) of the target file	
<TargetSha1>	The SHA1 hash (hexadecimal) of the target file	
<TargetMD5>	The MD5 hash (hexadecimal) of the target file	
<UserName>	The name of the user in whose context the blocked operation was initiated	"\MYCORP\rjones"

Conditional Messages for Block vs. Prompt

By using conditional tags within the same notifier text, you can show the user one message for block-only notifiers, when an action is simply blocked by a App Control rule, and a different message for prompt notifiers, when a user is prompted to block or allow an action. For example, you can create a single string of notifier text that displays a "block" message when a user in High Enforcement attempts to execute an unapproved file, but displays an "ask" message when a user in Medium Enforcement attempts to execute the same file. Similar prompt messages can be used for Custom, Registry or Memory rules in which the user is offered the option of blocking or allowing an action. [Table 76](#) shows the tags for different block conditions ("*message*" represents the variable text you use in the message).

Table 76: Conditional Notifier Tags

	Description
<BlockText:message>	Text to display if the rule blocks an action and the user has no choice to allow it.
<AskText:message>	Text to display if the rule prompts the user for a decision on whether to block or proceed with an action. This is the most generic “prompt” case.
<AskAllowText:message>	Text to display if the rule prompts the user for a decision on whether to <i>block or allow file execution</i> .
<AskRestrictText:message>	Text to display if the rule prompts the user for a decision on whether to <i>allow or restrict memory access</i> .
<AskApproveText:message>	Text to display if the rule prompts the user for a decision on whether to <i>block writing of a file or to approve the file and allow it to be written</i> .

For example, when an unapproved file is blocked, the notifier text might include the following:

```
An unapproved file attempted to run on this
computer<BlockText: and has been blocked. If you require
access to this file, please contact your system
administrator.><AskText:. Choose Allow to continue to let
this file run, or choose Block to prevent it from running at
this time.>
```

When a computer with an agent in a High Enforcement policy with this notifier text attempts to execute an unapproved file, the notifier message uses the *BlockText*:

```
An unapproved file attempted to run on this computer and has
been blocked. If you require access to this file, please
contact your system administrator.
```

However, when a computer with an agent in a Medium Enforcement policy with this same notifier text attempts to open an unapproved file, the notifier message uses the *AskText*:

```
An unapproved file attempted to run on this computer. Choose
Allow to continue to let this file run, or choose Block to
prevent it from running at this time.
```

You can nest other tags inside the conditional block/ask tags shown in [Table 76](#). For example, the following is the default notifier message for *blocked*, *unapproved* files:

```
<BlockText:App Control blocked an attempt by <ProcessName>
to run <TargetName> because the file is not approved. If you
require access to this file, please contact your system
administrator.><AskText:App Control identified and paused an
attempt by <ProcessName> to run <TargetName> because the
file is not approved. Choose Allow to let this file run, or
choose Block to stop it from running at this time.>
```

Notice that there are other tags nested inside both the BlockText and AskText conditional tags. The conditional block/ask tags are the only notifier *text* tags inside which you can nest other tags. In the notifier *link*, you can nest tags inside the “FriendlyText” tag.

Note

When you upgrade App Control Server from a previous release, your existing notifier messages are preserved, including those for Default and Template policies. Older notifiers might not include conditional text to provide different messages for “block” conditions and “ask” conditions and other special tags.

Informational Tags as Conditional Operators

In addition to the special “block-and-ask” conditional operators, notifier messages can include other conditional text based on any of the informational tags shown in [Table 75](#), except for the metatags, such as <DebugInfo>. You construct conditional text tags as follows:

```
<tagnameText:pattern-to-match:message-text>
```

You must append the word “Text” directly to the end of the tag name: the tag will not work without this addition.

For example, to set up notifier text that appears only if the computer on which an action is attempted is running App Control Agent 8.0.0, you would use the <Bit9AgentVersion> tag as shown in the following example:

```
<Bit9AgentVersionText:8.0.0.*:This will display only on
8.0.0 agents>
```

The asterisk wildcard character in “8.0.0.*” is used so that any build number of App Control Agent 8.0.0 matches the condition. The asterisk matches zero or more of any character; the question mark matches any one character (but not zero characters).

You could set up notifier text to appear if the hash for a target file matches a particular SHA-256 hash, using the `<TargetSha256>` tag. You could nest this conditional text within a generic “file blocked” notifier, as shown in the following example:

```
App Control blocked an attempt by <ProcessName> to run
<TargetName> because the file is banned.
<TargetSha256Text:c1c4eacd1fe39c93df477f335644902b3b83cc437b
fe4b641960f874af1e0708:This version of MyFavoriteApp has a
major security flaw.>
If you require a solution to this block, please contact your
system administrator. Scroll down for diagnostic data.

<DebugInfo>
```

Editing the Notifier Link

If you configure a notifier link, your users can click on it when an action is blocked to contact your in-house support desk or go to a web page that explains more about why the action was blocked. Although you can use the same notifier link for all conditions in which App Control blocks a file action, you have the option of providing a different link for each notifier, and as with notifier text, you can embed tags to provide more information about the event in the link.

A notifier link is one method for managing requests for access to a file or device, and may be a good choice if you already have IT policies in place for collecting and responding to these requests. App Control also provides its own Approval Request feature, which populates the notifier with the fields necessary for the user to compose and submit a request and manages these requests directly on the console. See [“Approval Requests and Justifications”](#) on page 566 for more information.

Platform Note: Notifier links display only on Windows notifiers.

Tags in Notifier Links


Using tags in the **Notifier link** field of the Add/Edit Notifier page can be helpful in two different ways:

- You can use tags to customize notifier mail messages or site URLs. This can be helpful for creating automated work-flow requests or making a website link automatically go to information about the file that caused the notifier to appear. [Table 75, “Informational Notifier Tags,”](#) on page 553 shows the complete list of these tags.
- You can create “FriendlyText” to display on the notifier dialog in place of the URL. The FriendlyText tag may appear anywhere in the notifier link text.

The following notifier link demonstrates both of these uses of tags:

```
mailto:it@mycorp.com?subject=Request approval of
<TargetName>&body=<UserName> on
<DomainName>\<ComputerName>has requested access to
<TargetName>.\%0AFile details available at https://
yourserver1/file-details.php?hash=<TargetSha256>
<FriendlyText:Please click here to request access to this
file.>
```

When the notifier text above is used in the “Block unapproved executables” notifier, an attempt to execute an unapproved file on a computer in High Enforcement displays a notifier similar to the following:



Notice that instead of displaying the notifier link URL (“mailto:mycorp.com...”), the link shows the “Submit Approval Request”), which provides an indication of why they would click on the link.

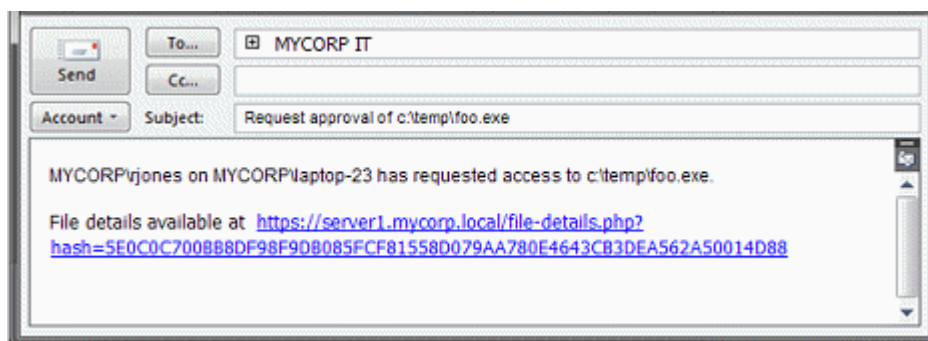
You can nest other tags inside a FriendlyText tag. For example, instead of the generic link text shown above, you could create the following link:

```
<FriendlyText>Please click here to request access to
<TargetName>.>
```

which would insert the name of the file that was blocked in the link text.

Whether you display a URL or friendly text, the resulting link text is displayed as one or two lines. The text will not interfere with the action buttons (“OK”, “Allow”, “Block”), and if link text is too long, it is truncated to fit on the dialog box.

In the example shown, when the user clicks on the link, a mail message similar to the following is initiated in the user’s default mail client:



The notifier link defined above used tags to make several customizations:

- It generated an email message to the organization's IT group requesting access to an unapproved file.
- It specified the name of the file in the message header.
- It identified the user, the computer, and the file in the message body.
- It provided a URL in the mail message that points directly to the File Details page in the console for the specific file in the request.

If this were a “block-and-ask” situation in which the end user could make his or her own judgment about a file, you could create a simpler notifier link that goes directly to the URL for the file details (without generating a mail message), similar to the following:

```
https://yourserver1/file-details.php?hash=<TargetSha256>  
<FriendlyText:Please click here for information about this  
file.>
```

Editing the Notifier Source Line

There is a line at the bottom of notifiers that identifies the source of the notifier. By default, this says *Protection by Carbon Black, Inc.* You can change this line by inserting the following tag into the Notifier Text field, substituting your own source identification for *text*:

```
<NotifierComment:text>
```

If you want to eliminate this line from the notifier, use a single space as your *text*.

Platform Note: The Notifier Source line displays only on Windows notifiers.

Specifying a Custom Notifier Logo

By default, the notifier displayed when files are blocked on an agent computer includes the Carbon Black logo. You also have the option of having no logo on a notifier, or of providing a custom logo. Logos are specified on a per-notifier basis.

Important

- If you create and configure a custom notifier logo, it will appear only on Windows systems. For Linux and Mac endpoints, you can either display the Carbon Black logo or no logo.
- Pre-7.0.0 implementations of a custom logo, including both special solutions provided by Carbon Black Support and the standard customization available in Bit9 Parity 6.0.2, are not maintained when you upgrade to this release. You must use the method below to implement custom logos. If you specified a custom logo in v7.0.0 or later, that will be maintained on upgrade.
- Pre-6.0.2 App Control Agents will not display a newly configured custom logo until they are upgraded.

To specify a custom logo for a notifier (Windows only):

1. On the console menu, choose **Rules > Notifiers**. The Notifiers page appears.
2. On the Notifiers page, either:
 - Click **Add Notifier** if you are creating a new notifier. The Add Notifier page appears.
 - or-
 - Click the View Details button next to the name of an existing notifier you want to edit. The Edit Notifier page appears:

Edit Notifier Block unapproved executables

Name:	Block unapproved executables
Notifier Title:	Security Notification - Unapproved File
Notifier Text:	<BlockText:Carbon Black App Control blocked an attempt by <ProcessName> to run <TargetName> because the file is not approved. If you require access to this file, please contact your system administrator or submit an approval request. Note that approval requests are processed based on priority and arrival time. Please be patient while your request is reviewed and processed.><AskText:Carbon Black App Control identified and paused an attempt by <ProcessName> to run
Notification Logo:	Carbon Black Logo
Notifier Link:	
Notifier Timeout:	0 seconds (0 = never timeout, -1 = never display)
Approval Request:	Approval Request and Justification

- On the Notifier Logo menu, choose **Custom**. A text box appears next to the menu.

Notification Logo:

Notifier Link:

Notifier Timeout: seconds (0 = never timeout, -1 = never display)

Approval Request:

- Put the file containing the logo you want to use in an accessible location, and enter that location in the Notifier logo text box. You have three options for specifying the location of the logo file:
 - UNC:** You can provide a network-based path specification to the logo file in the form **\\server\share\path\imagefile.gif**. The App Control Agent will attempt to make a local copy. If the file cannot be downloaded, the agent will continue to use the prior image (e.g., the default Carbon Black image) until the new image can be obtained. The agent will continue to attempt to download the image once per hour until the image is successfully downloaded or the image is explicitly changed or disabled.
Note: The **LocalSystem** account must have access to the UNC path you provide for the image to be accessible on agent computers. Also, you must not put the logo in a location that would require a password for access.
 - URL:** You can specify a web-based path in the form **http://path/imagefile.gif**. This path should be accessible to the App Control Agent process and allow anonymous, unauthenticated access. The App Control Agent will make a local copy of this file as described above.
 - Local:** You can specify a local file path (on the local computer) in the form **d:\path\imagefile.gif**. The target file must be locally accessible to the App Control Agent process. You must put the logo file on each agent computer that will use it. Any updates to this file take place the next time the notifier is displayed. If the specified path is not accessible, the Carbon Black logo is displayed instead and an event is generated once per App Control Agent session, just as with non-local paths.
- Click **Save**. Your changes are saved and the Notifiers page is displayed.
- Repeat the steps above for each notifier whose logo you want to change.

Image File Requirements

Windows systems on which the App Control Agent is installed include a blank sample notifier image called **GenericLogo.gif**, which is located in the App Control data directory (by default, **ProgramData\Bit9\Parity Agent\images**). Assuming that the agent is installed on the App Control Server, you can go to this folder on the server and use **GenericLogo.gif** as a starting point for creating your own logo image. Otherwise you can copy it from another system that has the agent installed.

The custom image you provide should meet the following requirements:

- The image size should be 60 x 60 pixels.
- The file format should be GIF, JPG or BMP.
- The image should use the same background as **GenericLogo.gif**; you *cannot* use a transparent background.

Logo-Related Events

If all App Control Agents successfully retrieve your custom logo, there will be no logo-related events generated. If an agent fails to retrieve its logo file, however, an event of subtype “Agent Error” will be generated, noting the computer name and the image file name. If (and only if) there was a failure to retrieve the logo, another event is generated if the computer later successfully retrieves the custom logo.

Changing the Logo Image

When you specify a non-local image as the notifier logo (i.e., using a UNC or URL path), that image is copied to each agent system, including the server if it has the agent installed. If you change the non-local image but do not change its name, App Control Agents will not update to the changed image.

To update the logo image for a notifier, change the name of the image file and update the Notifier logo path for that policy. For example, if you deploy a custom logo `\\server\share\mylogo.gif` and you then modify the logo, you could rename the file to `mynewlogo.gif` and edit the path in the notifier details to `\\server\share\mynewlogo.gif`. Agents in that policy would then update to the new image.

Image files downloaded to agents are not updated or deleted. Because of this, if you switched from `logo1.gif` to `newlogo.gif`, and then you switched back to `logo1.gif`, the originally downloaded version of `logo1.gif` would be used, even if you had modified the source image file at the location you entered for download.

Suppressing the Notifier Logo in a Policy

You can prevent display of the notifier logo for all notifiers in a policy. The *Suppress Logo in Notifier* checkbox on the Add/Edit Policy page suppresses the logo, regardless of what the notifier configuration in each notifier specifies.

Resetting a Notifier to Initial Settings

You can reset any default notifier to its initial settings. You reset a notifier by opening the Edit Notifier page for the notifier and clicking **Reset Notifier**. If there is no Reset Notifier button on the page, the notifier was not one of the default notifiers.

Important

If you use Reset Notifier, you will lose all of the customizations you may have made to a default notifier – there is no undo.

Resetting a Policy to Initial Notifiers

The Edit Policy page includes a *Reset Policy* button. When you click this button and choose **OK** to confirm, the Device Control and Advanced settings are reset to the *initial* settings of the Template Policy (i.e., the settings in effect immediately after you installed the App Control Server). This includes reverting to the default notifiers for each setting.

Disabling App Control Notifiers

There might be situations in which you want to disable notifiers for some or all of your agent computers. For example, if you are running single-purpose devices in High Enforcement, you might simply want to block unauthorized actions without feedback.

Block-only notifiers can be disabled without disabling the rules that would otherwise display them. You can disable notifiers on a per-setting basis in each policy. You also can disable notifiers for specific custom, memory, or registry rules.

You can disable notifiers only for *block-only* rules. Rules that *prompt* users for a response should always display a notifier.

Disabling App Control notifiers does not necessarily mean that actions will be blocked silently. Some App Control blocks cause the display of operating system notifiers. Also, except for a Custom, Registry, or Memory rule that has *Block Silently* as its action, events continue to be recorded for blocks even though the notifier is disabled.

To disable notification for a setting in a policy:

1. Open the Edit Policy page for a policy whose notifiers you want to disable.
2. Choose the tab on which you want to disable one or notifiers. For example, if you want to disable an Advanced Setting notifier, click **Advanced**.
3. For each setting on that tab whose notifier you would like to disable, choose <none> on the Notifiers menu.
Consider all conditions for a setting before setting its notifier to <none>. For example, if you choose <none> for *Block unapproved executables*, users in Medium Enforcement policies, who should be able to choose whether to block or allow execution of unapproved files, will not have the opportunity to make that decision. The file will be blocked without any notice from App Control.
4. Click the **Save** button to preserve your changes on this tab and select another tab if you want to disable additional notifiers.
5. Repeat steps 2-4 for each tab on which you want to disable notifiers in this policy.
6. When you have finished disabling notifiers for this policy, click the **Save & Exit** button.
7. Repeat this procedure for each policy whose notifiers you want to change.

To disable notification for a specific custom, registry, or memory rule:

1. On the console menu, choose **Rules > Software Rules**.
2. On the Software Rules page, click the tab for the rule type you want to modify.
3. In the table of rules, click the View Details button next to the rule whose notifier you want to disable.
4. On the Edit Rule page, *un-check* the Use Policy Specific Notifier box next to any actions configured in the rule.
5. In the Custom Notifier menu, choose <none>. Note that <none> is not an option for rules that prompt the user.
6. Click **Save** to preserve your changes.

For Block actions, events are still recorded even if the notifier is disabled. For some rules, you can choose Block Silently from the action menu to disable both notifiers and event recording.

Note

You also can disable a notifier everywhere it appears (rather than choosing not to use a notifier for a setting). You do this by entering minus one (-1) as the value for Notifier Timeout on the Add/Edit Notifier page.

Notifiers in Windows Session Virtualization

App Control supports special treatment of notifiers for hosted session virtualization environments, such as those provided by Citrix XenApp, Windows Server Remote Desktop Services, and Windows Server Terminal Services. In these environments, you can add special notifier tags that instruct your App Control Server to route notifiers in the following ways: If one user is logged into multiple sessions and attempts an action that triggers a notifier, the notifier is displayed to all logged in sessions for that user. For a prompt notifier, responding to any of those notifiers dismisses all of them. For a block notifier, the notifier must be dismissed in each session.

- If multiple users are logged in to one session each, and if one of them attempts an action that triggers a notifier, the notifier is displayed only to the user that triggered the block.
- If an action that triggers a notifier is initiated by the system and not a specific user, you can choose to display the notifier to a specified user or group, all users, or no users. No matter which option you configure, App Control logs a block event on the Events page.
- Even when you enable the special notifier behavior, users of agent-managed computers not using session virtualization see notifiers according to the normal rules.

Special treatment of notifiers applies only to *hosted sessions* on a terminal or application server (session virtualization). That is, they apply to a single system and users and applications on that system. Application virtualization that runs applications locally is not compatible with the feature.

Notifications are always directed to the session of the user taking the action that blocks, not necessarily the originating session. For example, if user A has access to user B's command prompt, and User A executes `runas /user:A cmd.exe` and then executes an unapproved file, the notifier is displayed in user A's remote session, not in the session where user A appeared to have executed the unapproved file.

Platform Note

Broadcast notifiers are available for Windows sessions only.

There are two tags that activate session virtualization notifier behavior:

- **<NotifierBroadcastMessage>** is required to enable special notifier routing. If present, notifiers are displayed on all sessions for the user that initiated an action, or for System actions, as specified by `NotifierBroadcastSystem`.

- **<NotifierBroadcastSystem:user|group|blank>** is used to determine what is done when a system-initiated action is blocked by a App Control rule. The default is **<NotifierBroadcastSystem>** with no other arguments. If you leave this tag out but have **<NotifierBroadcastMessage>** in the notifier, notifiers will be displayed to all logged in session users.

The following procedure assumes you want to modify notifier behavior for all settings in a policy. You can add the tags to individual notifiers through the Notifier page if you prefer.

To enable special notifier routing for session virtualization:

1. On the console menu, choose **Rules > Policies**.
2. Click on the View Details button next to the policy whose notifiers you want to edit.
3. Choose a setting whose notifier you want to change and click on the **Edit** button to the right of the Notifier field.
4. On the Edit Notifier page, enter **<NotifierBroadcastMessage>** in the Notifier Text field.
5. Also in Notifier Text, enter **<NotifierBroadcastSystem:>** with the option you want:
 - To route notifiers for blocks of system-initiated actions to a single user, enter a user name after the colon. For example, **<NotifierBroadcastSystem:MYCORP\jsmith>**
 - To route notifiers for blocks of system-initiated actions to members of a group, enter a specified or built-in group name after the colon. For example, **<NotifierBroadcastSystem:MYCORP\itgroup>**
 - To suppress notifiers for blocks of system-initiated actions, do not enter anything after the colon (the colon is optional in this case). For example, **<NotifierBroadcastSystem>**
Note that if you suppress the notifier in this case, users in Medium Enforcement Level policies will not have the option of allowing unapproved software – it will always be blocked.
 - If you leave the **<NotifierBroadcastSystem>** tag out of the notifier text area but include **<NotifierBroadcastMessage>**, notifiers will be displayed to all logged in session users.
6. **Save** your changes to the notifier.
7. Repeat for each notifier in the policy (and any others you would like to modify).

Approval Requests and Justifications

When a App Control rule blocks an action, it normally displays a notifier on the computer where the action was blocked. The Approval Request feature allows users to send feedback to administrators when they see a notifier:

- **Approval Requests** – When an action is *blocked* with no option to allow, users might want to request that a file or device is unblocked. Notifiers can be configured to allow users to submit a formal *approval request* for a blocked file or device.
- **Justifications** – When an action triggers a *prompt* notifier, which provides the user the option to block or allow access, you might want to allow (or require) the user to explain why they allowed the action. The approval request feature also includes an interface for submitting these *justifications*.

When submitted, both approval requests and justifications appear in the Approval Request table in the console and are recorded in the App Control events database. If you choose, you can enable a built-in alert that is triggered when someone makes an approval request. There also is an alert for justifications.

Throughout this chapter “Approval Requests” is the generic term used for the feature that includes both approval requests and justifications. A distinction is made where needed.

Notes

- Pre-7.0 agents cannot submit approval requests or justifications.
- Approval Requests and justifications can be used for actions blocked by different types of rules. A full set of features is included on the Approval Request and Approval Request Details page for managing rules that are based on file state (file bans and blocks of unapproved files). For Custom, Registry or Memory rules, links are provided to take you to the rule details pages but the request management features are more limited. There is not currently any direct way to manage requests involving Rapid Configs from the Approval Requests pages.
- If you already have your own request workflow in place, you can use *notifier links* to manage requests *outside of the App Control Console*. Links can be used to automatically open a blank email directed to the person or group responsible for approving files, or they can direct the user to a web page that you use to handle IT requests. See [“Editing Notifier Text”](#) on page 552 for steps to set up these links.
Platform Note: Notifier links appear on Windows computers only.

Enabling Requests and Justifications

Approval requests and justifications are enabled on a per-notifier basis. What (if anything) you do to enable them depends on whether you are upgrading from a pre-7.0 release, and also on whether you want to customize the appearance and behavior of the feature:

- **New Installations** – In new App Control installations beginning with version 7.0.0, Approval Requests are enabled for all file and device blocking settings in the Default and Template policies. New policies that you create from these policies will also be configured for approval requests and justifications, and will distinguish between the

two in the notifier interface. You do not need to follow the procedure below unless you want to further customize the notifiers.

- **Upgrades** – Upgrades from pre-7.0 versions of Bit9 Parity will not have Approval Requests enabled. You can enable them using the Approval Request menu in each notifier, and you can further customize their appearance by adding tags to the notifier text. For upgrades, any notifiers you customized prior to v7.0.1 do not distinguish between *approval requests* and *justifications* in the notifier labeling.

Platform Note: Disabling Approval Requests and/or Justifications prevents their interfaces from appearing on all Prompt notifiers and Windows Block-only notifiers. On Mac and Linux, the Justifications panel is grayed out in the Notifier History window when a block event is selected and Approval Requests and/or Justifications are disabled.

To enable Approval Requests and/or Justifications for a notifier:

1. Choose a notifier and open its Edit page.

Edit Notifier Block banned file hashes

Name:	<input type="text" value="Block banned file hashes"/>
Notifier Title:	<input type="text" value="Security Notification - Banned File"/>
Notifier Text:	<div style="border: 1px solid #ccc; padding: 5px; min-height: 100px;"> <p><BlockText:Carbon Black App Control blocked an attempt by <ProcessName> to run <TargetName> because the file is banned. If you require access to this file, please contact your system administrator or submit an approval request. Note that approval requests are processed based on priority and arrival time. Please be patient while your request is reviewed and processed.> Scroll down for diagnostic data.</p> </div>
Notification Logo:	<input type="text" value="Carbon Black Logo"/>
Notifier Link:	<input type="text"/>
Notifier Timeout:	<input type="text" value="0"/> seconds (0 = never timeout, -1 = never display)
Approval Request:	<input type="text" value="Approval Request"/>

Notifier Applies To

Type	Name
Policy	Template Policy: Block banned file hashes
Policy	Local Approval Policy: Block banned file hashes
Policy	Default Policy: Block banned file hashes
Policy	Contractor policy: Block banned file hashes

2. On the Approval Request menu, choose the option you want. The options are:
 - **Approval Request**
 - **Justification**
 - **Approval Request and Justification**
 - **None**

3. Click the **Save** button.

Note


You can enable automatic email notification that informs the requestor when an approval request is closed. See “[Reviewing and Resolving Requests and Justifications](#)” on page 572.

Submitting Requests and Justifications

When a file action is blocked without an option to allow it, if Approval Requests are enabled, the user can request approval of the file. The location for entering this request varies depending upon the platform.

On Windows computers, Approval Request are submitted through the block notifier. The users can read the notifier’s description of the block and why it happened. If the user still wants access to the file or device that was blocked, he or she can type an Approval Request of up to 512 characters into the Approval Request box in the bottom-left section of the notifier. The user can enter an email address if they want that included in the request, and can set a priority (Medium by default). Once the text of the approval is entered, the **Submit** button is activated and clicking it submits the request to the App Control Server.

On the Windows notifier, the *Submit button*, not the *Submit Approval Request link*, sends the request. The link *Submit Approval Request link* opens and closes the Approval Request panel at the bottom of the notifier.



Submitting a request from a block-only notifier closes the notifier – it is not necessary to click **OK** in this case.

On Mac and Linux computers, when an action is completely blocked, users can make approval requests from the App Control Notifier history window by selecting any block event from the history and entering the information as described above for Windows (limited to 512 characters). Unlike in Windows, Mac and Linux users can make a series of requests for different file approvals without closing the App Control Notifier history.

On all platforms, if a notifier prompts the user to Allow or Block a file action, the user can submit a *Justification* for choosing Allow. The interface is the same as for an Approval Request. After submitting the justification, the user must click either the Block button or one of the buttons that let the action happen (Allow or Promote).

Once a user submits a request or justification, there is no formal connection to the request from the agent. However, the user can send another request for the same file or device, and can change the comments or the priority (for example, if lack of access to a file is preventing them from accomplishing a task) in the resubmission. Response to the request, or lack of one, is at the discretion of App Control administrator reviewing it.

Managing Requests and Justifications

App Control Console users with default Administrator and PowerUser privileges can view and manage approval requests. In addition, custom User Roles can be created with permission to just view approval requests or to view and manage them. If a console user's primary activity will be addressing approval requests, that user might want to set the Approval Request page as the Default Starting Page on login. This is done on the User Settings page (*loginname* > **User Settings**).

Requests and justifications submitted by users appear on the Approval Request page.

To view the Approval Requests and Justifications table:

- On the console menu, choose **Tools > Approval Requests**.

Status Summary

Time Period: Show: Approval Requests Justifications Both

Submitted: 1 | Open: 0 | Escalated: 0 | Closed: 0 | Not Resolved: 1

Approval Requests and Justifications

Saved Views: (The Current View Has Unsaved Changes - Discard) Group By: (none) Ascending View As: Individual Requests Related f

Show Filters | Show Columns | Export to CSV | Refresh Table

Action: Automatically apply Showing 1 out of 1 item(s)

<input type="checkbox"/> Select 1	ID	Date Requested	Requestor	Reason	File Name	Assessment	Pu
<input type="checkbox"/>	1	Jul 20 2020 04:32:42 PM	APC-58609\bitt9qa	I need to test software in multiple browsers.	firefox installer.exe	<input checked="" type="checkbox"/>	M

Showing 1 out of 1 item(s) Showing all data

Approval Requests Summary

The Approval Requests and Justifications table initially lists all of the open items submitted by users. At the top of the page, it includes a Status Summary panel that shows the

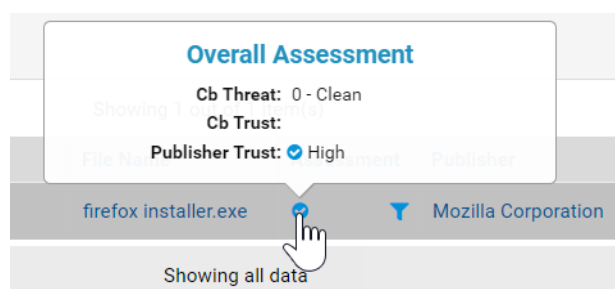
number of new, open, escalated, closed and unresolved requests over a given period. If you hovering the mouse cursor over a box in the summary, additional information about that status is displayed. These status types will be described later in this section.

The reporting period for the summary is “All Time” (no limit) by default, but you can choose a different period on the Time Period menu, and your choice will persist until you change it again. The Status Summary panel can be useful for tracking how well you are meeting the stated service level agreements within your company for user requests. In addition, clicking on any of the boxes in the summary filters the Approval Requests and Justifications table to show the requests in that category.

As with other table pages, you can add columns to those shown by default on the Approval Requests and Justifications table. In addition to the information shown in the table itself, details for certain fields are shown when you move the mouse cursor over those fields. This includes the Assessment, Computer, File, Policy, Publisher and Related Requests fields.

Assessment

The popup information for Assessment shows all of the threat and trust information available for the requested file and its publisher.



Assessment combines CB Reputation file and publisher trust and file threat values with the verdicts about this file from all connected network security devices. It uses the worst report from the available sources to rate the file. The possible values are:

- **Malicious (Red)** – This is the assessment if the CB Reputation file threat or any connector result reports that the file is malicious.
- **Potential Risk (Yellow)** – This is the assessment if the App Control file threat or any connector result reports that the file is a potential risk, or if the App Control file trust is zero.
- **Potentially Clean** – This is the assessment if the CB Reputation file trust is between 1 and 6 or the publisher trust is low or medium.
- **Clean (Green)** – This is the assessment if the CB Reputation file threat or connector results are clean, or if the publisher trust is high, or if the CB Reputation file trust is 7 or higher.

No assessment is displayed if none of the contributing information is available for the file. Clicking on the Assessment brings up a details screen.

Approval Request Details

The Approval Request Details or Justification Details page shows the complete details for a request or justifications. To find a specific request, you can click Show Filters and use

any of the filters in the Filters panel on the Approval Requests and Justifications page. You also can use the Search box on that page to enter text that matches any of the following:

- ID
- Computer Name
- Requestor
- File Name
- Publisher Name

Note that the Search box only matches objects that begin with the characters or string you provide.

To open the details page for one Approval Request or Justification:

- On the Approval Requests and Justifications table page, click on the View Details button next to a request.

Approval Request Details (firefox installer.exe)

Id: 1

Date Requested: Jul 20 2020 04:32:42 PM

Computer: WORKGROUP\APC-58609

Platform: Windows

Policy: Contractor policy

Enforcement Level: High (Block Unapproved)

Requestor: APC-58609\bit9qa

Requestor E-Mail: JaneDoe@mycorp.com

Priority: Medium

Rule Type: Custom:Internal

Reason: I need to test software in multiple browsers.

Comments:

Resolution: Not Resolved

Status: Submitted

Related Views

- Related Events for WORKGROUP\APC-58609
- Related Events for APC-58609\bit9qa
- Related Events for firefox installer.exe
- Related File Instances for firefox installer.exe

Actions

- Approve File Locally
- Approve File Globally
- Ban File Globally
- Approve File By Policy
- Ban File By Policy
- Open Request
- Escalate Request
- Close Request

Advanced

- View Carbon Black Reputation Data

▶ Platform Analysis

▶ File Information

▶ Process Information

▶ Rule Information

▶ Request History

On the Approval Request Details page, you can examine details about the request and the requested file or device. You also can edit the request, adding comments and indicating what you did to respond to the request. The Actions menu to the right of the page provides

shortcuts to some of the App Control rules you might change if you decide to provide access to the blocked file or device.

The Approval Request Details page is divided into the following panels:

- The top panel primarily describes the request itself, including the computer and user it came from, and the App Control rules and settings relevant to the request. It also includes the user's description of the request, and provides fields for the administrator's response. A complete description of the fields in this panel is available in [Table 77, "Details for a Request/Justification \(top panel\)"](#) on page 583.
- The **Platform Analysis** panel is initially blank. If you click the **Run Analysis** button, the panel shows information about the blocked file or device, the user requesting the approval, and other data related to the request. This is additional basic information about the file and the request, not Carbon Black File Reputation data or the results of an analysis by a connected device. A complete description of the information provided by this analysis is available in [Table 78, "Platform Analysis of Requests and Justifications"](#) on page 583. You can click Rerun Analysis to update the information.
- The **File Information** panel shows the name, hash, prevalence, publisher, state, and (if Carbon Black File Reputation is activated and the file is known) trust and threat level of a file that is blocked. You can click the **Analyze** button in this panel to get more Carbon Black File Reputation information about the file. For a description of each field in this panel, see [Table 79, "File Information in Approval Request/Justification Details"](#) on page 585. Note that for device and write blocks of non-executable files, not all information will be available.
- The **Related Requests** panel shows a table of related requests; that is, requests for the same file (by hash) from different users and computers. It appears only when multiple requests have been made for the file referenced in the request details. For a description of features for handling related and duplicate requests, see ["Managing Duplicate and Related Requests"](#) on page 579.
- The **Process Information** panel shows information about the process that attempted to initiate the action. For a description of each field in this panel, see [Table 80, "Process and Installer Information in Request/Justification Details"](#) on page 586.
- The **Rule Information** panel shows information about the rule that blocked an action. If it is a custom, memory, or registry rule, a link is included to the details page for that rule so that you can modify the rule to resolve the request, if you choose. For a description how you can use this panel, see ["Opening Rule Details from the Rule Information Panel"](#) on page 578.
- The **Installer Information** panel shows information about the installer (if known) that installed a blocked file. For a description of each field in this panel, see [Table 80, "Process and Installer Information in Request/Justification Details"](#) on page 586.
- The **History** panel shows any date and time of changes to the approval request, including when it was created, opened, modified and closed. It does not include the history of changes you might make to App Control rules in response to the request.

You may expand and collapse all but the top panel using the arrows next to the panel names.

Reviewing and Resolving Requests and Justifications

Initially, the request or justification Status is *Submitted* and the Resolution is *Not Resolved*. The Status field indicates where the request stands in your work flow between *New* and

Closed. The Resolution field indicates what you did to resolve the request, including approval by different rule types or rejecting the request.

Changing the request Status to *Open* helps indicate that you have begun working on it and is required before you can modify the editable fields in a request. You can Open the request using the Action menu on Approval Requests table page, or the Open Request button or Actions menu on the Approval Request Details page.

When you have reviewed the information in a request or justification and are ready to make a decision about what to do in response, take the following high-level steps:

- Open the request to indicate that you are working on it.
- If you are not rejecting the request, make any needed file state or rule changes.
- Use the information panels on the Approval Request Details page to investigate the request.
- Update the status of the request itself and optionally make comments about your decision and actions. For example, if you determine that the request is especially important but are not ready to resolve it, you could change its status to *Escalated*. Status is for auditing purposes and also can be used to provide feedback to the requester.
- Close the request to indicate that you have finished working on it. If automatic email responses are enabled, this also sends an email to the user that made the request, indicating the decision you made.
- If automatic responses are not enabled and you choose to do so, send mail to the user requesting the approval, indicating the outcome of the request.

To open, review, and resolve an approval request:

1. On the console menu, choose **Tools > Approval Requests** and click the View Details button next to the request you want to review. The Approval Request Details page opens.
2. On the Approval Request Details page, choose **Open Request** in the Actions menu or the button at the bottom of the first panel. This activates the Comments, Resolution, and Response E-mail fields.

?
Approval Request Details (firefox installer.exe)

Id: 1

Date Requested: Jul 20 2020 04:32:42 PM

Computer: WORKGROUP\APC-58609

Platform: Windows

Policy: Contractor policy

Enforcement Level: High (Block Unapproved)

Requestor: APC-58609\bit9qa

Requestor E-Mail: JaneDoe@mycorp.com

Priority: Medium

Rule Type: Custom:Internal

Reason: I need to test software in multiple browsers.

Comments:

Resolution: Not Resolved

Status: Submitted

Open Request
 Cancel

▶ Platform Analysis

▶ File Information

▶ Process Information

▶ Rule Information

▶ Request History

Related Views

- Related Events for WORKGROUP\APC-58609
- Related Events for APC-58609\bit9qa
- Related Events for firefox installer.exe
- Related File Instances for firefox installer.exe

Actions

- Approve File Locally
- Approve File Globally
- Ban File Globally
- Approve File By Policy
- Ban File By Policy
- Open Request
- Escalate Request
- Close Request

Advanced

- View Carbon Black Reputation Data

3. If you choose to allow access to a blocked file or device, use one of the command shortcuts on the Actions menu to change one or more of the App Control rules that caused the block. For example, you might locally approve a file, edit or remove a ban, or globally approve the file.

You are not limited to the commands on the Action menu - it is possible that your response to the request will involve changes to other rules.

Note: If you address the request by creating file approvals and bans or other rule changes through the standard rule interfaces, any remediation you make does not affect the Resolution or Status fields of the request itself. You must make these changes separately. However, the Approval Requests and Justifications table and the Approval Request/Justification Request Details pages now provide commands that reduce the number of steps needed to address a request and change the Resolution and Status. See [“Request Management Work Flow Shortcuts”](#) on page 576.

- Indicate what you did (or didn't do) in response to the request by choosing from the Resolution menu in the Approval Request Details. This is for informational purposes only and does not affect file or device state. If you are not allowing access to the requested item, choose **Reject**. Note that the request status must be *Open* for the Resolution menu to be activated.

1 request(s) opened.

Approval Request Details (firefox installer.exe)

Id: 1

Date Requested: Jul 20 2020 04:32:42 PM

Computer: WORKGROUP\APC-58609

Platform: Windows

Policy: Contractor policy

Enforcement Level: High (Block Unapproved)

Requestor: APC-58609\bit9qa

Requestor E-Mail:

Priority: Medium

Rule Type: Custom:Internal

Reason: I need to test software in multiple browsers.

Comments:

Resolution: Not Resolved ▼

Status: Not Resolved Approved Cancel

Resolution Menu:

- Not Resolved
- Rejected
- Resolved - Approved
- Resolved - Rule Change
- Resolved - Installer
- Resolved - Updater
- Resolved - Publisher
- Resolved - Other

Platform Analysis

File Information

File Name:

SHA-256:

File State:

Local State: Unapproved

- Add or modify the Comments for the request to provide more detail about what you did in response to the request and why.
- If the Response E-mail address is missing or incorrect and you intend to inform the requestor of the resolution, add or correct the address while the request is still Open.
- If you are finished working on the request, choose **Close Request** in the Action menu. For multiple requests related to one file, you can choose **Close All Requests for this file**. Closing a request is primarily useful for keeping track of request status, but it also sends request status email to the user that made the request, if automatic email responses are activated.

You can re-open a request if needed.

Note: Closing a request closes all current instances of the request but does not preclude future requests for the same file.

- If automatic email notification of requestors is not activated, you can click the Response E-mail address field to open your default email client with a message pre-addressed to the requestor. If you choose to do this, fill in any details you want them to have about your response before sending.

Request Management Work Flow Shortcuts

The Approval Request work flow involves three elements:

- the Resolution you choose for the request
- the Status of the request
- the specific action taken (if any) on a file or rule to address the request

You can take an action, such as locally or globally approving a file, without changing either the status or the resolution shown for an Approval Request. Similarly, you can choose a Resolution value without actually having made any changes to rules. In some cases, you might deal with all of these elements separately.

Beginning with App Control 8.0.0, shortcuts are available that address all three of these elements with one choice. These shortcuts appear on both the Action menu on the Approval Requests and Justifications table and on the right menu on the Approval Request (or Justification) Details page. They apply to requests that can be resolved by a change in file state – they are not applicable when an action is blocked by a Custom, Memory, or Registry rule. For requests involving those rules, see [“Opening Rule Details from the Rule Information Panel”](#) on page 578.

Approval Request Details Action Menu

Approval Request Details (firefox installer.exe) ?

Id: 1
Date Requested: Jul 20 2020 04:32:42 PM
Computer: WORKGROUP\APC-58609
Platform: Windows
Policy: Contractor policy
Enforcement Level: High (Block Unapproved)
Requestor: APC-58609\bit9qa
Requestor E-Mail: JaneDoe@mycorp.com
Priority: Medium
Rule Type: Custom:Internal
Reason: I need to test software in multiple browsers.
Comments:
Resolution: Not Resolved
Status: Submitted

Platform Analysis

Related Views

- Related Events for WORKGROUP\APC-58609
- Related Events for APC-58609\bit9qa
- Related Events for firefox installer.exe
- Related File Instances for firefox installer.exe

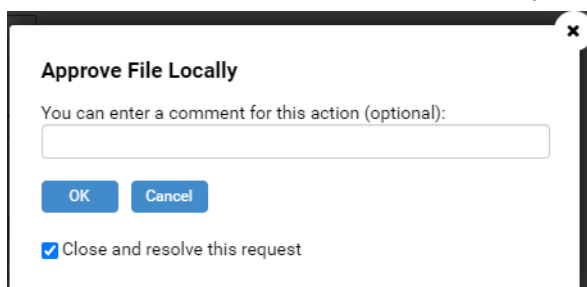
Actions

- Approve File Locally
- Approve File Globally
- Ban File Globally
- Approve File By Policy
- Ban File By Policy
- Open Request
- Escalate Request
- Close Request

On the Approval Request Details page, the right menu includes the commands Approve File Locally, Approve File Globally, and Ban File Globally. Each of these opens a dialog in which you can:

- Provide a comment describing what you did in response to the request and why you did it

- Check a box to close and resolve the request automatically



Approve File Locally

You can enter a comment for this action (optional):

Close and resolve this request

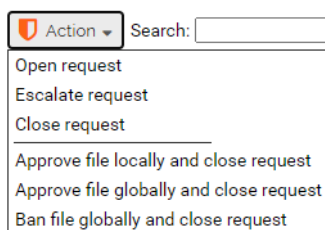
When you choose one of these commands and click OK, the changes you specified are made and the view returns to the Approval Requests and Justifications table. If you chose to close the request, email is sent to the requestor if Approval Request email is configured.

The menu choices Approve File by Policy and Ban File by Policy open the Add File Rule page where you can define the policies for which you want the Approval or Ban to apply and make any other changes available on that page. These "by policy" commands do not change the Resolution or Status of the request.

If you approve or ban a file referenced in a request, one or more of the following commands are added to the right menu on the Approval Request Details page:

- Remove Local File Approval
- Remove Global Approval/ Remove Global Ban
- Edit Global File Approval/Edit Global File Ban

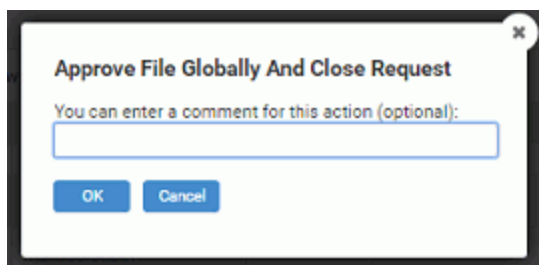
Approval Requests and Justifications Table Page Action Menu



Action Search:

- Open request
- Escalate request
- Close request
- Approve file locally and close request
- Approve file globally and close request
- Ban file globally and close request

On the Approval Requests and Justifications (table) page, the Action menu includes the commands Approve File Locally and Close Request, Approve File Globally and Close Request, Ban File Globally and Close Request. You can check the box next to one or more requests in the table and apply these commands to them. Each of these commands opens a dialog box in which you can provide a comment. If you choose OK on the dialog, the action you chose is taken and the request is closed and resolved.



Approve File Globally And Close Request


You can enter a comment for this action (optional):

Opening Rule Details from the Rule Information Panel


The Action menu on the Approval Requests and Justifications table page and the Actions menu on the Approval Requests Details page include shortcut commands to ban or approve a requested file. If an action in a request was blocked due to a Custom, Registry, or Memory rule, these shortcuts do not apply.

On the Approval Request Details page, a Rule Information panel shows the type of rule that blocked the action referenced in the request. This panel provides different information depending upon the type of rule involved:


- **Banned or Unapproved Files** – If the action was blocked because the file was banned or was unapproved in a policy that blocks unapproved files, a more generic description of the rule is provided, and there is no link. Any modifications in response to the request can be made using the shortcuts in the Actions commands in the right menu, described in “[Request Management Work Flow Shortcuts](#)” on page 576.




- **Custom, Memory, and Registry Rules** – If the rule that blocked the action was a Custom, Memory, or Registry rule, the panel shows the name and type of the rule, and the name is a link to the Edit Rule Details page.



If you open the details page for a rule and make changes, the rule information panel for the request that involved that rule shows that the rule has been modified. Modification might involve changing parameters such as the rule path, it could mean that the rule was limited to certain users or policies, or it could mean that the rule was disabled entirely.



If a rule is deleted, the rule name link is disabled (since there is nothing to link to) and “(Deleted)” is appended to the name.



Keep in mind that if you address a request by deleting or modifying a Custom, Memory, or Registry Rule, even if you do so through the Rule Information panel, you must separately change the Resolution for the request and change its status to Closed. Unlike requests that involve file state, there is no shortcut that combines all of the actions in one step.

Managing Duplicate and Related Requests

You may receive multiple requests for access to the same file. App Control separates these requests into two categories:

- **Duplicate Requests** – These are requests for approval of the same file (identified by hash) from the *same user and computer*. This might occur if a user is anxious for a resolution to a request but administrators have been delayed in providing feedback.
- **Related Requests** – These are requests for the same file (identified by hash) from *different users or computers*. This might occur if multiple users become aware of a particular application at the same time, or if a file multiple users have been using becomes blocked or unapproved.

Columns may be added to the Approval Requests and Justifications table showing duplicate and related requests.

Viewing and Resolving Duplicate Requests

For duplicate requests, since only one resolution should be necessary to handle all of them, only the first request is listed in the table and shown on the details page. If you open the details page for a request that has duplicates, an additional *Duplicates* field appears in the top panel, showing the number of duplicate requests for this file from this user on this computer. Moving the mouse over the number of requests in either the table or the details page displays a popup that gives the date, time, priority, and comments for the original and each additional request.

The screenshot shows the 'Approval Request Details (newapp.exe)' window. The main details pane lists the following information:

- Id:** 4
- Date Requested:** Nov 20 2016 07:11:45 PM
- Computer:** MYCORP\Laptop-8
- Platform:** Windows
- Policy:** Maximum Protection
- Enforcement Level:** High (Block Unapproved)
- Requestor:** MYCORP\rjones
- Requestor E-Mail:** rjones@mycorp.com
- Priority:** Medium
- Rule Type:** Unapproved executable

A popup window titled 'Original and Duplicate Requests' is overlaid on the details. It contains a table with the following data:

Date	Time	Priority	Comments
Nov 20 2016	2016 07:11:45 PM	Medium	I saw this app in a news story. It looks really cool.
Nov 29 2016	10:15:34 AM	Medium	Second request: Can I use this? I've read about it, people seem think it's safe.
Nov 30 2016	09:23:55 AM	High	Hi guys. I'd really like to use this. Any updates?

Below the table, it shows 'Duplicates: 3' and two buttons: 'Open Request' and 'Cancel'.

Viewing and Resolving Related Requests

For Related requests (same file, different user or computer), you might want different resolutions for different users, even though they involve the same file. If you include the

Related column in the table, moving the mouse over the number of related requests for any request in the table displays a popup similar to the one shown for duplicates, showing all requests related to the one in that row.

The Approval Request Details page for a request that has related requests includes a *Related Requests* panel. This panel allows you to choose one or more of the related requests to address, and includes an Action menu for that purpose.

Select	ID	Date Requested	Requestor	Reason	Priority	Computer Name	Status
<input type="checkbox"/>	10	Dec 8 2016 02:28:26 PM	MYCORP\vrjones	I need this to scan my ...	Medium	MYCORP\LAPTOP-8	Open
<input type="checkbox"/>	9	Dec 7 2016 08:20:08 PM	MYCORP\bgreen	Can I use this? It would...	Medium	MYCORP\LAPTOP-4	New

In addition, the right menu commands are modified to allow you to open, escalate, or close all of the related requests in one step. If you choose to resolve more than one of the related requests in one action and resolution email is enabled, all users who made the request receive mail when you close the request.

Notifying Users of Approval Request Resolution

You may choose to notify a user that an approval request they made has been resolved. App Control provides two ways to do this via email:

- **Manual** – You can click on the Response E-mail field on the Approval Request Details page to open a pre-configured email form in your default mailer.
- **Automatic** – You can add automatic notification to your request work flow. Automatic email notification is activated on the Mail tab of the System Configuration page. This is disabled by default.

For either method, the response mail will go to the email address (if any) that the requestor provided with their request.

Note

The automatic response features applies to Approval Requests only. No mail is sent automatically for Justifications.

To enable automatic approval request email responses:

1. On the console menu, click the configuration (gear) icon and choose **System Configuration**, and on the System Configuration page, click the **Mail** tab.
2. In the Approval Request Settings panel, check the **Mail Notification Enabled** box.

System Configuration

General Events Security Advanced Options **Mail** Licensing External Analytics Connectors Unified Management SAML Login

Mail Notification Configuration

Alert Settings

Mail Notification Enabled:

Global Subscriber Enabled:

Global Subscriber:

Approval Request Settings

Mail Notification Enabled:

Server Settings

Mail Server:

Mail Server Port:

Mail "From" Address:

Secure Mail (TLS):

Please validate settings by sending a test mail before updating the Carbon Black App Control Server.

Validate Server

Test Address:

- If you have not already configured a mail server for App Control, provide the necessary information in the Server Settings panel and validate the server by sending a message to a test address. See [“Configuring Alert and Approval Request Mail”](#) on page 754 for more details about mail server configuration.
- Click the **Update** button at the bottom of the page to save your settings.

When Notifications are Sent

After the server mail configuration is correctly configured and approval request notification mail is enabled, *closing* an Approval Request causes a mail notification to be sent in the following cases:

- The Resolution field is *changed to any Resolved option* from Not Resolved or Rejected.
- The Resolution field is *changed to Rejected* from any other option.
- The Resolution field is *Not Resolved* when an open request is closed.

Notification mail is *not* sent if the Resolution field is changed from one Resolved option to another (for example, from *Resolved - Approved* to *Resolved - Updater*).

Also, notification mail is not sent unless the Status is changed to *Closed*.

When approval request notification is enabled, notifications are not sent for requests that have already been closed. However, if a request is opened for the first time (or re-opened) after notification is enabled, the requestor will be notified if the Status and Resolution fields meet the criteria above.

The App Control Server keeps a record of request resolution mail, including a timestamp of when it was sent from the server. This is a record of mail being sent, not received. If the

email address for the recipient is incorrect, the server will still record that the message was sent. If there is no email address for the requestor, the server *does not* indicate that mail was sent.

The record of when a request response was sent appears in the Mail Sent field. In the Approval Requests table, this is an optional column that you can add using the Show Columns feature. On the Approval Request Details page, it always appears if a message was sent.

Approval Requests and Justifications							
Saved Views: All Open		Group By: (none)		View As: Individual Requests			
Show Filters Show Columns Export to CSV Refresh Table							
Action	Search:	Automatically apply		Showing 4 out of 4 item(s)			
Select	ID	Date Requested	Requestor	Reason	File Name	Mail Sent	Status
<input type="checkbox"/>	4	Nov 20 2016 07:11:45 PM	MYCORP\vjones	I saw this app in a news story. It looks really cool.	newapp.exe	Nov 20 2016 07:09:20 PM	Closed
<input type="checkbox"/>	3	Nov 20 2016 07:11:45 PM	MYCORP\dbyrd	Can I use this for organizing my mail?	mailassist.exe	Nov 23 2016 11:22:18 AM	Closed

Notification Mail Content

When approval request resolution mail is sent, it contains the following information:

- The filename for which the approval was requested
- The Resolution (i.e., the choice made on the Resolution menu)
- Any comments added by App Control administrator in the Approval Request Details.
- The reason for the request (if provided by the requestor).
- The requestor's email address
- The date of the request
- The hostname of the App Control Server

Thu 6/15/2017 11:13 AM

serveradmin@mycorp.local

Cb Protection Approval Request Response on cbp1.mycorp.local for mailassist.exe: Resolved - Approved

To  Diane Byrd

Request Approval Response

Request for "mailassist.exe" is Resolved - Approved

Approval Response: We have determined that this file is safe. Locally approved.

Request Reason: I would like to use this for organizing my mail.

Requested By: dbyrd@mycorp.com

Requested On: Jun 15 2017 3:06PM

Cb Protection Server: cbp1.mycorp.local

Approval Request and Justification Details

The following tables describe the fields on the Approval Request Details page. Other fields may be available as options in the Approval Request table.

Table 77: Details for a Request/Justification (top panel)

Field	Description
ID	A locally unique numeric identifier for the request.
Date Requested	The date and time this request was received.
Computer	The name of the computer on which the block occurred.
Platform	The platform of the computer on which the block occurred.
Policy	The Policy in effect for the agent computer at the time of the block.
Enforcement Level	The Enforcement Level of the Policy in effect for the agent computer at the time of the block.
Requestor	The user that made the request.
Requestor E-Mail	The email address (if any) provided by the blocked user.
Priority	The priority of the request (as set by the user). The options are High, Medium (the default), and Low.
Rule Type	The type of rule that blocked the action. For example, "Unapproved executable" indicates that execution of an unapproved file was blocked on a computer whose policy blocks such executions.
Reason	Approval request or justification text entered in the notifier.
Comments	Comments by an administrator reviewing the request. Can be modified and updated at any point.

Field	Description
Resolution	<p>How the request was resolved. The menu choices are:</p> <ul style="list-style-type: none"> • Not Resolved • Rejected • Resolved-Approved • Resolved-Rule Change • Resolved-Installer • Resolved-Updater • Resolved-Publisher • Resolved-Other <p>This field can be changed only when the request or justification is open. It is informational only and does not affect rules or file states.</p>
Status	<p>The status of the request. The values are:</p> <ul style="list-style-type: none"> • Submitted – A user sent the request; it has not been opened. • Open – The request has been opened by an administrator. Both Submitted and Closed requests can be opened. A request must be open for the Resolution field to be changed. • Escalated – The request has been escalated by an administrator. This might be done to draw greater attention to a high priority request. Other than the name, this is the same as Open. • Closed – The request has been closed, presumably because it has been resolved in some way. Requests can be closed even if no action has been taken to respond to them.
Mail Sent	<p>If automatic request responses are enabled and one was sent for this request, this field shows the timestamp for that mail.</p>

The Platform Analysis panel shows information resulting from clicking the **Run Analysis** button. It provides statistics about the blocked file and the user requesting access.

Table 78: Platform Analysis of Requests and Justifications

Link/Button	Comments
<number >blocks seen by this computer within 1 hour(s).	Number of blocks on this computer in one hour time period ending at the time analysis was run. Clicking this link displays Events page filtered to show all types of block events associated with this computer
<number> blocks from this process on this computer. within 1 hour(s).	Number of blocks by the given process on this computer in one hour time period ending at the time analysis was run. Clicking link displays Events page filtered to show block events associated with the process that attempted to perform the blocked action on this computer.

Link/Button	Comments
<number> files written by <the process that tried to execute this file> on this machine.	Link to Find Files page filtered to show files written by this process on this computer. Platform Note: This field appears only for files on Windows computers.
<number> files written by <the process that tried to execute this file> on the network.	Link to Find Files page filtered to show all instances of files written by this process on any computer. Platform Note: This field appears only for files on Windows computers.
File appears on <number> computers with <number> different hashes.	Search results for the name and path in the request, across all computers managed by your App Control Server. Clicking the link displays the Find Files page filtered to show all instances matching the file name and path.
<number> approval requests for this file.	The number of requests for this file, identified by <i>hash</i> . Clicking link displays the Approval Requests table filtered to show all requests for this file hash.
<number> total approval requests by this user.	Link to the Approval Requests table filtered to show all approval requests from this user.
<number> open requests by this user.	Link to the Approval Requests table filtered to show all <i>open</i> approval requests from this user.
Last Analysis Completed On <datetime> (Read Only)	Reports when the last analysis was run for this request, or if it has not yet been run.
Run/Rerun Analysis (button)	Runs an analysis that provides the information in this panel. If the analysis has already been run, reruns it to update any of the changed information, such as the number of requests from the user or the number of files written by the process that tried to write the blocked file.

Table 79: File Information in Approval Request/Justification Details

Field	Description
File Name	Clicking on link displays the File Instance Details page for the blocked file.
SHA-256	Clicking on link displays the File Instance Details page for the blocked file.
File State	The global state of this file in the File Catalog.
Local State	The local state of the blocked file instance on this computer.

Field	Description
Publisher	The publisher name and publisher approval state. Clicking on the publisher name opens the Publisher Details page for the blocked file's publisher.
File Prevalence	The number of computers on which the blocked file appears.
Trust Rating	Trust rating (if known) from Carbon Black File Reputation for the blocked file. Ranges from 0 (untrusted) to 10 (highly trusted).
Threat Level	Threat level (if known) from Carbon Black File Reputation for the blocked file. Values are 0 (Clean), 1 (Potential Risk) and 2 (Malicious).
(Security analysis results)	Assessments of the file (i.e., malicious, potential risk, or clean) from analysis on the blocked file from any connected security devices or services. This may include one or more of the following: CB Trust, CB Threat, Check Point, Palo Alto Networks WildFire, or a custom connection.

The Process tab and the Installer tab provide the same information for their subjects.

Table 80: Process and Installer Information in Request/Justification Details

Field	Description
Process	Full path to process that attempted to write or execute the blocked file.
Installer	Full path to the installer for the blocked file.
SHA-256	SHA-256 hash of the process or installer.
Trust Rating	Trust rating (if known) from Carbon Black File Reputation for the process attempting to run the blocked file or the installer that installed the file. Ranges from 0 (untrusted) to 10 (highly trusted).
Threat Level	Threat level (if known) from Carbon Black File Reputation for the process attempting to run the blocked file or the installer that installed it. Values are 0 (Clean), 1 (Potential Risk) and 2 (Malicious).

Table 81: Rule Information in Approval Request/Justification Details

Field	Description
Rule Type	<p>For actions blocked due to Custom, Memory, and Registry Rules, the rule type is composed of one of those three rule types plus the specific type chosen on the rule details page. For example, “Custom: Advanced”.</p> <p>For actions blocked due to file bans or blocking of unapproved files on agents at higher Enforcement Levels, the rule type is a generic description of the type of file blocked, for example, “Unapproved executable”.</p>
Rule Name	<p>For actions blocked due to Custom, Memory, and Registry Rules, this field displays the name given to that rule on its rule table and details pages, for example, “Protect MyApp Folder”. The name is also a link to the details page for the rule.</p> <p>For actions blocked due to file bans or blocking of unapproved files on agents at higher Enforcement Levels, this field displays the relevant setting name from the Advanced tab of the Policies page, for example, “Block unapproved executables”.</p>
(messages)	<p>If a rule was modified after the Approval Request was received, a message indicates that here. This may indicate that the rule was changed the rule in some way to allow the action indicated in the request to be completed.</p>

Customizing the Request/Justification Interface in Notifiers

You can change the text for the headings, links, and instructional text in the Approval Request panel.

Notes

- If you add any customization tags for Approval Requests and/or Justifications, you *must* enable the feature(s) using the Approval Request menu on the Edit Notifier page.
- **Platform Note:** The Approval Request/Justification interface in the Notifier History window can be customized only for Windows agents.

Table 82, “Approval Request and Justification Customization tags,” shows the tags that can be used to modify approval requests in notifiers. The example below, which is the

Notifier Text for Block unapproved executables in the Template Policy, shows where you would put tags to have different labeling for each of them.

```
<BlockText:App Control blocked an attempt by <ProcessName>
to run <TargetName> because the file is not approved. If you
require access to this file, please contact your system
administrator.><AskText:App Control identified and paused an
attempt by <ProcessName> to run <TargetName> because the
file is not approved. Choose Allow to let this file run, or
choose Block to stop it from running at this
time.<NotifierRequestLink:Submit
Justification><NotifierRequestText:Enter your reason for
access.><NotifierRequestHeading:Justification><NotifierReque
stProcessed:Justification has been submitted.> Scroll down
for diagnostic data.
```

Table 82: Approval Request and Justification Customization tags

Tag	Description
<NotifierRequestLink: <i>text</i> >	This text appears on the link that opens and closes the Approval Request panel in the notifier.
<NotifierRequestHeading: <i>text</i> >	This text appears above the text box into which the user types the request.
<NotifierRequestText: <i>text</i> >	This text appears inside the text box into which the user types the request. It disappears when the user begins entering the actual request.
<NotifierRequestProcessed: <i>text</i> >	After a request is submitted, this text appears in the text box to show that the request was processed. Note: This tag is for Justifications only. For notifiers that do not allow a choice of Allow or Block, entering an Approval Request dismisses the dialog, so this would never be seen.
<NotifierRequireSubmitOnAllow>	If present, the Allow or Approve button in a notifier is disabled until the user submits a justification.
<NotifierRequireSubmitOnBlock>	If present, the Block button in a prompt notifier is disabled until the user submits a justification.
<NotifierRequestMinLength: <i>n</i> >	If present, the Submit button in a notifier is disabled until the user enters at least <i>n</i> characters in the request/justification text box.

Chapter 21

Events, Alerts and Meters

This chapter explains how to use App Control event reports and alerts to monitor file activity and other key operations on your network. It also describes tools for detecting propagation of files on your network and for keeping track of the number of times a specified file executes.

There are many uses for these features, individually and in combination. For example, when you are allowing computers on your network to execute unapproved files, you can track the executions by file, computer, and computer user. If you are operating entirely at High Enforcement Level, you can use App Control monitoring features to be sure that files are being blocked or allowed as you want. And you can connect other monitoring features to *alerts* that will automatically tell you when certain actions occur or thresholds are passed.

See also [Chapter 22, “Monitoring Change: Baseline Drift Reports,”](#) for details on App Control’s ability to track changes in the overall inventory of files on your systems.

For information about analyzing App Control events and file information with your own tools, see [Appendix A, “Live Inventory SDK: Database Views,”](#) and the separate *App Control Events Guide* document available on the Carbon Black User Exchange. Also see [Appendix F, “Exporting Data for External Analysis,”](#) for information about exporting events to external data analytics tools.

Sections

Topic	Page
Monitoring Prerequisites	590
Event Reports	590
Viewing Reports on the Events Page	591
Taking Action on Files in Event Reports	596
Customizing Event Reports	596
Caching Events for Later Viewing	602
Using App Control Alerts	608
Creating Alerts	612
Alerts for File Prevalence	627
Monitoring Specific File Executions	629

Monitoring Prerequisites

Accurate event monitoring requires that client computers (laptops, desktops, and servers) are online and actively monitored by App Control Agents. This chapter assumes that:

- App Control policies have been created and configured.
- The agent is installed on the computers you want to monitor, and the computers have completed their initialization.
- All agents are at the latest available agent version for each OS platform in your environment.

For more information about these tasks, refer to [Chapter 5, “Creating and Configuring Policies,”](#) and [Chapter 4, “Managing Computers.”](#)

Although not a prerequisite for monitoring, if you intend to use an external event logging server, install the SQL Server on that system and configure the App Control Server to connect to the external server (see [“Setting up External Event Logging”](#) on page 734) so that you begin capturing events on the external server as soon as possible.

Event Reports

The Events page provides access to all recorded events related to App Control activities, including files blocked, unapproved files executed, system management processes and actions by console users. The App Control Server updates event data in near-real-time for connected computers, with minor variations due to event volume.

There are predefined reports, available on the Saved Views menu, and you also can create and save your own Saved Views using existing views as templates or starting with the full events table. For any event report, you can change the window of time for which you want results without having to create a new Saved View.

The Events page displays up to 200 events per page for the time period you specify. You can adjust the number of events displayed in a table by changing **rows per page** value in the bottom right of the page.

You can *cache* the events from a view you create for later examination, saving the time and processing needed to extract them from the database. When you issue a Cache command on the Events page, the events in current view are queued for processing overnight and become available on the Cached Events page when processing is complete. See [“Caching Events for Later Viewing”](#) on page 602 for details.

Notes

You can optionally choose to direct the App Control Syslog event output for postprocessing on another system. If you do so, event output also remains displayed on the Events page in the console. For more information, please refer to [Event Management Options](#) in the “System Configuration” chapter.

You also can export events to a folder for use by external data analytics products. See [Appendix F, “Exporting Data for External Analysis.”](#)

See *VMware Carbon Black App Control Events Guide*, a separate document available from Carbon Black, for a complete list of events and mapping instructions for output to supported Syslog formats.

Using the Home Page Event Reports Portlet

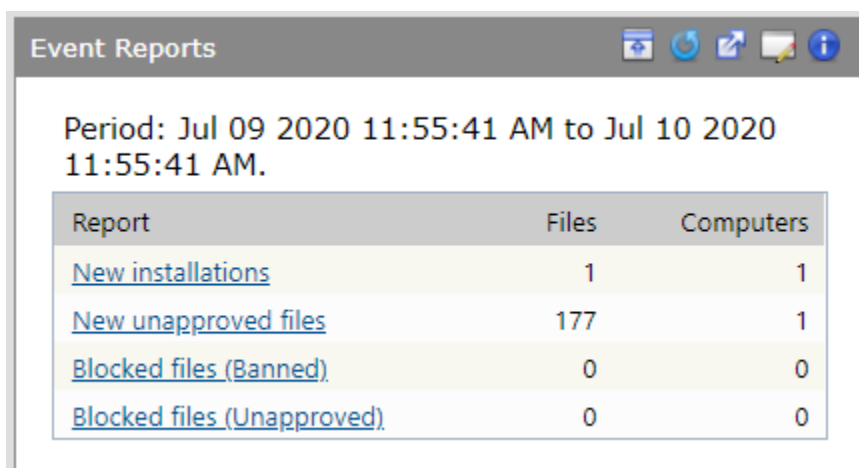
One way to monitor events is to use the Event Reports portlet on the console Home Page. The summary provides basic data from and links to the following four predefined Saved Views on the Events page – the views are described in more detail in [“Viewing Reports on the Events Page”](#) on page 591:

- New installations (*Windows only*)
- New unapproved files
- Blocked files (Banned)
- Blocked files (Unapproved)

The portlet shows the number of files and/or computers involved in events of each type over the previous 24 hours. This data is updated when you display or refresh the page, and you can get the full report by clicking on the report name.

To display the Home page daily event summary:

1. On the console menu, click **Home Page**. By default, the Event Reports portlet appears in the lower left of the page.



Report	Files	Computers
New installations	1	1
New unapproved files	177	1
Blocked files (Banned)	0	0
Blocked files (Unapproved)	0	0

2. From the Event Reports portlet, click a report name to see the full report on the Events page. See [“Viewing Reports on the Events Page”](#) for more information.

Note

You can create custom event portlets for display on the Home Page or another dashboard. See [“Using and Customizing Dashboards”](#) on page 680 for more details.

Viewing Reports on the Events Page

All event reports available on your App Control Server, whether built in to App Control or created at your site, appear as *Saved Views* on the Events page. [Table 83](#) lists the predefined Saved Views and the events they include.

Table 83: Saved Views on the Events Page

Menu Choice	Description
(none)	Displays an unfiltered view of all App Control events during the selected time period, with default columns.
Alerts and Meters	Displays a report that includes creation, modification, or deletion of alerts or meters, plus all activity that triggers an alert or increments a meter (during the selected time period).
Approval Requests	Displays a report that includes each time an approval request for a blocked file is created (on an agent computer) and opened or closed (in the console).
Blocked Files (All)	<p>Displays a report that includes all files blocked for any reason (or that would have been blocked but are in Report Only state) during the selected time period. This includes files that are explicitly banned, files in an unapproved state that were blocked because of a particular computer's Enforcement Level or policy, files that have not been analyzed yet, files on blocked devices, and files blocked because of custom rules.</p> <p>Actions blocked by registry or memory rules and certain built-in internal rules do not appear on this list.</p>
Blocked Files (Banned)	Displays a report that includes all files that have been blocked on computers running the agent during the selected time period due to an explicit ban on the file.
Blocked Files (Report Only)	Displays a report that includes all files that would have been blocked during the selected time period but are in Report Only state due to the combination of policy settings and Enforcement Level for the computer executing them.
Blocked Files (Unapproved)	Displays a report that includes all Unapproved files that have been blocked during the selected time period as a result of a policy's Unapproved Executables or Unapproved Scripts setting and its applied Enforcement Level.
Carbon Black EDR	When a Carbon Black EDR Server is integrated with a App Control Server, displays a report that includes all watchlist events from the Carbon Black EDR Server and status events from Carbon Black EDR sensors.
Computer Management	Displays a report that includes the events for the selected time period related to computers running the agent, including new and deleted computers; agent startup and shutdown; computers moved to a different policy; changes in policy's settings or Enforcement Level; and changes in the AD policy mapping rules (including their order).
Connectors	Displays a report that shows network security connector-related events, such as external notifications, malicious file detections, file analysis activity, and the addition, configuration, and removal of connector integrations.

Menu Choice	Description
Console Access	Displays a report that includes user logins and logouts, and creation, editing, and deletion of console login accounts during the selected time period.
Device Control	Displays a report that includes device-related events during the selected time period. These events include approving, banning or removing approvals or bans on devices, detection of a new device on the network, detection of attachment or detachment of a device on the network, and any device access covered by device-related policy settings. See Chapter 12, "Managing Devices," for information about devices controlled by these features.
Duplicate Computer Registrations	Displays a report that includes all events involving attempts to register more than one computer under the same agent id.
File Analysis	Displays a report that includes all events related to file analysis by external tools. This includes external notifications, file upload events, and reports of malicious or potential risk files from Carbon Black File Reputation or third-party tools.
Memory	Displays a report that includes all events related to memory (process protection) rules. Platform Note: Memory rules affect Windows systems only.
New Files (All)	Displays a report that includes all new files (i.e., not previously in the File Catalog) that have appeared on computers at your site during the selected time period.
New Files (Approved)	Displays a report of all files approved because of various reasons during the selected period. Does not include files approved because of initialization.
New Files (Banned)	Displays a list of all new banned files seen on the network.
New Files (Unapproved)	Displays a report that includes all new files that have appeared on the server during the selected time period and have not been approved or banned.
New Installations	Displays a report that includes each instance in which a file writes one or more files (creating a new file group) during the selected time period. Platform Note: Includes Windows installations only.
Registry	Displays a report that includes all events related to Windows Registry rules. Platform Note: Registry rules are applicable to Windows computers only.
Reputation	Displays a report that includes all reputation-related events, including adding or deleting a file or publisher approval based on reputation, or changes to file or publisher reputation properties.

Menu Choice	Description
Security Alert Events	Displays a report of security-alert-related events. Events include agent computers unprotected by App Control because of upgrade failures, detection or prevention of agent tampering, and a computer clock out of sync (potentially set back to attempt to defeat security measures.).
Server Management	Displays a report that includes any modifications to data on the System Configuration pages, data related to App Control database backup (success, failure, changes), server errors, Carbon Black File Reputation errors, database errors, and startup or shutdown of the App Control Server (during the selected time period).
System Health History	Displays a report that includes any changes in the severity of a health indicator as well as any creation, modification or deletion of health indicators.
Temporary Policy Overrides	Displays a report that includes each time a temporary policy override code is generated for an agent.
Threat Indicators	Displays threats detected by the ATIs in the Indicator Sets on agent-managed computers. If no Indicator Sets have been activated, this view will be empty. See Chapter 23, "Advanced Threat Detection," for more information about this and other threat-related event views.
Threat Indicators - Legacy	Displays threats detected by the ATIs that were installed in releases prior to v7.2.0. If you did not install the Detection Enhancement in a prior release, this view will be empty.
Threat Report - Suspicious Executable Created by Shell	Displays events in which certain executable files are created by cmd.exe or powershell.exe in locations such as the system directory, RecycleBin, or AppData.
Threat Report - Suspicious Files by Location	Displays events in which a file is first seen or executed on any computer, or first appears (unapproved) on at least one computer, in an unusual, suspicious location. An example would be unexpected file activity in the Recycle Bin.
Threat Report - Suspicious Files by Name	Displays events in which a file is first seen or executed on any computer, or first appears (unapproved) on at least one computer, with a suspicious name, often a name that appears similar to the name of a legitimate Windows file. For example, discovery of a file named svch0st.exe (using a zero in place of the lowercase 'o' in svchost.exe) would appear in this event view.
Threat Report - Suspicious Files by Parent	Displays events in which an unknown, or low prevalence, executable file is written by a program that should not normally be creating such files. An example of this would be an executable file created by Adobe Reader; this is often indicative of a malformed- or malicious-PDF-style attack.
Unified events	Displays events for which the source is a Unified Management server. See Chapter 27, "Unified Management of Multiple Servers" for more information about these events.

To view an existing event report:

1. On the console menu, choose **Reports > Events**. The Events page appears with the default view showing all events in the past hour:

The screenshot shows the 'Events' page with the following details:

- Saved Views:** (none) Cache Add
- Group By:** (none) Ascending
- Max Age:** 1 hour
- Actions:** Show Filters, Show Columns, Export to CSV, Access Event Archives, Refresh Table
- Search:** [Search box] Automatically apply
- Showing:** 2 out of 2 item(s)

Select	Timestamp	Severity	Type	Subtype	Source	Description
<input type="checkbox"/>	Jul 10 2020 11:58:28 AM	Info	Policy Management	Trusted Directory check	WORKGROUP\APC-79895	Trusted Directory 'c:\customer_testing\computer 'WORKGROL79895' is 'invalid'.
<input type="checkbox"/>	Jul 10 2020 11:03:22 AM	Info	Session Management	Console user login	System	User 'admin' logged in

2. Select a view from the Saved Views menu. The view appears. For views with many, and in some cases, wide columns, you might need to scroll left and right to see all the data for an event.

See [“Customizing Event Reports”](#) on page 596 for details about changing and saving reports.

Notes

- You can download event tables in CSV format.
- If an IP address is listed in an event table or description, it is the IP address of the agent computer at the time the event was reported, which is not necessarily the current IP address.

Object Previews in Events Tables

As in other tables, if an item in the Events table is highlighted, you can click on it for more details. You also can hover the cursor over many highlighted items to see an Object Preview, which provides summary information without navigating away from the page.

The screenshot shows an Object Preview popup for a file instance with the following details:

- File Name:** reader10manifest.msi
- File Path:** c:\programdata\adobe\arm\reader_10.1.9
- Computer:** MYCORP\Laptop-3
- Global State:** Approved
- Local State:** Approved

Type	Subtype	Description
Policy Enforcement	File approved (custom rule)	File 'c:\programdata\adobe\arm\reader_10.1.9\reader10manifest.msi' [EF4FF...5DD70] was approved due to custom rule.
Policy Enforcement	File approved (custom rule)	File 'c:\programdata\adobe\arm\reader_11.0.10\reader11manifest.msi' [18E93...24479] was approved due to custom rule.

Taking Action on Files in Event Reports

Whenever the details of an event identify a file, you can take action on that file directly from the Events page. To do this, you check the checkbox to the left of the event in the table and then choose an action from the Action menu. Only events containing file information can be checked.

The actions you can take on a file on the Events page are the same as those you can take on the Files page, including:

- Locally approve a file instance or remove local approval
- Globally approve or ban a file for all computers
- Create a custom approval or ban that applies to computers in specific policies
- Create a report-only ban that only reports that it *would have blocked* the file if fully enabled
- Remove an approval or ban
- View Carbon Black File Reputation data for the file

See [Chapter 8, “Approving and Banning Software”](#) for details on these file actions.

If the Connector option is installed and licensed, you also can upload files or analyze them with a third-party network security appliance. See [Appendix C, “App Control Connector”](#) and [Appendix E, “Uploading Files from Agents”](#) for details.

Customizing Event Reports

Several Saved Views are available on the Events page. In any view, you can use the Show Filter and Show Columns buttons to customize what you see, for instance, choosing to show events for a particular platform. You can also use the Show buttons to determine whether a table you are viewing has already been modified.

If you want a special report for one time use, you can simply make the customizations, view the results, and *not* save the changes. If you have made unsaved changes, a message next to the Saved Views menu reports that and offers you the option of discarding the changes. Depending upon the setting for Remember Page Settings on the console User Settings page, when you leave and return to the Events page, your view may be filtered to show these customizations even when not saved.

To save a custom report, use the Saved View panel and either save it under the an existing Saved View name (if it is not a built-in report) or under a new name.

You also can *cache* the events in a custom Saved View for later examination. Cache processing occurs overnight when it is less likely to compete with other activities. The resulting cached view loads load more quickly and efficiently than it would if you had to redo the search and extract the same data in real time. See [“Caching Events for Later Viewing”](#) on page 602.

For more information on console table features, see [Console Tables](#) in [Chapter 2, “Using the Console.”](#)


Using the Event Search Box

The Events page includes a Search box that helps you quickly locate events matching strings you enter. Search strings are matched against data in the following fields:

- File Hash
- Source

- Subtype
- Platform
- IP Address

If any data in these fields in the Events database matches the string, an auto-completion menu provides a list from which you can select the item you wanted to see.



When you choose an item from the list, the table is filtered in one of two ways:

- If you checked *Automatically apply* before entering the search screen, clicking on an option in the menu immediately filters the table to show only events matching that string in the appropriate field.
- If you did not check *Automatically apply*, clicking on an option in the menu opens the Filters panel with a filter configured to show only events matching that string in the appropriate field. You can add other filters if you choose before applying the changes to the table view.

Table 84: Event Report Fields

Field	Description
Saved View	Name for this report. If you are creating a new report, enter any text that indicates the purpose of the report in the <i>right</i> text box of Saved Views and then click Add . The report is saved and listed by its new name in the Saved Views menu with the other reports.
Maximum age	Time period of interest. You see events in the report between the time the report is run and a specified period in the past (hours, days, weeks, or months). Your choice takes effect immediately. Note that the Filters panel allows you more options for setting a time window, including Timestamp , for which the start and/or end date does not have to be the current date and time.

Field	Description
Rows per page	<p>Maximum number of events displayed on a single page in the Events table. This is controlled on a per-user basis by the <i>rows per page</i> menu in the bottom right below the table.</p> <p>Default value is 25. If your report includes more items than the <i>rows per page</i> setting, The console creates more pages and a page number panel for navigation.</p>
Group by	<p>Data field (column) by which you want to group results for default display and the sort order (ascending or descending). <i>Group by</i> creates expandable lists that initially only show the group name (for example, security policies) and number of items per group, but can be clicked to show the members of the group (for example, computers). Not all column names are available for grouping.</p> <p>The order of the groups in <i>Group by</i> (and <i>Subgroup by</i>) can be specified as one of the following:</p> <ul style="list-style-type: none"> • Ascending – Display the groups in ascending alphabetical order. • Descending – Display the groups in descending alphabetical order. • Ascending by count – Display the groups based on the number of results (rows) in each group, from fewest to most. • Descending by count – Display the groups based on the number of results (rows) in each group, from most to fewest.
Subgroup by	<p>Similar to <i>Group by</i> except it creates a second level of grouping within the first group.</p>
Show Filters	<p>Event fields you want to apply to the report. You can specify any combination of filters to determine which events are included in a report.</p> <p>Although most of the filters are for data clearly associated with the file or computer in the event, the following are special cases:</p> <p>Subtype – Subcategories of events for all event types. You can specify one or more event subtypes for display. If you select no subtype, the console searches for all.</p> <p>Severity – filter enables you to show or hide events based on standard Syslog message severity guidelines, categorized as follows:</p> <ul style="list-style-type: none"> • Critical – critical conditions • Debug – debug-level messages • Error – error conditions • Info – informational messages • Notice – normal but significant condition • Warning – warning conditions <p>Severity status for each log message is shown in the Severity column.</p> <p>Note: In previous releases, the column and filter now labeled Severity was called “Priority”.</p>

Field	Description
Show Columns	Information to be included as columns in the Events table. Use arrows to specify which columns are displayed and in what order: Items in the <i>Selected</i> list are displayed in the table. Items in the <i>Available</i> list are not displayed in the table.

To customize and save an event report as a Saved View:

1. In the console menu, choose **Reports > Events**. The Events page appears.
2. If one of the existing reports in Saved Views is similar to the report you want, choose it from the Saved Views menu. Otherwise, choose **(none)**.
3. Click in the right box of the Saved Views panel, type in a report name, and click **Add**. The new report now appears as the current Saved View and is added to the menu. Note that you also can wait until you have made all of your changes to create the new view.

Events

Saved Views: (none) Report123 Add Group By: (none)

Show Filters | Show Columns | Export to CSV | Access Event Archives | Refresh Table

Action Search: Automatically apply

4. Choose one or more filters to limit your results to certain values in specified fields. For example, you might want to have a report that only shows events whose Severity is Critical. You can specify filter fields and values in several ways:
 - Click the **Show Filters** link, choose a field (column head), and specify the value you want to filter on.
 - Click on the funnel icon in a column head and fill in the value in the Filters dialog.
 - Click on the funnel icon in a table cell to filter on that specific value.

Events

Saved Views: Report123 Delete Save Cache Add

Hide Filters | Show Columns | Export to CSV | Access Event Archives | Refresh Table

Filters

Add filter

Severity is Critical

Apply Cancel Reset

- You can add as many filters as you need. Click **Apply** when you are finished configuring filters.
5. Click the **Show Columns** link and use the arrow buttons to choose which types of data you want to display in your report, and the order in which you want them to appear. Click **Apply** when you are finished adding and removing columns.
 6. If you did not choose the time range for your report during filter configuration, choose time span from the *Maximum Age* menu.
 7. If you would like a different number of rows per page than currently shown, use the *rows per page* dropdown menu in the bottom right of the page.
 8. If you would like the data in your report collapsed into expandable group, use the *Group by* menus to choose a group and sorting order (ascending or descending by either alphabetical order or by the number of rows included in each group). For example, if you Group by Policy, the Events page initially shows Policy names, and you click on the Policy name in the table to show the events for computers in that policy.
 9. If you would like to further organize the data into expandable subgroups, choose a group and sorting order in the *Subgroup by* menus.
 10. When the report is formatted as you want it, make sure the name you want to use for it is showing in the Saved Views menu and click the **Save** button in the Saved Views panel. Your report is saved with the changes you specified.

Saved Views: *(The Current View Has Unsaved Changes - Discard)*

Report123 ▼ Delete Save Cache Add

[Show Filters](#) | [Show Columns](#) | [Export to CSV](#) | [Access Event Archives](#) | [Refresh Table](#)

You can request that events in a Saved View you create be *cached*. See [“Caching Events for Later Viewing”](#) on page 602 for more information.

Editing Event Reports

Editing a report is similar to creating one, except that you keep the same report name.

Note

The pre-defined Saved Views provided with the App Control Server are Read Only. You cannot modify them and save them under the same name; you can modify them and save them under a different name.

To edit an existing event report:

1. In the console menu, choose **Reports > Events**. The Events page appears.
2. From the Saved Views menu, select the report you want to edit. The report appears.
3. Make all of the changes you want in the report (see [Table 84, “Event Report Fields”](#) on page 597) and then click the **Save** button.

Adding Command Line Information to Event Reports

You may be interested in the command lines for processes referenced in events generated by agents. Although it is not part of the default Event Page views, a Command Line column can be added to the Events page using the Show Columns panel.

When there is a process associated with an 7.2-agent-generated event, the Command Line field will show the first 512 characters of a process command line. Pre-7.2.0 agents will not provide this information.

<input type="checkbox"/> Select 75	Timestamp	Severity	Type	Subtype	Command Line
<input type="checkbox"/>	Jul 10 2020 07:52:55 AM	Notice	Discovery	New unapproved file to computer	C:\WINDOWS\system32\svchost.exe -k wsappx -p -s AppXSvc
<input type="checkbox"/>	Jul 10 2020 07:52:55 AM	Notice	Discovery	New unapproved file to computer	C:\WINDOWS\system32\svchost.exe -k wsappx -p -s AppXSvc
<input type="checkbox"/>	Jul 10 2020 07:52:55 AM	Notice	Discovery	New unapproved file to computer	C:\WINDOWS\system32\svchost.exe -k wsappx -p -s AppXSvc

The command line shows the process that attempted the action, not the file that was acted upon. In the example above, the first two lines show that execution of a script was blocked. In the first case, a user attempted to run the script from a command prompt. In the second, the user double-clicked on the script.

To capture command line data for actions that do not normally produce events, you can add a Custom Rule to report for those actions. On the Add Custom Rule page, you choose **Advanced** as the Rule Type, **Execute** (or Execute and Write) as the Operation, and **Report Process Create** as Execute Action. Then enter the Process and Path or File information for the process that may be created by the initiating process. Actions matching the rule will report events (including command line information) upon process creation.

Important

- Command line data may include sensitive information such as passwords. While the Command Line column heading will appear to all users if added to a view, only users with specific permission will see any data in the column or in any data exported to a CSV file. This permission, which is called *View process command lines*, is not enabled by default for any of the console login account groups, and should be enabled only for users that need it. See [“User Roles and Permissions”](#) on page 89 for details about changing the permissions for a user account.
- This permission has no effect on events in Syslog. There is a separate parameter on the System Configuration/Events page that can be used to add command line data to Syslog output (off by default).
- Live Inventory SDK output always includes command line data if available.
- The potential for revealing password data in this field should be kept in mind when using the agent management commands. If you configured a password for these commands (as described in [“Configuring Agent Management Privileges”](#) on page 727), putting the command and password on one line means that the password will be included in command line field for an event. Carbon Black recommends that you enter the agent management command alone, and then provide the password at the prompt that follows.

Caching Events for Later Viewing

You might find it useful to review a time-boxed collection of events meeting certain characteristics, either on a regular basis or just once. For example, you might want to review the executable files most often accessed in your organization and determine which ones need to be approved to expedite movement to High Enforcement. You can use filters and Saved Views to specify the events you want to review, but if you have a particularly large number of events or a complex set of filters, fetching these events for viewing could be time consuming and possibly cause the server to time out.

To make the review of a set of events more efficient, App Control allows you to *cache* the events in a custom Saved View. Cache requests are queued for overnight processing (beginning at approximately 12:30AM local time), when there is less likely to be as much of a load on the server as during prime working hours.

The results of cache processing appear as a new, named view on the Cached Events page (separate from the Events page). Because these events are in a cache rather than being fetched from the database in real time, access is much faster and further filtering of the cached view should also be much faster. Once you request an event cache, the view you defined is used to create a new cache every night until you remove it from the Cached Events page.

The time period covered by the events in a cache depends upon several factors:

- Any time periods you define are based on when the event was recorded on the server (the server “timestamp”), not on when an event occurred on an endpoint.

- If no other time constraints are defined, the cache includes events recorded on the server *up to the time the cache processing begins*. For example, if you click the Cache button at 4PM and processing begins at 12:30AM the following day, all events prior to 12:30AM are included.
- If you specify beginning and ending timestamp filters in the view you cache, these determine the events included in the cache. However, event pruning may remove older events that are included in your timestamp filter period. See [“Managing the App Control Event Database”](#) on page 733.
- If you include a Max Age property for the view, the cache includes events recorded on the server *up to the time the cache processing begins* and going back for whatever time period you choose. For example, if you set Max Age to ‘1 Day’, the resulting event cache includes all events from 12:30AM yesterday to 12:30AM today. Each successive processing deletes the oldest day and adds the latest day to the results.

To create an event cache for later viewing:

1. In the console menu, choose **Reports > Events**.
2. Configure the view that includes the events you want to cache, using any of the following tools that help to refine that view:
 - **Existing Saved Views** – If any existing Saved View matches or is similar to the view you want, you can start with that view to create your cached event view. For example, you might choose the “New Files (Unapproved)” view if you are in Low or Medium Enforcement and want to see files on your users’ endpoints that should be approved.
 - **Filters** – Use either the Show Filters link or the funnel icon in a table column or cell to add or modify the view using any of the filter categories.
3. Determine the time period for events you want to cache. You can do this in one of two ways:
 - **Max Age** – You can use the Max Age field to designate the length of the time period for which you will cache events. If you use Max Age and no other Timestamp filters, the end of the time period is always the time of cached event processing (not the time that you clicked the Cache button). For example, if you choose 1 day for Max Age and the cache processing occurs at 12:30AM, your cached events will include events from 12:30AM the previous day until 12:30AM the day the cache is processed.
 - **Timestamp Filters** – If you have a more specific time period (both beginning and ending) that you want to cache events for, use the Filters panel and set both before and after Timestamps.
4. If you added any filters to the view, click the **Apply** button in the Filters panel.
5. When you have defined the event data you want cached, create a Saved View for that data by entering a name in the box to the left of Add and clicking **Add**. This name will appear on the Saved Views menu on the main Events page and will also be the name for these events on the Cached Events page.
6. While this view is still showing, click the **Cache** button.

The events in the view are queued for caching. Cache requests in the server queue are run at each night at approximately 12:30AM local time and the results are available on the Cached Events page the next day.

Once you create an event cache, the view you defined for that cache is run every night until you remove it from the Cached Events page. If you defined that view using the Max


Age property, each new processing deletes the oldest day and adds the latest day to the results.

You cannot edit or rename a cached view. If you need to modify the view, create a new one on the events page and delete the old one.

Viewing and Taking Action on the Cached Events Page

When the initial processing of a cached event view has completed, its name appears on a menu on the Cached Events page.

Info: Actions such as banning and approving files on this page will not be reflected on this page until the cache task runs again.



Choosing one of the views displays the latest results for that view. A legend at the top of the view indicates the time period covered by these results.

Timestamp	Type	Subtype	Description
Sep 15 2019 11:45:37 PM	Discovery	New unapproved file to computer	Computer MYCORP\DESKTOP-12 discovered new downloads\firefox installer.exe' [37D0A...C0704]. Di FileCreated[9/15/11:45:11 PM] Discovered[9/15/2019 11:45:37 PM] Publisher[Mozilla Corporation (Eigib)]
Sep 15 2019 11:15:01 PM	Discovery	New unapproved file to computer	Computer MYCORP\LAPTOP-7 discovered new file local\microsoft\windows\inetcache\ie\eu20s9he\arm DiscoveredBy[Kernel:Write] FileCreated[9/15/2019 11:15:01 PM (Hash: 9/15/2019 11:15:01 PM)] Publi

You can use the Show Filters and Show Columns panel to further customize the data displayed in a cached view. There are several differences in behavior for these features between a Cached Events page view and a view on the main Events page:

- You can use the Show Filters to further filter the cached view. If you add a filter, however, keep in mind that you are filtering only on the data included in the cached view.
- You can use the Show Columns panel to add or delete columns in a cached view. The column changes persist while you remain on the page – if you navigate off of the page and then reload the cached view, it returns to the column layout it had when it the cache request was first made.
- You cannot modify the filters that were used to create the cached view – they are grayed out and inoperative.
- You can use the Timestamp filter to further restrict the data shown in the cached view, but you cannot extend it. Also, there is no Max Age option on the Cached Views page.

The Cached Events page includes an Action menu similar to the one on the regular Events page. You can take actions such as approving or banning files shown in events

listed in the cached view. As with filters, there are some differences between the Events page and the Cached Events page.

If you take actions while on the Cached Events page, keep in mind that the events shown in the cached view reflect conditions at the time the cache request was made. Event information, such as file state, rules being enabled or disabled, or computer properties, does not change due to an action you take on the Cached Events page.

For example, if an event on the Cached Events page shows that a file was blocked because it was unapproved, you can approve that file, but in the Cached Events view the event will still show it as unapproved because that was its state when the event occurred. If a file name or hash shows as a link in the Cached Events view, however, you can see the new file state in the object preview if you hover over the file hash, and its new state will appear on the Files and File Details pages. In addition, actions affecting file state you take on the Cached Events page will appear on the main Events page as new events.

The screenshot displays the 'File' details page. On the left, a table cell shows the state 'Unapproved'. A red box highlights this cell, with a red arrow pointing to the 'Global State: Approved' field in the file details. Another red box highlights the 'Unapproved' state in the table, with a red arrow pointing to the 'Global State: Approved' field. The file details include:

- File**
- SHA-256: 972091534687ce264181333doda7d732a972081334a4bc72afd1d69ef0e33de0
- SHA-1: f926047a616f7edc0df46f9260f0dbf216f1ebeb
- MD5: d46885749f10d84be620d468899264ac
- First Seen Name: installme.msi
- Publisher: Global State: Approved
- First Seen Computer: MYCORP\LAPTOP-8
- Prevalence: 0
- Trust: Unknown
- Threat: Unknown

Below the details, a table row shows the event state as 'Unapproved'.

Removing an Event Cache

Once you create an event cache, the view for that cache is run every night until you remove it from the Cached Events page. Removing caches you no longer need eliminates unnecessary nightly processing and the additional data footprint of the cache.

To remove an event cache:

1. In the console menu, choose **Reports > Cached Events**.
2. On the Cached Events page, use the **Cached View** menu to choose the cache you want to remove.

- Click the **Remove Cache** button. The cache is removed from the Cached Events page and will no longer be processed.

Cached Events

Cached Views: Cert Events **Remove Cache** Group By: (none) Ascending Subgroup By: (none)

This view is showing all data before Sep 13 2019 12:30:37 AM

Show Filters | Show Columns | Export to CSV | Refresh Table

Action Search: Automatically apply Showing 75 out of 161 item(s)

Select All	Timestamp	Severity	Type	Subtype	Source
<input type="checkbox"/>	Sep 12 2019 01:26:50 PM	Info	Discovery	Certificate checked	System
<input type="checkbox"/>	Sep 12 2019 01:26:50 PM	Info	Discovery	New certificate on network	System

- Click the success banner to refresh the Cached Events page (without the removed view).

Viewing Install Event Details

If an event subtype is highlighted, the event has other events associated with it. Clicking on a highlighted event subtype brings you to an Install Event Details report, which shows all of the sub-events associated with the event you clicked (per computer). The Details report is useful primarily to show the *connections* between a root event and the events it generates.

Note

It is *events* generated by the root installation *event* that are reported here, not *files installed* by an *installer*. Whether installation of a file generates an event depends on the approval status of the installer, and may also depend upon the security policy on the computer where the files are being installed and other rule and configuration settings that can exclude file tracking. Events include information such as process name and user running the process.

Approved installers generate locally approved files, and approved files do not generate sub-events on the Install Event Details page. *Unapproved* installers generate unapproved files (unless previously approved by some other means), and unapproved files do generate sub-events. Also, any newly installed files that are blocked generate Install Event Details.

Viewing Event Archives

The **Access Event Archives** link on the Events page opens a table of daily archives for App Control events. These events are archived in CSV files.

Event Log Archives		
Events are archived daily to a compressed CSV file format.		
Archive name	Archive date	Archive size
20170522.csv	May 25 2017 04:07:37 PM	490.3 MB
20170518.csv	May 22 2017 12:48:05 PM	500.0 MB
20170512.csv	May 18 2017 03:44:31 AM	500.0 MB
20170509.csv	May 12 2017 08:15:29 PM	500.0 MB

You can open or download any day's event archive by clicking on the CSV file name and making your choice of action from the dialog box. These archives are located in the "archive logs" folder under your App Control Server installation directory.

To return to the Events page, choose **Reports > Events** in the console menu.

Notes

- Archiving can be enabled or disabled on the Events tab of the System Configuration page. See ["Managing the App Control Event Database"](#) on page 733 for more information.
- Unlike event times shown in the console, timestamps for the archived events listed in the CSV files are shown in UTC time.

Using App Control Alerts

Alerts notify you of important App Control-monitored activities, such as the appearance or spread of risky files on your computers. When conditions specified in an alert are met, notifications can be provided in the following ways:

- **Notification in Console Banner** – If any alerts are triggered, indicators appear on all console pages, in the upper right above the console menu. There are three different symbols that can appear in the banner, each representing a different alert priority, and the number of triggered alerts in each category is shown to the right of its symbol. See [“Alert Priority”](#) on page 618 for more on details on priority.

Hovering the mouse cursor over a symbol or the number to its right shows a tooltip describing either the type of alert (if there is only one in that priority) or its priority. Clicking on the symbol or number opens the Alert Instances page if one alert is triggered or the Alerts page filtered to show those alerts if more than one alert is triggered at that priority level.



- **Email Notification** – Email notification about the event(s) triggering the alert goes to a list of subscribers.
- **Alerts Page Row Highlighting** – On the Alerts page, the row for each triggered alert is highlighted, with the highlight color indicating the alert priority (red for high, orange for medium, yellow for low).
- **Home Page and other Dashboards** – All currently triggered alerts appear in the Alerts portlet, which is part of the default console Home Page and can be added to other Dashboards. This portlet also uses the color and symbol coding for alert priority.

You can *reset* an alert when you no longer want to be notified about it. This removes the warning banners on the Alerts and Home pages (and any dashboard with the Triggered Alerts portlet), and if you have enabled automatic re-sends of alert email, it stops those. If the conditions that triggered the alert occur again, another alert will be triggered. If the conditions that caused the Alert cease to exist, the Alert will be auto-reset to a non-triggered state (see [“How Alerts are Triggered”](#) on page 619 for details).

An Alert History is kept for each alert, and this history is modified as alerts are triggered and reset.

Note

Access to alert features is determined by the *View alerts* and *Manage alerts* permissions on the Login Accounts Add/Edit Group pages.

There are two top-level classes of alerts:

- **Built-in Alerts** – [Table 85](#) shows the alerts preconfigured and listed by default in the console.
- **User-Created Alerts** – You can create and edit alerts through the Alerts page. This is described in [“Creating Alerts”](#) on page 612.

The Alerts page lists all currently available alerts, including built-in and user-created, and both enabled and disabled.

Alerts					
Group By: Priority ▾ Ascending ▾ Show Filters ▾ Show Columns ▾ Export to CSV Refresh Page					
Action ▾ ⏪ 1 2 ⏩ + Add Alert					
<input type="checkbox"/>		Name	Type	Enabled	Priority
Priority: High					
<input checked="" type="checkbox"/>		System Health OER Alert	System Health Alert	Yes	High
<input type="checkbox"/>		System Health Infrastructure Configuration Alert	System Health Alert	Yes	High
<input type="checkbox"/>		System Health Rules Alert	System Health Alert	Yes	High
<input type="checkbox"/>		System Health Environment Alert	System Health Alert	Yes	High
<input type="checkbox"/>		System Health Product Configuration Alert	System Health Alert	Yes	High
<input type="checkbox"/>		System Health Backlog Alert	System Health Alert	Yes	High
<input type="checkbox"/>		Database Limit Alert	System Alert	Yes	High
<input type="checkbox"/>		Backup Missed Alert	System Alert	Yes	High
<input type="checkbox"/>		Database Verification Failed	System Alert	Yes	High
<input type="checkbox"/>		Agent Install Package Generation Disabled	System Alert	No	High
<input type="checkbox"/>		Host Package Not Found (Windows)	System Alert	Yes	High
<input type="checkbox"/>		Host Package Not Found (Mac)	System Alert	Yes	High
<input type="checkbox"/>		Host Package Not Found (Linux)	System Alert	Yes	High
<input type="checkbox"/>		Default Rules Not Found	System Alert	Yes	High
<input type="checkbox"/>		Malicious File Detected	File Security Alert	Yes	High
<input type="checkbox"/>		Cb Collective Defense Cloud Unavailable Alert	System Alert	No	High
Priority: Low					
<input type="checkbox"/>		Updater Modified Alert	System Alert	Yes	Low
<input type="checkbox"/>		Justification Alert	Approval Request Alert	No	Low
<input type="checkbox"/>		New Certificate Alert	Certificate Alert	No	Low
<input type="checkbox"/>		Indicator Set Alert	Event Alert	No	Low

Table 85: Built-in Alerts

Alert	Description
Database Limit Alert	Alerts subscribers when SQL Express database size reaches its specified limit (varies depending upon SQL edition). Only active if you have installed SQL Server Express edition (not a full SQL version). Always enabled (cannot be disabled).
Backup Missed Alert	Alerts subscribers when Database backup was scheduled but missed. Enabled by default, but can be disabled.
Database Verification Failed	Alerts subscribers when the App Control database is found to be corrupt. If triggered, contact Carbon Black Support. Always enabled (cannot be disabled).
Potential Risk File Detected	Alerts subscribers when a file on an computer monitored by an agent is considered potentially malicious by Carbon Black File Reputation or another connected security device or service. Disabled by default.

Alert	Description
Malicious File Detected	Alerts subscribers when a file on a computer monitored by an agent is considered malicious by Carbon Black File Reputation or another connected security device or service. Can be configured to ignore banned and/or approved files. Disabled by default.
Elevated Privilege: Install Mode	Alerts subscribers when any computer remains in local approval mode longer than a specified time period. The default is 1 hour, but can be modified. No computer should remain in approval mode longer than is necessary to install software.
Carbon Black File Reputation Unavailable Alert	Alerts subscribers when expected Carbon Black File Reputation tasks are not performed during a period of time specified in the alert. The default period is three hours, but you can modify this. Enabled by default if integration with Carbon Black File Reputation is activated (and cannot be disabled). Disabled if Carbon Black File Reputation integration is not activated. Once triggered, the alert remains in effect until all standard Carbon Black File Reputation tasks are restored to normal operation. It can be manually reset, but will trigger again after the specified period if the conditions that caused the alert still exist. The conditions that trigger this alert also add a notification that Carbon Black File Reputation is unavailable to the System Configuration/Licensing page.
Approval Request Alert	Alerts subscribers when more than the specified number of approval requests are in New or Open state. Requests older than one week and Closed requests are not considered when triggering the alert. Once triggered, the alert remains in place until it is manually reset or enough requests are Closed to bring the total below the threshold. Enabled by default.
Justification Alert	Alerts subscribers when more than the specified number of justifications are created for files that endpoint users chose to allow to run. Justifications older than one week are not considered when triggering the alert. Once triggered, the alert remains in place until it is manually reset or enough justifications are Closed to bring the total below the threshold. Enabled by default.
Updater Modified Alert	Alerts subscribers when an updater is created, modified or deleted by Carbon Black File Reputation. Always enabled (cannot be disabled). Note: Automatic updater management by Carbon Black File Reputation must be enabled on the Advanced Options tab of the System Configuration page.
Rapid Config Alert	Alerts subscribers when a Rapid Config is created, modified or deleted by Carbon Black File Reputation. Disabled by default. Note: Automatic Rapid Config management by Carbon Black File Reputation must be enabled on the Advanced Options tab of the System Configuration page.

Alert	Description
Computer Security Alert	<p>Alerts subscribers when suspicious behavior is detected on a computer. Triggering conditions include detection of a computer that is unprotected due to an upgrade failure, agent tampering detected or prevented, and a computer clock out of sync with the App Control Server. Always enabled (cannot be disabled).</p> <p>See “Detecting Agent Issues with Computer Security Alerts” on page 625 for more details on these alerts and the conditions that cause them.</p>
New Certificate Alert	<p>Alerts subscribers when a file with a certificate for a publisher not yet listed in the console is discovered, or a new certificate is imported directly into the App Control Server. By default, this alert is triggered when a new certificate for <i>any</i> publisher is detected. However, it can be configured to trigger only for new certificates for specific publishers.</p> <p>If set to Specific Publisher, you must provide a string that matches all or part of the name of the publisher for which you want alerts. For example, if you provide “Apple” as the string, it will alert you about new certificates whose publisher is identified as “Apple”, “Apple, Inc.”, “Big Apple, Ltd.”, etc.</p> <p>You can add multiple publishers (or partial names) to the alert.</p> <p>Requires v7.0.1 or later agent. Disabled by default.</p>
Revoked Certificate Alert	<p>Alerts subscribers when a certificate known to this App Control Server is revoked. By default, this alert is triggered when a certificate for <i>any</i> publisher is revoked. However, it can be configured to trigger only for specific publishers.</p> <p>If set to Specific Publisher, you must provide a string that matches all or part of the name of the publisher for which you want alerts. For example, if you provide “Apple” as the string, it will alert you about revoked certificates whose publisher is identified as “Apple”, “Apple, Inc.”, “Big Apple, Ltd.”, etc.</p> <p>You can add multiple publishers (or partial names) to the alert.</p> <p>Requires v7.0.1 or later agent. Disabled by default.</p>
Indicator Set Alert	<p>Alerts subscribers when a detection indicator set is created, updated, or deleted.</p>
System Health OER Alert	<p>Alerts subscribers when the environment for this server is out of compliance with App Control <i>Operating Environment Requirements</i>, which can indicate immediate or potential performance issues.</p> <p>Note: This alert only appears and can only be triggered if System Health Indicators are enabled on the Advanced tab of the System Configuration page and this indicator has been downloaded to the server. If present, it is always enabled.</p>

Alert	Description
System Health Infrastructure Configuration Alert	Alerts subscribers when the conditions in your server environment trigger a Health Indicator on the Infrastructure Configuration tab of the System Health page. Note: This alert only appears and can only be triggered if System Health Indicators are enabled on the Advanced tab of the System Configuration page and this indicator has been downloaded to the server. If present, it is always enabled.
[Sample] Windows File Properties	Alerts subscribers when an the <i>Report write (custom rule)</i> occurs and triggers the Windows File Properties Indicator Set for threat detection. Disabled by default.

Creating Alerts

You can create and configure alerts of the following types:

Table 86: User-Creatable Alert Types

Alert Type	Description
File Activity: Propagating File	Alerts subscribers when a <i>locally unapproved</i> file appears on more than a percentage of computers for the policies and time period you specify. If you are not operating in High Enforcement, propagating files can indicate a spreading virus.
File Activity: Blocked File	Alerts subscribers when the same file is blocked on more than a specified percentage of computers for the policies and time period you specify.
Baseline Drift Alert	Alerts subscribers when baseline drift of files reaches the specified threshold.
File Prevalence Alert	Alerts subscribers when a <i>specified</i> file is present on more than a specified number of computers.
Event Alert	Alerts subscribers when specified events occur, or a specified event rule is triggered, more than a threshold number of times in the specified time period.

To create an alert:

1. In the console menu, choose **Tools > Alerts**. The Alerts page, which lists all currently available alerts (both enabled and disabled), appears:
2. From the Alerts page, click the **Add Alert** button. The Alert Information page appears:

Add Alert

General

Alert Name:

Message:

Priority:

Status: Enabled Disabled

Type

Type:

Description: Alerts subscribers when a specified event(s) or event rule(s) triggers it

Mail Template:

Criteria

Threshold:

Time Period:

Trigger On: Event(s) Event Rule

Select Event Properties

Add filter

Subtype

Subscribers

Note: Alert must be created before email recipients can be specified

Reminder Mail

Status: Enabled Disabled

Remind Every:

Auto Reset

Status: Enabled Disabled

Reset After:

3. In the Alert Information panel, enter the information requested. See [Table 87](#) below for details on the fields you can specify.
4. When you have finished defining the alert, click either **Create**, to create the new alert and stay on this page, or **Create & Exit** to return to the table of alerts. You might use Create if you want to add subscribers to this alert.

Once created, the new alert appears on the Alerts page. If the alert is Enabled, it begins monitoring activity on your network and will trigger if it finds conditions matching the definition you set up.

Table 87: Alert Fields

Section	Field	Description
General	Alert name	Name for the Alert as you would like it to appear in the Alerts table.
	Message	Message to be sent when alert is triggered. You can add tags to the message for an Event Alert so that it provides data specific to the alert instance. See “Informational Tags for Event Alert Messages” on page 617
	Priority	Priority level assigned to this alert. The choices are: High, Medium, Low. Priority level determines the color assigned to the alert in the user interface, and allows you to group alerts by priority to highlight the most critical items.
	Status	Specifies whether the alert is enabled (on) or disabled (off). Note that if you disable an alert after it is triggered, this does not automatically reset the alert.
Type	Type	Type of alert you want to configure: <ul style="list-style-type: none"> • File Activity: Propagating File • File Activity: Blocked File • Baseline Drift Alert • File Prevalance Alert • Event Alert
	Description	Read-only text with more information about the specified alert Type.
	Mail Template	Template you want to use to determine the format and content of the email you send subscribers of this alert. The default template can be used for any alert, but the other standard templates may be more appropriate for the alert type they represent: <ul style="list-style-type: none"> • Default • Template for File • Template for Elevated Privilege • Template for Approval In addition, you can create custom templates if you choose. You may also find guidance on template creation on the Carbon Black User Exchange.
Criteria: File Activity and Prevalance alerts	Threshold	Threshold of affected computers required to trigger the alert. Appears only if applicable to the alert type. This can be a percentage or an absolute number.
Criteria: File Activity alerts	Time Period	Minimum time period within which activity must occur to trigger the alert. Appears only if applicable to the alert type.

Section	Field	Description
Criteria: Baseline Drift alerts	Drift Report	Name of the drift report whose data you want to analyze to trigger alerts. Appears only if applicable to the alert type.
	Alert When	The drift parameter you want to measure and the threshold at which it triggers an alert. Appears only if applicable to the alert type.
Criteria: File Prevalence alerts	Specify File By	The way you want to identify a file – the choices are Hash and Filename .
	File Name	Filename to monitor for the alert. Appears only if you chose filename for <i>Specify file by</i> . Note: You cannot use wildcards in the file name for a prevalence alert.
	Publisher Contains (optional)	The name of the publisher (if any) identified as the source of the file. Appears only if you chose filename for <i>Specify file by</i> .
	Hash Type	The type of Hash (MD5, SHA-1 or SHA-256) you are using to identify the file. Appears only if you chose Hash for <i>Specify file by</i> .
	Hash Value	The hash value of the file. Appears only if you chose Hash for <i>Specify file by</i> value type.
Criteria: Event Alerts	Threshold	Number of times an event or event rule must match the properties defined in this rule during the specified time period to trigger an alert.
	Time Period	Time period during which the conditions defined in this rule must be met to trigger an alert.
	Trigger On	Specifies whether the alert is triggered by Event(s) or an Event Rule .
	Select Event Properties	If you chose to trigger on Event(s), the properties of the event(s) that will trigger this alert. The properties include: <ul style="list-style-type: none"> • Subtype – A rule set to trigger on events must include at least one subtype, and may contain more than one. • Other properties – The Add filter menu includes other event parameters that may be added to more narrowly specify the conditions under which an alert is triggered.
	Select File Properties	If you chose to trigger on Event(s), you may optionally add properties that a file mentioned in the event must meet to trigger this alert. It is not necessary to include file properties, but if specified, the alert will not trigger if the property specified does not match the rule or if the value of property is unavailable for the event.

Section	Field	Description
	Select Process Properties	If you chose to trigger on Event(s), you may optionally add properties that the parent process of the file specified in file properties must meet to trigger this alert. It is not necessary to include process properties, but if specified, the alert will not trigger if the property specified does not match the rule or if the value of property is unavailable for the event.
	Event Rule	If you chose to trigger on Event Rule, an Event Rule menu lists the existing rules.
Policies (appears only for appropriate alert types)	Rule Applies To	Click the radio button to activate this alert for All policies or Selected policies . For <i>Selected policies</i> , check the box next to each policy for which you want the alert enabled.
	Selected	Policies that will be subject to this alert. Select policies and use the arrow buttons to move them into the appropriate column.
Subscribers	Email	Note: You cannot add subscribers (the fields do not appear) until after the alert is created. Add all email addresses to which you want alert notifications sent. Enter each address in the <i>Email address</i> box, and click the Add button each time to create a subscriber list. Add is enabled when you enter a qualified email address. The dropdown menu to the right of the address box specifies the format of notification email. The choices are: text , HTML , or Auto . Auto allows the recipient's mail server to define the format.
Reminder Mail	Status	Reminder Mail status determines whether alert email is resent after a specified period of time when the alert has <i>not</i> been reset. The choices here are Enabled or Disabled .
	Remind Every	When Reminder Mail is enabled, the amount of time between alert email re-sends for alerts that are not reset.
Auto Reset	Status	Auto Reset determines whether an alert will be reset automatically, either after a specified time period or, for certain alerts, when conditions that triggered it are no longer in effect. When Enabled , alerts may be auto-reset. When Disabled , alerts must be reset manually.
	Reset After	If Auto Reset is enabled, this setting determines the time period after which a triggered alert instance will auto-reset if it has not already been reset for another reason. The default value is 4 weeks. It may be changed to a different period, ranging from minutes to weeks.

Informational Tags for Event Alert Messages

The Alert Message can provide additional documentation for you and others about the conditions that triggered an alert. For Event Alerts, you can add tags to the message so that it provides data specific to the alert instance. [Table 88](#) shows the available tags.

Table 88: Informational Tags for Event Alert Messages

Tag	Description
<FileName>	Name of the file from the event initiating the alert. If multiple files led to the alert, contains a comma-separated list.
<Sha256>	SHA-256 hash of the file from initiating event. If multiple files led to the alert, contains a comma-separated list.
<Md5>	MD5 hash of the file from the initiating event. If multiple files led to the alert, contains a comma-separated list.
<Sha1>	SHA-1 hash of the file from initiating event. If multiple files led to the alert, contains a comma-separated list.
<RootSha256>	Root SHA-256 hash of the file from the initiating event. If multiple files led to the alert, contains a comma-separated list.
<HostName>	Name of computer from the initiating event. If multiple computers led to the alert, contains a comma-separated list.
<UserName>	Username from initiating event. If multiple users led to the alert, contains a comma-separated list.
<EventRuleName>	If an event rule initiated the alert, the name of the rule.
<EventRuleDescription>	If an event rule initiated the alert, the description of the rule.
<EventSubtype>	The subtype of the initiating event. If multiple events led to the alert, contains a comma-separated list.
<EventDescription>	Description field from the initiating event.
<AntibodyId>	ID of the file from initiating event. If multiple events led to the alert, contains a comma separated list.
<HostId>	ID of the host from the initiating event. If multiple events led to the alert, contains a comma separated list.

Editing Alerts

You may need to modify an alert to change its threshold, the time period it covers, its subscribers, or some other parameters. In addition, you may need to enable or disable the alert. All of this is done through the Alert Information page.

To edit, enable or disable an alert:







1. If you are not already on the Alerts page, choose **Tools > Alerts** in the console menu.
2. Click the View Details button next to the alert you want to modify. The Alert Information page appears.
3. If you only want to enable or disable the alert, click the appropriate button in the General section of the Alert Information panel and then click the **Save** button at the bottom of the page.
4. If you want to make other changes, edit the appropriate parameters (see [Table 87](#)) and then click **Save**. The alert is updated and you return to the Alerts page.

Although you can't create new instances of built-in alerts, you can edit some of their settings. For example, you can change the number of approval requests necessary to trigger an Approval Request alert. You also can modify which actions (creation, editing, deletion) trigger an Updater Modified alert.

Alert Priority

Each alert is assigned a Priority, which can be High, Medium, or Low. Alert priority determines the icon shape and color used to represent an alert in the console banner and dashboard, and the color of the rows for triggered alerts on the Alerts page.

Table 89: Alert Priorities

Alert Priority	Icon	Row Color
High	 or 	Red
Medium	 or 	Orange
Low	 or 	Yellow

In addition to providing a visual cue that one alert is more important than another, alert priorities also allow grouping on the Alerts page, making it easier to give attention to the most important alerts first. When you choose Priority on the Group by menu, alerts are sorted first by Priority and then by Date Triggered, in descending order.

System alerts have predefined priority levels that cannot be changed:

- Database Limit Alert – High
- Database Verification Alert – High
- Carbon Black File Reputation Unavailable Alert – High

For other alerts, you can change priority using the Action menu on the Alerts table page or the Priority menu on the Add/Edit Alert page.

Deleting Alerts

When you delete an alert, you delete the definition of the alert and end any monitoring you have been doing with it. As an alternative, you can disable an alert if you don't want it to be active but might use it in the future. You cannot delete the some pre-defined alerts provided by the App Control Server, and these do not have a Delete button.

To delete an alert:

1. On the Alerts page, click the Delete (x) button next to the alert you want to delete.
2. On the confirmation dialog box, click **Yes**.

How Alerts are Triggered

Any alert shown on the Alerts page, whether it was built-in or created by you, can be considered an *alert class*. Each time conditions exist that meet the triggering condition of that alert class, an *alert instance* occurs. For some alert classes, it is only possible to have one instance. For example, there is only one database for a App Control Server, and so the Database Limit Alert can have only one instance at a time. For other classes, there can be many instances simultaneously. For example, there might be multiple malicious files on a network, and so there could be multiple Malicious File Detected alert instances.

When any triggered instances of an alert class exist, the alert is highlighted on the Alerts page using the color-coded severity level, and a Reset button is added next to the alert name. The Date Triggered column shows when the alert was triggered, and the Instances column shows the number of triggered instances and links to the Alert Instances page. By default, triggered alerts appear at the top of the page, in descending order of when they were triggered. This includes alerts that have been reset.

The console does not display new banners for each alert *instance* during a console login session, but the number of instances is shown. The view below shows triggered alerts without grouping and sorted by Date Triggered.

Alerts						
Group By: Priority Ascending						
Show Filters Show Columns Export to CSV Refresh Page						
Action Add Alert						
	Name	Type	Priority	Date Triggered	Instances	
<input type="checkbox"/>	System Health OER Alert	System Health Alert	High	May 08 2017 09:03:11 AM	2	Reset
<input type="checkbox"/>	Updater Modified Alert	System Alert	Low	May 08 2017 08:28:46 AM	2	Reset
<input type="checkbox"/>	Potential Risk File Detected	File Security Alert	Medium	May 08 2017 08:17:45 AM	1	Reset
<input type="checkbox"/>	Malicious File Detected	File Security Alert	High			
<input type="checkbox"/>	Approval Request Alert	Approval Request Alert	Low			
<input type="checkbox"/>	File Propagation Alert	File Activity Alert	Medium			
<input type="checkbox"/>	Block Propagation Alert	File Activity Alert	Low			

In addition, triggered alerts appear on the Alerts portlet, which is on the default Home Page, and a count of triggered alerts appears in the console banner.

Name	Type	Priority	Date Triggered	Instances	
System Health_OER Alert	System Health Alert	High	May 08 2017 09:03:11 AM	2	Reset
Potential Risk File Detected	File Security Alert	Medium	May 08 2017 08:17:45 AM	1	Reset
Updater Modified Alert	System Alert	Low	May 08 2017 08:28:46 AM	2	Reset

Mail Notification for Triggered Alerts

When an alert is triggered, notification mail is sent to each subscriber to that particular alert and, if configured, to the global alert subscriber.

While the console shows one banner per triggered alert *class*, the App Control Server sends alert email for every *instance*. Instances are defined as distinct cases that match the alert conditions. In the case of malicious files, for example, if the same malicious file shows up 20 times before you reset the alert, it only counts as one instance. But if 20 *different* malicious files appear before the alert notification is reset, each one counts as an instance and each one generates a new email message to alert subscribers.

Mail notifications contain basic information about the alert such as the time of this action for this instance alert, the system(s) on which an action took place, the logged in user, and the file hash. The File Propagation Alert mail shown below is typical of file-related alerts – the exact information provided varies by alert type.

File Activity Alert

File Propagation Alert

Priority: **Medium**

The file 'test.exe' with hash [a789c3f1318cc2491e4319543622e9e8463988095f5e6624f8b195ee8e93e0ae](#) propagated on 2 of 10 computers with policy Standard Protection within 1 hour. Computer(s) and user(s) affected: ORG123\LAPTOP-3 (rjones) and ORG123\DESKTOP-5 (dgoomez)

Cb Protection Server:	cbpserver.org123.local
Triggered On:	May 29 2017 9:54AM
Created By:	System
Message:	Propagation of a new unapproved file
File State:	Unapproved

Tell me more: [Alert Details](#)

[File Details](#)

[Event Details](#)

As the example above shows, mail notifications also include links to console pages that display information relevant to the alert. This includes Alert Details (the list of instances for this alert), and in this case, the File Details page for the triggering file, and Event Details related to the file (hash) that is the subject of the alert. File and Event Details are not included for non-file alerts. There also may be a Manage Computers link to the Computers table for events that involve App Control settings such as the policy for the computer.

Each email generated by a new instance of the same alert class is tracked in the same Alert History and has a link to a list of instances of that alert. When you reset an alert, the instance history is cleared, but a record of when it was first triggered during this session remains. See [“Viewing Alert Instances and History”](#) on page 623 for an example of the history and instance list for one triggered alert.

Note

The details provided in an alert notification email describe a particular *instance* of the alert. When you click the Alert Details link in email, it opens the Alert Instances page, which shows the details for *all* instances of the triggered alert.

Reminder Mail for Triggered Alerts

If you enable Reminder Mail for an alert, a new email notification of that alert is generated on a schedule you specify as long as the alert has not been reset (manually or automatically). For example, if a Carbon Black File Reputation Unavailable Alert is triggered, email is sent immediately. If Reminder Mail is not enabled, no subsequent email will be sent about this alert unless it is reset and then the condition reappears.

If Reminder Mail *is* enabled, and is set for 30 minutes, subscribers to this alert receive a new email about it every 30 minutes until connectivity is restored or the alert is reset.

Manual and Automatic Alert Resets

Resetting an alert means taking it out of the "triggered" state and clearing the history of all the current instances that caused it to be triggered in the first place. When an alert is reset, it no longer appears on the Triggered Alerts portlet or as a highlighted item on the Alerts page. If the conditions that match the alert return, a new alert will be triggered, new email will be sent to subscribers, and the alert will appear in the usual places in the console.

An alert may be reset manually or automatically:


- **Manual reset** - You manually reset an alert by clicking its Reset button on the Triggered Alerts portlet, the Alerts page, or the Alert History page. In addition to resetting the alert, this adds a "Reset" event to the alert history, with a time stamp and the account name of the console user doing the reset.
- **Automatic reset due to a time limit** – If Auto Reset is enabled for an alert, a time period can be set for an automatic reset. If the alert has not be reset manually or because of change in conditions by the time this time period expires, it will be automatically reset. The default value is 4 weeks. If you want to allow automatic resets for changes in alert conditions but do not want an alert to auto reset based on time, you can use a very large number of weeks as the value in this field. A time-based automatic reset adds an "Auto-Reset" event to its history, with a time stamp. Alert email is not sent for automatic resets.
- **Automatic reset due to changed conditions** – If Auto Reset is enabled for an alert, changes in the conditions that triggered the alert may automatically reset the alert. If the conditions that trigger an alert instance no longer exist, that instance is removed from the list of triggered instances for the alert class it is in. If *no* triggered instances currently exist for an alert class, the alert notification is reset automatically. The conditions that trigger resets differ from one alert type to another, and some types do not auto reset in this way (although they still can auto reset by time period). An automatic reset of an alert adds an "Auto-Reset" event to its history, with a time stamp and user making the change listed. Alert email is not sent for automatic resets.

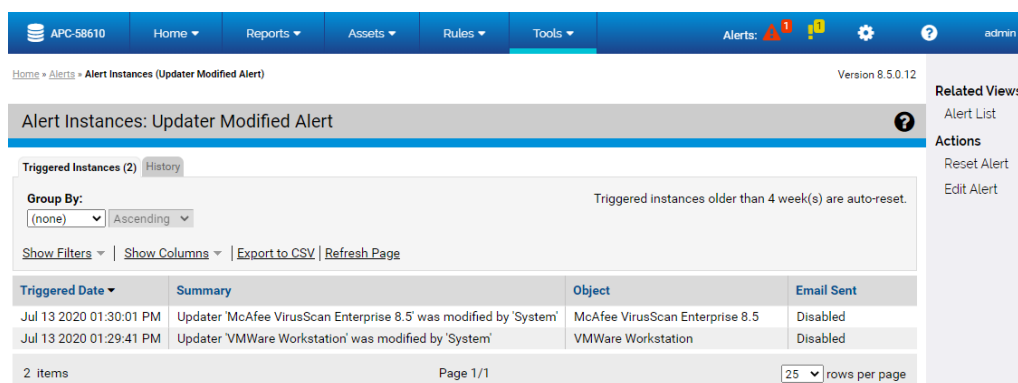
Table 90: Reset Conditions for Different Alert Types

Alert Type	Reset Condition
Backup Missed Alert	Resets when backup is successful
Database Limit Reached	Resets when database size falls below the threshold
Database Verification Failed	Resets when database verification succeeds
Potential Risk or Malicious File Detected	Resets when <i>none</i> of the files that triggered the alert (or would have if they had been detected first) are present
Carbon Black File Reputation Unavailable Alert	Resets when your App Control Server reconnects to Carbon Black File Reputation and synchronization of CDC data with the server is operating properly; generates an event.
Local Approval Alert	Resets when no machines are in Local Approval mode
File Prevalence	Resets if the prevalence of the specified file falls below the specified threshold
Baseline Drift	Resets when the drift in the report falls below the threshold for the specified parameter (user, computer, or policy)
Computer Security	Resets when the conditions leading to it are no longer met (if this change is detectable)
Approval Request Alert	Resets if enough approval requests are Closed that the total number in New or Open state goes below the triggering threshold
Justification Alert	Resets if enough justifications are Closed that the total in New or Open state goes below the triggering threshold
File Propagation and Block Propagation Alerts	No conditional reset because they are time-based alerts. For example, if an alert determined that a particular file propagated to 20 percent of your machines in a one hour period, no future event can change what happened during the one hour period in the past, so the alert remains triggered. Automatic reset by Auto Reset time period only
Updater Modified Alert	No conditional reset; once an updater is modified it remains modified. Automatic reset by Auto Reset time period only

Alert Type	Reset Condition
Rapid Config Alert	No conditional reset; once a Rapid Config is modified it remains modified. Automatic reset by Auto Reset time period only
New Certificate Alert	No conditional reset. Automatic reset by Auto Reset time period only
Revoked Certificate Alert	No conditional reset. Automatic reset by Auto Reset time period only
Event Alert	No conditional reset. Automatic reset by Auto Reset time period only
System Health OER Alert	Resets when no OER indicators on the System Health page show an issue

Viewing Alert Instances and History


If an alert is currently triggered, you can view the instances that triggered it by clicking its View Instances button () on the Alerts page. The Alert Instances page shows the date, summary explanation, the object of the alert (what it was and the action taken, such as created, modified, deleted, etc.), and whether email was sent for each instance.



The screenshot shows the 'Alert Instances: Updater Modified Alert' page. It features a navigation bar at the top with 'APC-58610', 'Home', 'Reports', 'Assets', 'Rules', and 'Tools'. The main content area includes a 'Triggered Instances (2)' section with a 'History' tab. Below this is a table with columns for 'Triggered Date', 'Summary', 'Object', and 'Email Sent'. The table contains two rows of data. At the bottom, there are pagination controls showing '2 items', 'Page 1/1', and '25 rows per page'.

Triggered Date	Summary	Object	Email Sent
Jul 13 2020 01:30:01 PM	Updater 'McAfee VirusScan Enterprise 8.5' was modified by 'System'	McAfee VirusScan Enterprise 8.5	Disabled
Jul 13 2020 01:29:41 PM	Updater 'VMWare Workstation' was modified by 'System'	VMWare Workstation	Disabled

When an alert is reset, details of the instances that triggered it are deleted.

You can view the history of any alert by clicking its View Instances button () on the Alerts page and clicking on the **History** tab, or clicking that tab from the Alert Instances page. If the alert is not triggered, the History tab is the only tab available when you click on the View Instances button.

The Alert Instances History page includes information about when the alert was created and modified (and by whom), when it was triggered and reset, subscriber additions, and if it was enabled or disabled.

Alert History: Updater Modified Alert

Triggered Instances (2) | History

Group By: (none) Ascending History entries older than 4 week(s) might get deleted.

Show Filters | Show Columns | Export to CSV | Refresh Page

Change Date	Changed By	Change
Jul 13 2020 01:29:41 PM	System	Triggered
Jul 13 2020 06:32:37 AM	System	Created

Alert List

Actions

- Clear History
- Reset Alert
- Edit Alert

Both the Alert Instances and the Alert History views have menus for taking actions on alerts, and most of the commands on these menus appear for both tabs:

- **Alert List** – Returns to the Alerts table page.
- **Clear History** – (History page only) Clears all of the history for this alert.
- **Reset Alert** – (Triggered alerts only) Resets the alert from its triggered state and deletes the current instances.
- **Delete Alert** – (Only if alert can be deleted) Deletes the alert itself from the list of available alerts.
- **Edit Alert** – Opens the Edit Alert page so you can modify the configuration of the alert.

Important

Reset Alert eliminates the detailed history of *instances* between the most recent triggering of the alert and the last time you reset it, but leaves all other information in place, including the date and time that the alert was triggered. **Clear History** deletes *all* of the alert's history, including information about its creation, modification, subscribers, and all triggering and reset events. Be sure you do not need this information before clearing the alert history.

Managing Alert Email Subscriptions

There are two types of subscriptions for alerts email:

- **For specific alerts** – On the Alert Information page, you can add subscribers to the email notifications specific to that alert.
- **For all alerts** – On the System Configuration page, you can set up *one* global subscriber for alerts email. See [“Specifying a Global Alert Subscriber”](#) on page 758.

Important

Subscribers receive alert email only if alerts email is properly configured and enabled on the System Configuration page. See [“Configuring Alert and Approval Request Mail”](#) in the “System Configuration” chapter for more information.

Subscription to individual alerts is the normal means of setting up email notification. This allows you to decide which alerts are of interest to a particular user and avoid burying them in other alert email. Users can always watch the Triggered Alerts portlet or the Alerts page for alerts not critical enough to require email notification.

To add a subscriber to the email notification list for one alert:

1. On the Alerts page, click the View Details button next to the alert you want to modify.
2. On the Alert Information page, scroll down to the Subscribers panel, click in the **Email Address** text box, and paste or type the subscriber name.
3. Choose the email type (**Auto**, **Text**, or **HTML**) from the dropdown menu. The default is Auto, which allows the server to determine the best format for the recipient based on information about the recipient's email system.
4. Click **Add** to add the subscriber. The new subscriber name appears in the list below the subscriber entry line.
5. Add any other subscribers you want to receive notifications when this alert is triggered.
6. Click **Save** at the bottom of the Alert Information page. The new subscribers are added to the distribution list for this alert.

You can edit the email address or delivery format of existing subscribers by opening the Alert Information page as you did to add the subscriber and then clicking **Edit** next to the subscriber name. When you have finished editing the subscriber information, click **Update** next to the name, and then click **Save** at the bottom of the Alert Information page. *Be sure to click both buttons.*

You can delete a subscriber from the email notification list for an alert by opening the Alert Information page and clicking **Remove** next to the name. Note that there is no confirmation for this action – the name is removed immediately.

Detecting Agent Issues with Computer Security Alerts

Although many alerts are related to computer security, there is one built-in alert that is specifically designed for this purpose. The Computer Security Alert, which is disabled by default, is triggered by events that may indicate suspicious behavior.

The screenshot shows the 'Edit Alert' configuration page for 'Computer Security Alert'. The page is divided into three sections: General, Type, and Criteria.

General

- Alert Name:** Computer Security Alert
- Message:** Suspicious behavior detected
- Priority:** Medium
- Status:** Enabled Disabled

Type

- Type:** Security Alert
- Description:** Alerts subscribers when a suspicious behavior is detected
- Mail Template:** Default

Criteria

Alert When:

- Computer not protected
- Agent tampering detected
- Agent tampering prevented
- Computer clock out of sync

Criteria Triggering a Security Alert

There are four triggering criteria that can be enabled in the Computer Security Alert - by default, all are enabled when you enable the alert itself. Which one of these criteria triggers a security alert is identified in Summary field on the Alert Instance page, and in the email notification (if enabled) sent due to the alert.

The criteria for triggering a Security Alert are:

- **Computer not protected** – This condition occurs if an agent upgrade fails. It means that the agent is not running on the identified computer, and so the computer is not protected by App Control (the Connection status indicator for this computer on the Computers page will be red). Restoring the agent to proper operation automatically resets the alert when this is the triggering condition.
- **Agent tampering detected** – If agent tamper protection is accidentally disabled through the console and a user on a computer running the agent modifies the agent folder (for example, by adding a new file), the Computer Security Alert is triggered with the summary description "Agent tampering detected". As soon as an administrator re-enables the tamper protection for the agent, this alert is automatically reset.
- **Agent tampering prevented** – If a user on an agent-managed computer attempts to tamper with the agent and fails, the Computer Security Alert is triggered with the summary description "Agent tampering prevented". An example of this might be a user attempting to copy files to the agent folder (Bit9\Parity Agent) but failing because of tamper protection. Another example might be unauthorized attempts to run special agent management commands (i.e., without a correct password). When this condition triggers the alert, the alert must be reset manually.
- **Computer clock out of sync** – One way to attempt to run malware or other unauthorized files without detection is to change the clock on the targeted system to create an invalid timestamp. The agent still detects and reports a file execution under these circumstances, but generates a Computer Security Alert with the summary description "Computer clock out of sync" as soon as the discrepancy between the App Control Server clock and the agent clock is detected. Correcting the system time on the computer that is the source of the unauthorized activity will allow this alert to be reset by the next event received by the App Control Server.

When a Computer Security Alert is enabled, *any* of the enabled criteria on any computer will trigger it. While the alert is triggered, additional cases of the triggering condition on the same computer are recorded in the history, but do not create another alert instance. If the same computer reports an event that meets a *different* triggering condition, however, another instance is displayed. For example, two failed attempts at tampering do not create two alert instances unless the alert is reset between them. However, an attempt to tamper followed by a clock out of sync on the same computer does create two different alert instances.

As with all alerts, each instance results in an email notification, if notification is enabled and properly configured. Both the Alert Instance displayed in the console and the email

notification of the alert contain the security event description, the name of the computer on which it happened, and the time of triggered instance.

Note

Because the Computer Security Alert is based on agent events, a disconnected agent will not produce an alert when the triggering conditions are met. In addition, in environments with a large number of agents, files and changes, this alert might be delayed if a large number of events is being processed by the App Control Server when the agent reports the security events.

Alerts for File Prevalence

On the File Catalog tab of the Files page, there is a *Prevalence* column that shows you how many computers a file is on (based on periodic updates).

The screenshot shows the 'Files' page with the 'File Catalog' tab selected. The table displays the following data:

Select	First Seen Date	First Seen Name	Product Name	Prevalence	Trust
<input type="checkbox"/>	Jul 16 2012 01:26:20 PM	winsta.dll	Microsoft® Windows® Operating System	124	10
<input type="checkbox"/>	Jul 16 2012 01:26:20 PM	ieframe.dll	Windows® Internet Explorer	17	10
<input type="checkbox"/>	Jul 16 2012 01:26:20 PM	msctfmonitor.dll	Microsoft® Windows® Operating System	126	10
<input type="checkbox"/>	Jul 16 2012 01:26:20 PM	wintrust.dll	Microsoft® Windows® Operating System	99	10
<input type="checkbox"/>	Jul 16 2012 01:26:20 PM	imagehlp.dll	Microsoft® Windows® Operating System	99	10
<input type="checkbox"/>	Jul 16 2012 01:26:20 PM	ielowutil.exe	Windows® Internet Explorer	92	10
<input type="checkbox"/>	Jul 16 2012 01:26:20 PM	vmtoolsd.exe	VMware Tools	0	10
<input type="checkbox"/>	Jul 16 2012 01:26:21 PM	vmwaretray.exe	VMware Tools	0	10
<input type="checkbox"/>	Jul 16 2012 01:26:21 PM	php-cgi.exe	PHP	1	4
<input type="checkbox"/>	Jul 16 2012 01:26:21 PM	cmd.exe	Microsoft® Windows® Operating System	126	10

When Prevalence is listed in a table, you can sort the table by prevalence or set Filters on the page to show a report of only those files with a prevalence greater than or equal to a number you specify. If a file was seen by an agent and reported to this App Control Server at one time but now has a prevalence of zero, it is removed from the table, although you can view it by choosing **Removed Files** from the Saved Views on the Files page.

Prevalence Alerts

Prevalence alerts are triggered when the prevalence of a particular file reaches a threshold you set. You can go to the Alerts page and type in information about the file you want to create an alert for, but the easiest way to create a prevalence alert is from the File

Details page of the file you want to track. See [“Using App Control Alerts”](#) on page 608 for more information about alerts.

Notes

- You cannot use wildcards in the filename for a prevalence alert.
- Provide a name, not a path, for prevalence alerts.

To create a prevalence alert for a file from its File Details page:

1. On the Files page, click on the View Details button next to the name of the file whose propagation you want to track. The File Details page opens.

The screenshot displays the 'File Instance Details' page. The main content area shows the following details for the file 'chrome.exe':

- File Name:** chrome.exe
- Date Created:** Jul 13 2020 07:22:26 AM
- File Path:** c:\program files (x86)\google\chrome\application\
- Computer:** WORKGROUP\APC-58609
- Platform:** Windows
- User Name:** (none)
- Local State:** Approved
- Local State Details:** Locally Approved
- Detached Publisher:** (none)
- Executed:** Yes
- Present At Initialization:** Yes
- Top-Level File:** No
- Deleted:** No
- Root File Name:** (none)

The 'General' section below shows:

- First Seen Name:** chrome.exe
- First Seen Date:** Jul 13 2020 07:25:12 AM
- Last Updated:** Jul 13 2020 07:25:12 AM
- First Seen Path:** c:\program files (x86)\google\chrome\application\
- First Seen Computer:** WORKGROUP\APC-58609
- First Seen Platform:** Windows
- Extension:** exe
- Global State:** Unapproved

The right-hand sidebar contains an 'Actions' menu with the following options:

- All File Instances
- File Events
- Computers with this file
- Computers without this file
- Carbon Black EDR Details
- Actions**
- Remove Local Approval
- Approve Globally
- Ban Globally
- Approve by Policy...
- Ban by Policy...
- Delete All File Instances from WORKGROUP\APC-58609
- Delete File from All Computers
- Add Meter
- Add Alert** (highlighted with a red box and mouse cursor)
- Advanced

2. On the File Details page for that file, click **Add Alert** in the Actions menu. The Alert Information page opens with the name of the file and its hash automatically filled in.

Add Alert

General

Alert Name:

Message:

Priority:

Status: Enabled Disabled

Type

Type:

Description: Alerts subscribers when a specified file is present on more than specified number of computers

Mail Template:

Criteria

Specify File By: File name Hash

Hash Type:

Hash Value:

Threshold:

Subscribers

Note: Alert must be created before email recipients can be specified

Reminder Mail

Status: Enabled Disabled

Remind Every:

Auto Reset

Status: Enabled Disabled

Reset After:

Create & Exit
 Create
 Cancel

3. Set the remaining parameters you want for this alert including:
 - a. Threshold number of computers on which this file must appear to trigger the alert.
 - b. Reminder mail specifications if you want periodic email reminders to be resent after a certain period of time if the alert is not reset or the condition not remedied.
4. Click **Create**, if you want to stay on this page or **Create & Exit** to go to the Alerts table page. You now have a prevalence alert for this file, visible on the Alerts page.
5. To add email alert subscribers, click the View Details button for the alert and add the addresses in the Subscribers section of the Alert Information page.

Monitoring Specific File Executions

Software metering enables you to track the number of times users run specified files. When you create a meter, you specify a file to be tracked. Each time the specified file runs on a computer, the server records its execution. Configurable reports enable you to display cumulative execution events by time of execution, user, computer, and policy. You can create as many meters as you need and centrally manage them (view reports, edit, and delete) in one place. Monitoring begins almost immediately after you create the meter.

Software metering is useful for the following purposes:

- Gathering data about how often applications are used
- Determining which computers are running an application
- Locating computers running older versions of software for upgrade or completely retiring obsolete applications

Notes

- App Control Agent is one of the first processes to start when you start your computer. It is normally configured so that a user cannot log in to an agent-managed computer until the agent has started up, or a specified timeout period expires. However, if a service or process is configured to start before the agent, its activity is not monitored or controlled until the agent starts.
- You can locate *all* executed files on your network, or on a subset of your computers, using Filters on the Find Files page or the Files on Computers tab on the Files page. See [Defining a Search on the Find Files Page](#).

You can create a meter from scratch, as shown in the procedure immediately below, or you can create a meter for a file directly from its File Details page – see [“Creating a Meter from the File Details Page”](#) on page 633.

To meter execution of specified file(s):

1. On the console menu, choose **Tools > Meters**. The Meters page appears:

2. On the Meters page, click **Add Meter**.

3. On the Add Software Meter page, select the type of identification (file name or hash) you want to use for this file. Additional fields appropriate for the selected type appear.
4. In the Software Meter panel, specify information about the file to be monitored.

Table 91: Software Meter Parameters

Field	Description
Meter name	Text description of the software to be metered.
Type	<p>To meter a file you must know the name of the file or its hash (data signature). Choose either one, as appropriate. Note that File Name meters are platform-specific; hash meters apply to all platforms.</p> <p>A meter created directly from a File Details page automatically has that file's SHA-256 Hash (if available) entered as the file identifier.</p>
Platform	<p>For file name meters, the platform (Windows, Mac, or Linux) for which the meter is in effect. File name meters can be used for one platform only.</p> <p>(Field does not appear for hash meters.)</p>
File Name	<p>File name (or path) to which this meter applies. If you provide just a file name, execution of that file in any location is metered. If you provide a path that ends in a file name, only executions of the file in the specified location are metered.</p> <p>If the path you enter ends with a directory, the meter counts all executions in that directory and all of its subdirectories.</p> <p>If you create a meter for a file name that does not currently exist, a message appears noting that. You can still create the meter in advance of the file appearing on one of your systems.</p> <p>Platform Notes:</p> <ul style="list-style-type: none"> • For Windows paths, you can specify a local drive name (for example, C:\dir\subdir\application) or a UNC path (for example, \\dir\subdir\application). You cannot specify mapped drives (for example, Z:\application) for network access. • For all paths, you must use the correct directory delimiters for the platform you choose. • You can switch platforms after a meter is created, but keep in mind platform differences, such as directory delimiters and drive letters, that might make a path invalid on a different platform.
Hash Type	<p>Cipher algorithm used to create the hash you want to monitor (MD5, SHA-1, or SHA-256). Note that App Control returns SHA-256 hashes by default for Files or Find Files searches, but cross-references it so you can monitor, approve or ban by the other hash types. If you create a meter directly from a File Details page, that file's SHA-256 Hash (if available) is used as the file identifier.</p>

Field	Description
Hash Value	Hash (data signature) for the file. Monitors file execution on computers even if the hash has not been previously identified. If you enter a hash from an external source, its execution is reported when first encountered on agents. To locate hashes on your network, use the Files page or Find Files utilities. Note that you can create a meter directly from the File Details page for any file identified on the App Control Server.
Description	Optional text that further describes the metered file. To display this information, add the Description column to the Meters table.

For example, a meter to monitor executions of Microsoft Excel by its name might be specified as shown in the screen below:

- To add the file to the table of metered files, click **Save**. The meter is created and activated, and the name of the meter, the metered file, and execution information appears in the Meters table on the Software Meters page:

Name	File or Hash	Executions	Computers	Users	Date Created	Created By
Microsoft Excel	excel.exe	45	10	8	Jul 14 2020 10:12:31 AM	admin

- To change meter information, click the View Details button next to the meter name.
- To display a report of meter events, click the View Report button to the far left of the report name.

Note

By default, meter events are grouped by computer. To view all executions of files on that computer, expand the computer name. Alternatively, you can eliminate the grouping by choosing **None** on the *Group by* menu.

Report Parameters ?

Basic

Meter Name: Microsoft Excel
File Name: excel.exe

Time Range

Ever
 In the past...
 During range

Past:

Meter Report Details ?

Group By:

[Show Filters](#) |
 [Show Columns](#) |
 [Export to CSV](#) |
 [Refresh Page](#)

Timestamp	Filename	User	Computer	Policy
Computer: 3 items				
May 29 2017 07:30:24 AM	excel.exe	MYCORPIrjones	MYCORPILaptop-12	Standard Protection
May 29 2017 07:15:36 AM	excel.exe	MYCORPIgomez	MYCORPIDesktop-3	Standard Protection
May 29 2017 06:59:11 AM	excel.exe	MYCORPIsmith	MYCORPIDesktop-6	Standard Protection
Computer: 2 items				
Computer: 6 items				

11 items in 3 groups Page 1/1

rows per page

- To delete a meter, click the Delete (x) icon next to its name on the Meters page.

Creating a Meter from the File Details Page

If you know you want to monitor executions of one specific file, you can create a meter directly from its File Details page. This has the advantage of pre-configuring most of the information required for the meter, including the hash value – meters created in this way are automatically Hash type meters.

To create a software execution meter from a File Details page:

- Open the File Details page for the file you want to meter.
- In the Actions menu to the right of the File Details page, click **Add Meter**. The Add Software Meter page appears, with the Hash value of the file already entered and the file name as the default meter name.
- If you choose, change the Meter name and add a description.
- Click **Save** to save and enable the meter.

Chapter 22

Monitoring Change: Baseline Drift Reports

This chapter explains how to use Baseline Drift Reports, which allow you to track changes in the inventory of files on systems running the App Control Agent. [Chapter 21, “Events, Alerts and Meters,”](#) describes other monitoring features.

Sections

Topic	Page
Baseline Drift Overview	635
Viewing and Managing Baseline Drift Reports	637
Responding to Drift Report Results	643
Creating and Editing Reports	646
Drift in Multi-Platform Environments	652
Managing Snapshots	653
Displaying Baseline Drift Reports in Graphs	656
Creating Baseline Drift Alerts	657

Baseline Drift Overview

App Control's Live Inventory of files on computers reporting to your App Control Server gives you the ability to measure baseline drift, the difference between a baseline of files and the current files on a target you specify. This difference is available as a baseline drift report that you can view either in detail in dynamic tables or as graphic charts on a App Control dashboard. Baseline drift reports provide not only simple numbers of file differences but also risk analyses related to those changes.



Once it is set up, a drift report runs automatically every few hours, giving you an up-to date record of changes in your file inventory. You can create different baseline drift reports for different targets and baselines, and App Control includes some reports pre-configured for your use. By default, only Power Users and Administrators can create, modify and delete reports. However, custom account groups can be configured to allow viewing only or viewing and management of drift reports and snapshots.

Table 92: Baseline Drift Terminology

Term	Description
Target	A collection of current files that you want to analyze. This might be all the files on a particular computer, on computers with a particular security policy, or on all computers. It also can be a custom filtered table of files from one or more computers.
Baseline	The reference against which you compare the target. It can be a set of files captured as a "snapshot," multiple snapshots, a set of one or more computers, or a custom baseline generated by filters and other parameters you define. You also can have no baseline, in which case a report shows you new files appearing over time.
Snapshot	A set of files collected from one or more computers. It can be <i>all</i> files from the selected computer(s), files selected based on custom-defined filter, or file lists captured from other pages in the console. Each snapshot is named, and can be used as the baseline for a drift report.
Baseline drift report	A report that contains information about the differences between a baseline and a target. A drift report can show differences simply in the number of changed files as well as the risk indicated by those changes.

How Drift and Risk are Measured

For the designated target, baseline drift reports can provide several different types of data about the computers or files in the report. [Table 93, “Basic Drift Values”](#) describes this information.

Table 93: Basic Drift Values

Term	Description
Drift	The amount of drift measured simply in terms of files added, changed, and (if configured for a report) deleted in the target. Files are identified by their hash value. An added file, a changed file, or a modified file each have a drift value of 1. See “Advanced Baseline Drift Report Options” on page 649 for more on how App Control determines whether a file has been modified.
Weighted drift	A calculation based on the drift value and adjusted by several factors that might increase or decrease the significance of the drift for each file. Among the adjustment factors are trust level, threat level, file type and associations with other files. For example, the weighted drift for files that have valid digital signatures, have high trust, or were installed by files with high trust will be reduced from what it would be without these factors.
Risk	A calculation similar to weighted drift, but adjusted so that files believed to pose no threat show a risk of zero.
% Weighted drift	The percentage of total weighted drift in the current report contributed by the item in a row.
% Risk	The percentage of total risk in the current report contributed by the item in a row.

Other key factors in determining the total drift and risk reported in a baseline drift report are:

- **File Filtering:** You can decide which files in the baseline and in the target participate in the comparison. For example, the pre-configured drift reports compare Unapproved files, but ignore Banned or Approved files – you can change this if you choose. There are several other file categories you can include or exclude from the comparison. See the [“Using Filters in Target and Baseline Definitions”](#) and [“Advanced Options: File Filter Options”](#) sections below for more detail.
- **File Comparison Method:** By default, if a file hash found in the baseline is also found *anywhere* in the target, it is considered a matching file, and no drift is reported. This is called the *File Content* method. The alternative is the *File Location* method, in which the same hash in different locations in the baseline and the target is considered a drift. See [“Advanced Options: File Comparison Method”](#) for more detail.

Viewing and Managing Baseline Drift Reports

All baseline drift reports appear on the Manage Baseline Drift Reports page. Two pre-configured baseline drift reports appear in the console: Drift of all computers, and Daily drift of all computers. These are disabled by default. These pre-configured reports provide a useful way to view the configuration options for baseline drift and view their results in a report. You can copy any existing report and use it as a starting point for new reports.





To view the table of Baseline Drift Reports:

- On the console menu, choose **Reports > Baseline Drift**.
The Manage Baseline Drift Reports page appears.

	Name	Date Created	Created By	Date Last Completed	Status
<input type="checkbox"/>	Drift of all computers	Jul 13 2020 06:32:35 AM	System	Jul 14 2020 10:21:59 AM	Available
<input type="checkbox"/>	Daily drift of all computers	Jul 13 2020 06:32:35 AM	System	Jul 14 2020 10:23:00 AM	Available

The Manage Baseline Drift Reports page gives you access to the existing reports as well as the ability to create a new report. On the Manage Baseline Drift Reports page, you can use any of the standard buttons and tools available on a console table page, including filtering, adding or removing columns, and grouping the items in the table. The following table describes the buttons, columns, and tabs on the drift page.

Table 94: Manage Baseline Drift Reports Page parameters

Item	Description
Reports and Snapshots tabs	The Reports tab (default) shows the table of all available drift reports and key information about them. It also provides an Add Report button for creating new reports. The Snapshots tab shows the table of all available snapshots and key information about them. See “Managing Snapshots” for more information.
Add Report button	Opens the Add Baseline Drift Report page, on which you can enter the details for a new Baseline Drift Report.
 View Report Results button	Shows the most recent results of the report in its row.
 View Details button	Opens the Baseline Drift Report Details page for the report in its row. You can view and edit the report details on this page.
 Schedule Run button	Schedules the report in its row to be run as soon as possible rather than waiting for the normal report period.
 Delete button	Deletes the report in its row.
Name field	The name of the report. Clicking this name shows the most recent results of the report.
Date Created field	The date and time this report was created.
Created by field	The console user who created this report – reports showing System in the <i>Created by</i> field were built in to App Control.
Date Last Completed field	The date and time the report was last run. If blank, the report is either disabled or is new and has not completed its first run.
Status field	Shows the current status of the report. The possible values are: <ul style="list-style-type: none"> • Available – Updated report is ready and available for viewing • Available (Updating) – New report is currently being generated. Previous report will be available for viewing until current report generation completes. • Disabled – Report is disabled and is not generating results. Last generated results are deleted. • Not available – Report is new; results have not been generated yet.

Viewing Baseline Drift Report Results

If a report listed on the Manage Baseline Drift Reports page shows that it is *Available*, you can view the most recent report results.

To view a baseline drift report:

1. On the console menu, choose **Reports > Baseline Drift**.

- On the Baseline Drift page, click the name of the report you want to see in the Manage Baseline Drift Reports table. By default, the initial view shows drift by computer.

Report Results: Computer View

The figure below shows the initial view of the built-in *Drift of all computers* report. The results show a table of all computers that have had agent-tracked files added or modified in the past 24 hours (deleted files are not tracked by default), and the amount of drift contributed by each computer. Note that the View Mode panel has Computers selected.

Drift of all computers

Saved Views: (none) Add View Mode: This report was generated on May 10 2017 11:59:16 AM. Show results by: Computers Files Group By: (none) Ascending

Show Filters Show Columns Export to CSV Refresh Page

K < 1 2 3 4 5 >>

	Computer	Drift	Risk	Policy
	MYCORPDESKTOP-5	87663	158220.4	Research Team
	MYCORPDESKTOP-10	32567	52143.8	Research Team
	MYCORPLAPTOP-2	21402	22461.7	IT Group
	MYCORPLAPTOP-7	12099	7194.4	Standard Protection
	MYCORPDESKTOP-3	10966	54.6	Standard Protection
	MYCORPDESKTOP-12	10707	16555.5	Standard Protection

Report Results: File Views

Files views of Baseline Drift Reports provide more detail than Computers views since the key elements of drift are based on the files themselves. There are three primary File views available for drift reports:

- All Top-level Files** – This is the main Files View of the drift report you choose. It shows the drift, risk, and other data for each top-level file in the report.
- Files Associated with One Top-Level File** – This is a drift report for the files associated with one top-level file. You can view an associated files report by clicking on a highlighted name in the Top-Level Files report.
- Files on One Computer** – This is a drift report for all the files on one computer that contribute to drift. You can view a computer-specific files report by clicking on the name of a computer in the Computer view.

In addition to the primary views, there are pre-configured **Saved Views** that give you a different perspective on the information in drift-by-files tables:

- Drift Contributing to Risk** – This shows the standard report on drift by (top-level) files, except that files with drift risk of 0 are filtered out.
- Drift by Category** – This view is the equivalent of choosing *Category* in the Group by menu or Filters list. It shows a list of file categories, as reported by Carbon Black File Reputation, in the left column of the table. Clicking on the plus sign next to a category expands the view to include all files in that category and the Drift and Risk levels for each file.
- Drift by Publisher/Company** – This view is the equivalent of choosing Publisher or Company in the Group by menu or Filters list. It shows a list of the identifiable




Publisher/Company names for the files in the left column of the table. Clicking on the plus sign next to a Publisher/Company name expands the view to include all files with that Publisher or Company, and the Drift and Risk levels for each file.

- **Drift by Installed Program** – This view is the equivalent of choosing Installed Program in the Group by menu. It shows total drift of all files associated with an installer program.

Platform Note: This view is useful only for Windows agents.

The table below shows the controls and default fields on the Files view of a drift report.

Table 95: Drift Report Results Elements

Item	Description
 View Report Results button	In Computer View mode, drills down to the Baseline Drift report for the computer in its row.
 View Details button	In Files views, opens the File Instance Details page for the file in its row.
 Find Files button	(In Files views only) Goes to the Find Files page and shows all file instances matching the hash of the file in its row, on all computers.
File Name	Shows the name of a file in the target that is contributing to drift. If the file is highlighted in blue, it is a link, indicating that it is a top-level file with associated files. Clicking on the link drills down to a Baseline Drift report for the files associated with the named top-level file.
Publisher or Company	Shows the publisher (if available) or company (if available and there is no publisher information).
Drift	In Computer View mode, the sum of drift for all drifted files on the computer in this row. In File views, the sum of drift for this file (if it has no associated files) or for files associated with this file (if it is a top-level file). For views with grouped information, the sum of the drift for each instance of the group parameter. Expanding the group shows drift for each member of the group.
Risk	The sum of the risk for all drifted files on the item in this row. See “How Drift and Risk are Measured” on page 636 for more details.
Threat	A threat level for the file in this row based on a weighted analysis of malware threats known to Carbon Black File Reputation. Threat levels are Malicious (red ! icon), Potentially Malicious (yellow ! icon), Unknown (no icon), or Clean (green ✓ icon).
Trust	On a scale of 0-10, the level of trust for the file in this row. Zero is the lowest level of trust and 10 is the highest. Trust is computed from a variety of factors, including file source, publisher, and identification in Carbon Black File Reputation (e.g., is it malware or some other undesirable category of file).
Computer	Shows which computer the file in this row is on. Clicking on the name opens the Computer Details page for that computer.

Item	Description
User Name	User logged into the computer when the installation was started or top-level file was created.
View Mode	Clicking on Files in the View Mode box changes the view from drift by computers to drift by files, and lists the top-level files in the report. Clicking on Computers in the View Mode box changes the view from drift by file to drift by computers, and lists all of the computers in the drift report. Note: Clicking on <i>Show individual files</i> in the lower right of the table causes the Files view to show both top-level files and any files associated with them.
Saved Views	Files View mode has three saved views. To return to a full list of files in the report, choose none on the Saved Views menu.
Action menu	Allows you to take action on checked files in the drift report. See “Responding to Drift Report Results” on page 643 for details.

Drift by Files: Individual or Top-Level

The default view for drift reports that show files is to show all of the individual files in the drift report. This shows both files installed by (or copied from) other files and files not generated by other files in the report.

In some cases, you might find the report for *top-level files* is often the most useful in tracking drift and risk since many of these files are the ones that install other files on computers. These are files not generated by other files in the report.

To display the files view of a baseline drift report:

1. On the console menu, choose **Reports > Baseline Drift**.
2. On the Baseline Drift page, click the name of the report you want to see in the Manage Baseline Drift Reports table.
3. In the View Mode box, click the **Files** radio button.
4. In the bottom right corner of the page, make sure the *Show individual files* box is checked (the default). All drift files are shown.

Drift of all computers

Saved Views: (none) Add View Mode: Show results by: Computers Files This report was generated on Jul 14 2020 10:21:59 AM. Group By: (none) Ascending

Show Filters | Show Columns | Export to CSV | Refresh Page

Action | K < 1 2 3 4 5 >>

	Date Created	File Name	Publisher or Company	Drift	Risk	Threat	Trust
<input type="checkbox"/>	Sep 27 2016 03:20:48 PM	robocopy.exe	Microsoft Corporation	20263	21971.8	<input checked="" type="checkbox"/>	10
<input type="checkbox"/>	May 30 2017 10:34:54 AM	tortoiseproc.exe	Open Source Developer,	12008	7185.3	<input checked="" type="checkbox"/>	10
<input type="checkbox"/>	Dec 17 2016 09:34:35 AM	atlassian-jira-soft...	Atlassian Pty Ltd	9130	16216.2		
<input type="checkbox"/>	Oct 17 2014 07:22:00 PM	svn.exe	http://subversion.apache	3980	7175.8	<input checked="" type="checkbox"/>	4

K < 1 2 3 4 5 >> Show individual files

20793 items Page 1/832 25 rows per page

- If you want the report results to show top-level files only, *uncheck* the *Show individual files* box in the far right bottom corner of the page.

K < 1 2 3 4 5 >> Show individual files

7400 items Page 1/296 25 rows per page

Drift by Files: Associated Files Report

A name highlighted in blue in a drift report indicates that more information is available if you click on the name. On the top-level files report, clicking on a highlighted file name gives results for files *associated with* the file you clicked. Associated files are files that either were installed by the top-level file or are copies of it (i.e., have the same hash).

To view files associated with a top-level file in a drift report:

- Click on the name of the top-level file in the drift report results.

Drift of all computers

Saved Views: (none) Add This report was generated on May 31 2017 01:14:14 PM

Group By: (none) Ascending

Show Filters | Show Columns | Show Snapshot | Export to CSV | Refresh Page

Action | K < 1 2 3 4 5 >>

Files associated with 'windbg.exe' on computer MYCORPLT-3 [Back to report]

	Date Created	File Name	Publisher or Company	Drift	Risk	Threat	Trust
<input type="checkbox"/>	Aug 19 2016 04:52:51 PM	kernel32.dll	Microsoft Windows	1	0.0	<input checked="" type="checkbox"/>	10
<input type="checkbox"/>	Nov 29 2016 09:38:43 PM	cdd.dll	Microsoft Corporation	1	0.0	<input checked="" type="checkbox"/>	8
<input type="checkbox"/>	Aug 19 2016 12:36:14 PM	ntdll.dll	Microsoft Windows	1	0.0	<input checked="" type="checkbox"/>	10
<input type="checkbox"/>	Nov 29 2016 09:36:46 PM	gdi32.dll	Microsoft Windows	1	1.3	<input checked="" type="checkbox"/>	10

To return to the top-level files view from an associated files report:

- In the *Files associated with* line above the table, click **[Back to report]**.

Drift by Files on a Single Computer

You can get a report of drift by files on a single computer. This can be useful in a number of situations; for example, it can help you locate a computer that has significantly more drift than others so that you can take remediation steps.

To display the drift by files for a single computer:

1. On the console menu, choose **Reports > Baseline Drift**.
2. On the Baseline Drift page, click the name of the report you want to see in the Manage Baseline Drift Reports table.
3. If the Computer View mode is not displayed, click the **Computers** button in the View Mode box.
4. Click the View Details button next to the name of the computer for which you want to see a files report. A report showing only the drifted files on that computer appears.

Drift of all computers

Saved Views: (The Current View Has Unsaved Changes - Discard)
 (none) [Add] This report was generated on May 31 2017 01:14:14 PM.
 Drift of computer: MYCORPILT-3 [Back to report]

Show Filters | Show Columns | Show Snapshot | Export to CSV | Refresh Page

Action [K+1] [2]

<input type="checkbox"/>	Date Created	File Name	Publisher or Company	Drift	Risk	Threat	Trust
<input type="checkbox"/>	Dec 17 2016 09:34:35 AM	atlassian-jira-software...	Atlassian Pty Ltd	9130	16216.2		
<input type="checkbox"/>	Mar 23 2016 05:33:11 PM	dllhost.exe	Microsoft Windows	682	11.7	✓	10
<input type="checkbox"/>	Jun 07 2016 02:54:12 AM	vsixautoupdate.exe	Microsoft Corporation	336	0.0	✓	10
<input type="checkbox"/>	Mar 23 2016 05:00:50 PM	68bef.msi	Intel Corporation-Wireles	73	0.0	✓	10
<input type="checkbox"/>	Mar 23 2016 05:00:51 PM	9c33.msi		37	45.5	✓	0
<input type="checkbox"/>	Mar 23 2016 05:00:49 PM	9c52.msi		36	46.8	✓	0

To return to the top-level Computer view from computer drift details view:

- In the *Drift of computer* line above the table, click **[Back to report]**.

Responding to Drift Report Results

You can use the results of a Baseline Drift Report for a wide variety of purposes, ranging from simply noting the level of drift to changing the security policy for some or all of your computers. Most of the actions you take can be done in the console, although some of them must be done manually, most notably, restoring missing files. In general, you check the checkbox next to files you want to act on. Many of the choices for responding are on the Action menu.

You can remediate drift in following ways:

- **Add Files to Snapshot:** If the baseline drift report was based on one or more snapshots, you can click the **Show Snapshot** link and add all files or just selected files in the report to a snapshot. The files you add are immediately removed from the report and will not become part of subsequent reports. Note that when a file group is checked, all files in the group are added to the snapshot you choose.
- **Locally Approve Files:** Using the Action menu, you can choose **Approve Locally** for checked files in a drift report. In addition to allowing the file to execute on the computer it was found on, this excludes the file from future drift reports if the report excluded all approved files (the default).
- **Remove Local Approval:** Using the Action menu, you can **Remove Local Approval** on checked locally approved files in a drift report.
- **Globally Approve or Ban Files:** Using the Action menu, you can Globally Approve or Globally Ban checked files in a drift report.
- **Create Custom Approvals or Bans:** Using the Action menu, you can choose Approve by Policy or Ban by Policy to create custom approvals or bans for checked files in a drift report. For approvals, you can approve by policy and/or choose to Mark the checked files as installers. For bans, you can ban by policy and choose to block files banned or just report that they would have been blocked if the ban had been fully enforced.
- **View and Act on Members of a File Group:** If you want to see the details of a file group, you can click on the file name or the View Details button, which shows a page with files in the group that contribute to drift. Here, you can approve or ban files on an individual basis.
- As on other pages in the console, from a drift report you can drill down to the File Details page for access to many of the actions described above.
- **Approve or Ban Files by Group or Trust Methods:** Rather than approving or banning individual files, you can approve the root package that installs a group of files. You might also want to approve files by Publisher, Updater, or User (via the Software Rules page) if you notice that a large number of files from the same source appear in your drift reports and you are willing to trust that source. While making this kind of change will not affect the current report, it will make sure the files covered by the change do not appear in future generations of the report (or other, similar reports) as long as you are not including approved files in the report.
- **Add or Remove Files:** Outside of the console, you can add or remove files from one or more of your systems based on the information in the drift report, reducing the drift shown in future reports.

Adding Drift Results to a Snapshot

When you view Baseline Drift Report Results, you might see files in the report that you do not want to track for drift. If the drift report uses one or more snapshots as a baseline, you can add files from the drift report to one of the baseline snapshots. You also can create a new snapshot and then add the new snapshot to the baseline.

This type of remediation essentially means you want to ignore certain drift results *in the future*. Nothing is sent to the agents to remove this drift (i.e. change their file inventory),

and existing report results remain the same. However, files you add to a snapshot that is part of the baseline will not be included in future drift report results.

Drift of all computers

Saved Views: (The Current View Has Unsaved Changes - Discard) View Mode
 (none) Add Show results by: Computers Files

Group By:
 (none) Ascending

Show Filters | Show Columns | Hide Snapshot | Export to CSV | Refresh Page

Add Files to Snapshot

Files to add: All files Checked files

Choose existing snapshot: (none) Add

Create new snapshot: Create

Action | < 1 2 3 4 5 >

<input type="checkbox"/>	Date Created	File Name	Publisher or Company	Drift	Risk	Threat	Trust	Global State
<input type="checkbox"/>	Oct 28 2015 12:12:48 PM	ntoskrnl.exe	Microsoft Corporation	3964	796.6	<input checked="" type="checkbox"/>	10	Approved
<input type="checkbox"/>	May 24 2014 08:20:07 PM	1.7.29-setup-x8...		3343	887.6	<input checked="" type="checkbox"/>	1	Unapproved
<input type="checkbox"/>	May 31 2017 11:48:30 AM	expand.exe		3314	1070.0	<input checked="" type="checkbox"/>	10	Approved
<input checked="" type="checkbox"/>	Nov 28 2016 03:46:31 PM	7zg.exe	Igor Pavlov	2383	10.4	<input checked="" type="checkbox"/>	10	Approved
<input type="checkbox"/>	Dec 05 2016 11:09:00 AM	svn.exe	Open Source Developer,	2267	2523.0	<input checked="" type="checkbox"/>	10	Approved

To add files to a snapshot from a baseline drift report:

1. In the report, unless you plan to add all of the files to a snapshot, check the checkboxes for any files you want to add.
2. Click the **Show Snapshot** link to display the *Add Files to Snapshot* panel.
3. In the *Add Files to Snapshot* panel, choose the radio button for **All files** or **Checked files** in the *Files to add* line.
4. Specify the snapshot to which you want to add the files:
 - a. If you want to add the files to an existing snapshot, pick one from the *Choose existing snapshot* menu and click **Add**. Note that this menu includes all available snapshots, not just those used as a baseline for the current report.
 - b. If you want to add the files to a new snapshot, type a name in the *Create new snapshot* box and then click the **Create** button.
5. If the report is more than one page long and you are adding checked files, repeat the procedure for each page containing files you want to add.

Note

The procedures above assume you are adding to a snapshot to affect future results when the *current* drift report is run, but there are no restrictions on how you use the snapshot. You may save files to a snapshot for another purpose.

Creating and Editing Reports

The Add Baseline Drift Report and Edit Baseline Drift Report pages share most of the same parameters. Here, we describe creating a report. The procedure for editing the report is essentially the same, except that you start with an existing report.

To create a Baseline Drift Report:

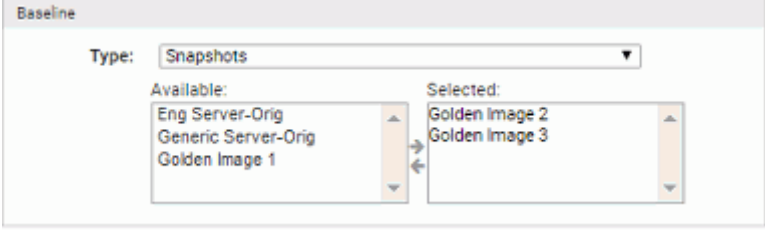
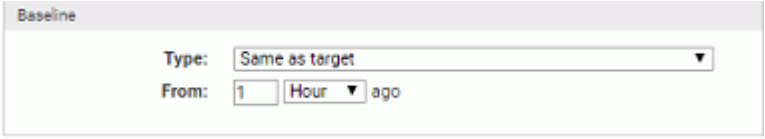
1. On the console menu, choose **Reports > Baseline Drift**.
2. On the Manage Baseline Drift Reports page, click on **Add Report**. This opens the Add Baseline Drift Report page.

3. Fill in the details of the report you want to create, referring to [Table 96, “Add/Edit Baseline Drift Report Details”](#) on page 646 for the parameter settings. Click **Show Advanced Options** if you want to configure options not shown in the current view.
4. When you have finished entering the report settings, click the **Enabled** status button if you want the report to begin comparing drift immediately.
5. Click **Save** to save the report and return to the Manage Baseline Drift Reports page.

Table 96: Add/Edit Baseline Drift Report Details

Item	Description
Copy settings from menu	(Available on the Add page only) Copy settings from an existing report to populate the details of your new report. You can make whatever changes you want to the copy. When you choose a report on this menu, the default name of your new report is <i>Copy of <the name of the existing report></i> .
Report name	The name that will appear on the Manage Baseline Reports page and the window banner for this report.
Description	Optional text that will help identify the purpose of the report.

Item	Description
Status radio buttons	Enabled means that the report results are automatically generated. Disabled turns off report generation <i>and</i> deletes the entire history of the report.
Target menu	What is to be analyzed in the report. The target Type options are: Computer – Track all file changes on the selected computer. Computers in policy – Track all file changes on all computers in the selected policy. Computer Filter – Track all file changes on computers that match the criteria specified in the filter. Advanced Filter – Track all file changes that match the criteria specified in the filter, which can include both file and computer criteria. All computers – Track all file changes on all of your computers. For each target Type except <i>All computers</i> , additional fields appear to allow you to complete the specification of the target.

Item	Description
Baseline menu	<p>What the target is compared to. The baseline options are:</p> <p>Computer – Compare target to the files found on the named computer at report run time.</p> <p>Computers in policy – Compare target to the files found on all computers (at report run time) in the policy selected from this menu.</p> <p>Computer Filter – Compare target to files from computers that match the criteria specified in the filter.</p> <p>Advanced Filter – Compare target to files that match the criteria specified in the filter, which can include both file and computer criteria.</p> <p>Snapshots – Compare target to the files in one or more selected snapshots.</p>  <p>Same as target – Compare the files on the target computer(s) to the files on the same computer(s) at a specified point in the past.</p>  <p>None – Calculate total drift of all computers without any baseline comparison. This choice generates a report that simply monitors all changes on a target set of machines since the agent was installed. This option does not allow tracking of <i>missing</i> files. If you keep the default Advanced Options, this choice essentially gives you the table of all unapproved files on your target systems, along with additional Drift and Risk information only available in Baseline Drift Reports. You can filter or sort by Risk if you choose to determine whether action is necessary on any of these unapproved files, and also see whether any particular group, user, or computer is contributing disproportionately to total Risk.</p> <p>For each Type choice except <i>None</i>, additional fields appear allowing you to complete specification of the baseline.</p>
Save button	Create the Baseline Drift Report by saving the parameters you have entered. Once created, the report is scheduled to run, unless you disable it.
Cancel button	Cancel the creation or editing of the report.
Show/Hide Advanced Options buttons	Shows or hides additional parameters for the report. See “Advanced Baseline Drift Report Options” for more details.

Advanced Baseline Drift Report Options

The Advanced Options section includes options that change the file types considered in a baseline drift analysis, the method of comparison between a baseline and its target, and the level of detail, and therefore the size, of the report when it is generated. Changing these options may affect performance, and also may create reports with considerably more detail for you to examine.

Advanced Options

Filter Options:

- Include approved files
- Include banned files
- Include initialized files
- Only include applications
- Only include executed files

Comparison Method: File content File location

Report Size: Full details Summary only

Save Cancel Hide Advanced Options

Advanced Options: File Filter Options

Filter options allow you to choose different types of files to include in the drift report. All of these options are off by default. They are essentially shortcuts for some of the more common options you can set by choosing Advanced Filters in either the baseline or target Type menu. The choices are:

- **Include approved files** – Files with a Local State of *Approved* are included in the baseline drift comparison.
- **Include banned files** – Files with a Local State of *Banned* are included in the baseline drift comparison.
- **Include initialized files** – Files initialized from a newly installed agent are included in the baseline drift comparison.
- **Include missing baseline files** – Baseline drift analysis includes tracking of files that exist in the baseline but are missing on the target systems (does not appear if baseline is Same as Target).
- **Only include applications** – Only files on your network that are executable (e.g., .exe or .com, but not Packages) are included in the baseline drift comparison.
- **Only include executed files** – Only files that actually have executed on your network are included in the baseline drift comparison.

Deciding which of the Filter Options to use depends upon your purpose in running a Baseline Drift Report. Although only unapproved files are included by default, you can run baseline drift reports that include locally Approved and/or Banned files. When both of those options are used, the drift report shows *every* new file of interest, which can be very useful if you want to see whether your systems have “drifted” from a golden image or known baseline. You might discover that some files you have approved should not have been, or that there is a large proliferation of banned files, which, although they cannot execute, indicate a problem.

Another situation in which including locally banned and approved files as well as missing baseline files might be useful is in an environment where systems must be absolutely standard, for example, point-of-sale systems. You can use drift reports to determine whether all your systems *exactly* match your golden disk image.

Advanced Options: File Comparison Method

Baseline drift reports use both file content (its hash) and file location (its full pathname) to identify added, missing, and changed files. The Advanced Options in Baseline Drift Report Details allow you to change *how* they use these factors:

- **File content** – By default, baseline drift reports use the *File content* method for comparisons. When this option is in effect, if a file in the baseline has the same hash as a file in the target, no drift is reported, regardless of the pathname (location) of the two files. A file in the same location (i.e., same path and filename on baseline and target) but with different hashes is considered modified on the target and so counts as drift. Baseline hashes not found anywhere on the target are reported as *missing files*, and target hashes not found on the baseline are considered *added files*.
- **File location** – If you choose *File location*, no drift is reported for the same hash found with the same path and filename on both baseline and target. Different hashes found at the same location (path and filename) are considered *modified files* and add to the drift number. And if the same hash is found in different locations, it is *not* considered a match. In that case, Baseline Drift Reports may report a new file (if the baseline had no file at the location where the file exists on the target), missing (if the target has no file where the baseline had one), or modified (if there is a file with the same name but a different hash on the baseline and target).

In some cases, the different comparison methods will have no effect on total drift. This is especially likely if you activate tracking of missing files as part of the drift report. If you maintain the default setting, however, and do not track missing files, the different comparison methods can produce different drift results, as the example in [Table 97](#) shows.

Table 97: Example: How different comparison methods affect drift

Files in Baseline	Files in Target	Drift by content	Drift by Location
C:\folder1\file1 (hash A)	C:\folder1\file1 (hash A)	None	None
C:\folder1\file2 (hash B)	C:\folder1\file2 (hash F)	1 new (hash F)	1 changed (file2)
C:\folder2\file3 (hash C)	C:\folder2\file3 (hash B)	1 changed (file3)	1 missing (hash C)
C:\folder2\file4 (hash D)		1 missing	1 missing
	C:\folder2\file5 (hash G)	1 new	1 new
Total Drift Including Missing Files		4	4
Total Drift Not Including Missing Files (default)		3	2

Advanced Options: Report Detail Level


The Advanced Options provide a choice of size for the baseline drift report. The default choice is *Full details*, which generates a drift report that includes details of top-level files and all individual files associated with them. The other choice is *Summary only*, which generates reports that include details at the top (file group) level and shows details of individual files only when requested (i.e., when you click on the file group to get details. The table shows some of the considerations in choosing one or the other of these options.

Table 98: Report Size Options

Differences	Summary Only Report	Full Details Report
Level of Detail	Initially reports results by file groups. Individual-file-level report is generated on demand when you click on a file group.	Contains individual files
Database size	Small size in database	Large size (approx. 10x larger than Summary)
Creation Speed	Faster to generate	Slower to generate
Report Access Speed	Slower to view	Faster to view
Compatibility with Dashboard	Not suitable for graphing (portlets) and extensive analysis because it lacks file-level details such as threat, trust, and publisher/company	Suitable for graphing and analysis by grouping, filtering, etc.

Using Filters in Target and Baseline Definitions

There are two types of filters on the Type menu for Target and Baseline definitions: Computer Filters and Advanced Filters. Advanced Filters includes all the filters types in Computer Filters. Once you choose the type, you can add as many different filters from its menu as you like. You also can add multiple filters of the same type.



Computer Filters are useful if you know that the only criteria you plan to use for specifying a baseline or target are computer-related. You have the following Computer Filter options:

- Computer
- Computer Tag
- IP Address

- Platform
- Policy

Although two of these duplicate choices on the Type menu, by using the Computer Filters type, you allow yourself to set multiple filters for computers. For example, you can specify that you want your baseline to include all computers in Low enforcement policies that have a Computer Tag of “Sales” or “Marketing”.

Advanced Filters are useful when you need to include criteria not available on the Computer Filters menu in your specification of a baseline or a target. You can still include computer filters, but Advanced Filters also allow you to use a large set of file criteria, including hash values, file prevalence, and threat level.

While most of the filter choices are self-explanatory, the File Type choice might not be. With the File Type filter, you can specify that your target or baseline includes *or excludes* the following choices:

- **Application:** Any executable (e.g. .exe or .com) except for Packages
- **Supporting File:** Any library loaded by an executable (e.g., .dll, .ocx, .sys)
- **Package:** Any installer (.exe with contents, such as a self-extracting zip or setup program)
- **Script File:** Any script or batch file (e.g., .bat, .vbs, .wsf)
- **Other:** Reserved for future types
- **Unrecognized Executed File:** A file that was not identified as an executable by App Control during initialization or later analysis, but that some process attempted to execute. The execution attempt causes the file to be added to the file lists in the console for tracking and management.
- **Unknown:** Files reported by older App Control Agents that don't provide file type information

Drift in Multi-Platform Environments

App Control supports installation of agents on Windows, Mac, and Linux computers. Because of the different platform software and applications found on different operating systems, it does not make sense to mix these different computers in a drift measurement. The “noise” level will make extraction of useful data difficult. Targeting all computers or all computers in a policy (unless the policy is platform-specific) in a drift report is not recommended.

If you have a multi-platform environment, possible ways to define a report that produces useful results are:

- Choose **None** as the Baseline Type. This will produce a report that monitors all changes on a target set of machines since the agent was installed, without tracking of missing files. By default, it lists all unapproved files on your target systems, along with additional Drift and Risk information.
- Choose **Same as Target** as the Baseline Type. This will produce a report that shows only the drift of each computer compared to itself.
- For other baseline types, you can create one drift report for each platform by choosing an Advanced Filter or Computer Filter on the Target menu and specifying the platform in that filter.

See [“To create a Baseline Drift Report:”](#) on page 646 for more information about specifying the parameters in a drift report.

Managing Snapshots

A snapshot is a listing of files (including their name, hash, and location) from one or more computers. You can use a single snapshot or a combination of snapshots as the baseline for a drift report. You can use filters to generate exactly the file list you want and then take a snapshot of that list of files. There are several locations in the console from which you can create snapshots. Once a snapshot is created, you can add or remove files from it as necessary.

Only Power Users, Administrators, and users in custom groups with view and manage snapshot permissions can create, modify and delete snapshots.

Platform Note: Mixing files from different operating system platforms (e.g, Windows, Mac, and Linux) in a single snapshot is not recommended.

Creating and Modifying Snapshots

There are two main ways to create a new snapshot:

- using all files on a particular computer
- using a file table, filtered or not, on a console page that includes the Snapshot button

To create a snapshot (or add to one) from all files on a computer:

1. On the console menu, choose **Assets > Computers**.
2. In the Computers table, click on the name of the computer whose files you want to use as a snapshot. The Computer Details page appears for that computer.
3. In the Actions menu on the right of the details page, click **Add Files to Snapshot**. The Add Files to Snapshot dialog appears:

The screenshot displays the 'Computer Details' page for a computer named 'WORKGROUP\APC-58609'. The 'General' section shows the IP address as 10.210.65.31, connection status as 'Connected', health check as 'Failed', and platform as 'Windows'. The 'Policy' section shows the policy as 'APC-58609-xicowo', mode as 'Control', and enforcement as 'Low (Monitor Unapproved)'. The 'Actions' menu on the right includes options like 'Change Policy', 'Delete Computer', and 'Add Files to Snapshot'. The 'Add Files to Snapshot' dialog is open, allowing the user to either choose an existing snapshot or create a new one.

4. To create a new snapshot, in the dialog, type in the name for the snapshot in the *Create new snapshot* box and click **Create**.

- or -

To add all of the files on the computer to an existing snapshot, choose an existing snapshot from the *Choose existing snapshot* menu and click **Add**.

A message appears confirming the creation or modification of the snapshot.

5. If you want to view the contents of your snapshot, choose **Reports > Baseline Drift** on the console menu and then click on the **Snapshot** tab. Your new or modified snapshot is displayed in the snapshots table.

Note

A snapshot of the files on a computer is static – it is the list of files that were on the computer when the snapshot was taken. You also can use a computer itself as a baseline for comparison, in which case the files on the computer when you run the report are the baseline.

To create a snapshot (or add to one) from a file table:

1. Go to the console page from which you want to create the Snapshot. For example, choose **Assets > Files** on the console menu to go to the Files page, and then click on **File Catalog**.
2. Choose the tabs, filters, columns, and/or Saved View you want to get the list of files you want in the snapshot.
3. Click the **Show Snapshots** link to show the Snapshot panel

The screenshot shows the 'Files' page in the VMware Carbon Black App Control console. The 'File Catalog' tab is selected, and the 'Files on Computers' sub-tab is active. The 'Add Files to Snapshot' panel is highlighted with a red box. This panel contains the following elements:

- Saved Views:** A dropdown menu set to '(none)' and an 'Add' button.
- Group By:** A dropdown menu set to '(none)'.
- Show Filters**, **Show Columns**, **Hide Snapshot**, **Export to CSV**, and **Refresh Table** links.
- Add Files to Snapshot** section:
 - Files to add:** Radio buttons for 'All files' (selected) and 'Checked files'.
 - Choose existing snapshot:** A dropdown menu set to '(none)' and an 'Add' button.
 - Create new snapshot:** A text input field and a 'Create' button.

Below the panel, there is an 'Action' dropdown and a status indicator 'Showing 75 out of 14380 item(s)'. A table displays a list of files with the following columns: 'Select 75', 'First Seen Date', 'First Seen Name', and 'Publisher or Company'. Two rows are visible:

Select 75	First Seen Date	First Seen Name	Publisher or Company
<input type="checkbox"/>	Jul 14 2020 08:30:38 AM	am_delta_patch_1.319.1375.0.exe	Microsoft Corporati
<input type="checkbox"/>	Jul 13 2020 07:57:30 AM	presentationhostdll.dll	Microsoft Corporati

4. If you want to individually select the files being added to the snapshot, check the box to the left of the file for each file you want to add, and click the **Checked files** radio button in the *Files to add* line of the *Add Files to Snapshot* panel. Otherwise, all files on the page are added to the snapshot.
5. To create a new snapshot, in the Snapshot box, type in the name for the snapshot and click **Create**. A new snapshot is created from the current table of files – it includes the files on *all* pages in the table, not just the currently displayed page.
- or -
To add all of the files in the current table to an *existing* snapshot, choose an existing snapshot from the *Choose existing snapshot* menu and click **Add**.
6. If you choose **Checked files**, you must check and add files for each page in the table – only the files checked on the currently visible page are added.
7. If you want to confirm that a new snapshot was created, choose **Reports > Baseline Drift** on the console menu and then click on the **Snapshot** tab. Your new snapshot appears in the snapshots table.

Viewing and Editing Snapshots

Once created, a snapshot may be viewed on the Snapshot tab of the Baseline Drift page.

To view a snapshot:





1. On the console menu, choose **Report > Baseline Drift**.
2. On the Baseline Drift page, click the **Snapshots** tab.


Baseline Drift: Manage Snapshots

Reports Snapshots

Group By: (none) Ascending

Show Filters | Show Columns | Export to CSV | Refresh Page

	Name	Date Created	Created By	Number of Files
 	SW Pilot Files	Jul 14 2020 01:45:56 PM	admin	6
 	APC-58609-baseline	Jul 14 2020 12:31:57 PM	admin	26503

3. Click either the name of the snapshot you want to view or the View Details button  in its row. A table of all of the files in the snapshot appears.













Home > Baseline Drift > Snapshots > Find Files

Find Files

Group By: (none) Ascending

Show Filters | Show Columns | Show Snapshot | Export to CSV | Refresh Table

Action Search: Automatically apply

<input type="checkbox"/> Select 75	File Name	Publisher or Company	Trust	Threat	File State
<input type="checkbox"/>  	vgauthservice.exe	VMware, Inc.	10		Unapproved
<input type="checkbox"/>  	vmacthlp.exe	VMware, Inc.	10		Unapproved
<input type="checkbox"/>  	vmtoolsd.exe	VMware, Inc.	10		Unapproved
<input type="checkbox"/>  	vsock.sys	VMware, Inc.	10		Unapproved

From the Snapshot Contents page, you can use any of the standard table tools (filters, column controls, etc.) to change your view of the files in the snapshot.

Managing Files in Snapshots

You can check one or more files in the snapshot and take the following actions:

- **Remove the checked file(s) from the snapshot** – Files you have checked when you choose **Remove from Snapshot** on the Action menu will be removed from the snapshot, but not from any computers on your network.
- **Approve or Ban the file(s)** – The Action menu provides commands for creating global or custom approvals or bans for checked files in the snapshot. Note, however, that there might be more efficient and flexible approval methods for handling a

particular file – for example, approving it by approving its publisher, or by approving the installer that generated the file.

- **View Carbon Black Reputation Data** – This retrieves information from Carbon Black File Reputation (if available) for the files checked when you choose the command.

Deleting Snapshots

On the Snapshot tab of the Baseline Drift Reports page, you can delete snapshots you no longer need. Before doing so, consider whether the snapshot is really no longer useful, or whether you can make it useful by adding files to or deleting them from it. You cannot recover a deleted snapshot.

To delete a snapshot:

1. On the console menu, choose **Reports > Baseline Drift**.
2. On the Baseline Drift page, click the **Snapshots** tab. This tab does not appear until you have saved at least one snapshot.
3. Click the Delete button in the row of the snapshot you want to delete, and in the confirmation box, click **OK**.

Displaying Baseline Drift Reports in Graphs

The tables on the Baseline Drift pages provide the greatest detail and flexibility in viewing drift results, but you might want a graphic representation of drift to use as a quick reference indicator of changes in files on your network. You display graphs of Baseline Drift Reports as graphic *portlets* on a *dashboard* page in the console.

The console includes pre-configured portlets, the individual graphs that make up a Dashboard, that provide baseline drift information. You can choose any portlet with “Drift” in its title to see an example of graphic presentation of drift information.

If you plan to create your own drift portlets, consider the following tips for making the information you display usable:

- The horizontal size of portlets varies according to the layout of the dashboard they are displayed on. You may need to move the portlet on the dashboard or change the dashboard layout to accommodate the data in a baseline drift portlet. You also can choose the data and the type of graph you display it in so that the portlet is appropriate for the presentation format.
- Consider how many items will appear on the X axis. The Portlet Editor does allow you to limit the items displayed on the X-axis to the 5, 10, or 15 with the highest or lowest values, but this means you are not seeing all of the data from the report. So if, for example, you have 1000 computers, you might choose to show drift by *policy* instead of by *computer* – you can always drill down to the more detailed information in console tables. (Similarly, if you use the “Split by” feature in a portlet, limit the number of items that will split the bar, column or other element in your portlet.)
- Use the Preview feature in the Portlet Editor to see how your data will appear. You can try out as many display options as you would like before you Save the portlet.
- If a Baseline Drift Report has a Report Size of *Summary Only* (an option in the Advanced Options panel), it will not have sufficient data for use in the Dashboard. Only reports that have a Report Size of *Full Details* can be displayed graphically.

The example below shows the same information presented in a Baseline Drift Report Results table, and then again in a graphic portlet. On a demonstration system with 5 or fewer computers running the agent, you will be able to easily view drift by computer in a graph. This is less likely to be useful in a production environment.

In tabular form, the drift report might look something like the following figure.

Daily drift of all computers


Saved Views: (none) View Mode: This report was generated on May 30 2017 03:03:55 PM.
 Show results by: Computers Files Group By: (none) Ascending

Show Filters | Show Columns | Export to CSV | Refresh Page

1 2 3 4 5

	Computer	Drift	Risk	Policy
	MYCORPILAPTOP-1	777	58.8	Research Team
	MYCORPILAPTOP-4	213	0.0	Research Team
	MYCORPILAPTOP-3	5	5.7	Standard Protection
	MYCORPILAPTOP-2	4	1.3	Standard Protection

The same information in a Dashboard would appear as shown in the next figure. Clicking View Details in the dashboard view brings you back to the full report table.



For more on dashboards, see [“Using and Customizing Dashboards”](#) on page 680.

Creating Baseline Drift Alerts

You can create an Alert to notify you and any other console users that baseline drift has crossed a threshold that you have set. When you enable a baseline drift alert, the triggering conditions are evaluated each time the report generation is complete.

Each time baseline drift conditions exist that meet the triggering conditions, the console highlights the alert in color and adds a Reset button, both on the Home page dashboard and the Alerts page. It also sends alert email to all subscribers to this alert. You can reset the alert manually by clicking the Reset button next to its name on the Alerts page. Drift alerts automatically reset when the drift in the specified drift report falls below the threshold for the specified parameter (user, computer, or policy).

See [“Using App Control Alerts”](#) on page 608 for more on alert behavior.

To create a baseline drift alert:

1. On the console menu, choose **Tools > Alerts** to display the Alerts page.
2. On the Alerts page, click the **Add Alert** button. The Alert Information page appears:

Add Alert

General

Alert Name:

Message:

Priority: Medium ▼

Status: Enabled Disabled

Type

Type: Baseline Drift Alert ▼

Description: Alerts subscribers when baseline drift factor reaches specified threshold

Mail Template: Default ▼

Criteria

Drift Report: Drift of all computers ▼

Alert When: Total risk ▼ per computer reaches ▼

Subscribers

Note: Alert must be created before email recipients can be specified

Reminder Mail

Status: Enabled Disabled

Remind Every: day(s) ▼

Auto Reset

Status: Enabled Disabled

Reset After: week(s) ▼

Create & Exit
 Create
 Cancel

3. In the General panel of the Alert Information window, enter an Alert name and a Message (what will be sent to subscribers when the alert is triggered).
4. In the Type panel, choose **Baseline Drift Alert** from the Type menu.
5. In the Criteria panel, choose the drift report whose data you want the alert to monitor. **Note:** If no drift reports have been created yet, the Drift Report line will display a message to that effect instead of the menu.
6. In the *Alert when* line, choose the threshold at which you want an alert to be triggered.
7. Click **Save** to create the alert.
8. On the Alerts page, click the View Details button next to the name of your new alert.
9. On the Alert Information page, in the Subscribers section, enter each email address to which you want alert email sent and click **Add** after each one.

10. To specify the email format, choose one from the menu to the right of the address box.
11. To resend alert emails periodically as long as the alert is not reset, set Reminder Mail to **Enabled** and choose a time interval.
12. Click **Save**.

Chapter 23

Advanced Threat Detection

This chapter describes how you enable and use App Control's Advanced Threat Indicators, and how you can monitor threats through events, file details, and alerts.

Sections

Topic	Page
Overview	661
Indicator Sets for Threat Detection	662
Indicator Set Exceptions	666
Monitoring Threat Reports	671
Threat Views on the Events Page	672
Threat Events in Syslog Output	674
Threat Views on the Files Pages	675
Threat-Related Alerts	676
Responding to Threats	677

Overview

App Control includes many features that help you monitor activities on your endpoints. To enhance these capabilities, App Control provides a set of advanced detection features, including:

- **Advanced Threat Indicators (ATIs)**, which are rules grouped in Indicator Sets that aid in detecting particularly threatening or suspicious activity on systems reporting to your App Control Server
- **Detection Views** into your App Control database that highlight detection-related data provided by the ATIs and other App Control features

Advanced Threat Indicators may indicate malicious activity based on an event or sequence of conditions on an endpoint. This has the potential to provide broader coverage and earlier warning than a detection system relying solely on a snapshot of a point in time. A conventional Indicator of Compromise (IOC) might report on the existence of a suspicious file or registry setting only after the fact. Because App Control's advanced detection feature also uses dynamic events as part of its implementation, it can provide real-time indication of suspicious activity and capture metadata for related events, such as the creation of a suspicious file.

While ATIs are strictly for reporting purposes, you may be able to remediate a detected threat using other App Control capabilities, or by actions outside of the console. For example, you could create a ban for a file reported as a threat or create a custom rule that bans an action in a particular location when conducted by a certain process. Also the App Control Event Rule capability allows you to immediately ban or delete *any* file that appears in a threat-related event.

The summary steps for using Advanced Threat Detection are:

- **Enable Indicator Sets for Detection** – On the console Indicator Sets page (**Rules > Indicator Sets**), enable the Indicator Sets that you want activated. Once the Indicator Sets are enabled on the server, the ATIs are committed to all the agents. Then, when the conditions specified by any of the ATIs occur, new detection events are sent to the server. See [“Indicator Sets for Threat Detection”](#) on page 662.
- **Monitor Threat Reports** – Periodically check for suspicious or threatening events or files using the Saved Views on the Events and Files pages. See [“Monitoring Threat Reports”](#) on page 671.
- **Fine-tune Reporting** – If you see detection-related events that you do not want reported, either disable the Indicator Set that detected them (if you are sure you do not want any reports from that Indicator Set) or create an Indicator Exception for the specific file reported in the event. See [“Indicator Set Exceptions”](#) on page 666. On the other hand, if you see detection-related events that you consider high priority, consider creating alerts for those events. See [“Threat-Related Alerts”](#) on page 676.
- **Remediate Threats** – If you see a threat that must be remediated, consider creating a App Control rule (for example, a ban, custom rule, or event rule) to prevent malicious action by the threat and/or take action outside of App Control (for example, deleting files or creating firewall rules) to respond to the threat. See [“Responding to Threats”](#) on page 677.

ATIs work with agents at any Enforcement Level (other than Disabled), although the conditions that lead to threat detection should be less likely in High Enforcement.

Upgrade Notes

- If you used the separately installed Advanced Detection features in Bit9 Platform 7.0.0 or 7.0.1, be aware that in this release, ATIs are grouped in Indicator Sets rather than Updaters. This means that you view, enable, and disable Indicators on the Indicator Sets page. Also, if you created an ATI-related Event Rule in one of these prior 7.0.x releases, you must create a new Event Rule that reflects the current implementation. There is, however, a Saved View on the Events page – *Threat Reports - Legacy* – that will show you threats reported by the previous ATIs. Event Rules created in v7.2.0 and later releases should not require any changes.
- Indicator Sets are disabled after an upgrade. If you are upgrading from a pre-7.2.0 release, to re-enable threat indicators, choose **Rules > Indicator Sets**, check the box next to each Indicator Set you want to enable, and choose **Enable Indicator Sets** on the Action menu.

Indicator Sets for Threat Detection

An Indicator Set is a group of related ATIs (detection rules) for the platform specified by its name. To view and manage Indicator Sets, a console user must have *Manage indicator sets* permission enabled. This permission is enabled by default for Administrators and Power Users. See [“User Role Permissions”](#) on page 106 for details on enabling user permissions

The following list describes default Indicator Sets provided with the initial release of v8.0.0 and the types of ATIs they contain. Note that Indicator sets may be added, removed, or modified in cloud-based updates or future versions of App Control. See [“Updates to Indicator Sets”](#) on page 670 for more details.

- **Windows Admin Tool Tracking** – The ATIs in this group are designed to track legitimate admin tools commonly abused by attackers. A good example of this is psexec activity. While more often than not, this is used legitimately, its use in attacks is so prevalent, and tracking of this activity so helpful in forensics investigations, that alerting on this activity was included in advanced detection. However, unlike other groupings, the rules contained here are likely to generate events that are not indicative of malicious activity. As such, this updater may be a good candidate for disabling in certain environments.
- **Windows Application Behavior** – The ATIs in this group detect behavior not normally expected from the type of application performing it. For example, one ATI in this group, called “Possible exploit of document handling application”, reports an event if an application such as Microsoft Excel creates an unknown executable (e.g., foo.exe).
- **Windows POS Indicators** – The ATIs in this group are specific to file and registry artifacts created during attacks on point-of-sale (POS) style systems. They are based on publicly released information about these types of attacks.
- **Windows Process Injection** – The ATIs in this group detect injection of suspicious code into specific system processes. For example, one ATI in this group, “Possible password hash tool execution”, reports an event if a process tries to harvest cached

- password hashes on a system. In general, this indicator set reports issues involving memory rules.
- **Windows Ransomware Indicators** – The ATIs in this group are designed to identify signs of ransomware-type malware such as cryptolocker/cryptoblocker and associated variants. Both registry- and file-based activities are included. Indicators are based both on malware analysis and published reports.
 - **Windows Startup Configuration** – The ATIs in this group detect suspicious changes to the Windows startup configuration.
 - **Windows Suspicious Based on File Name** – The ATIs in this group detect files whose names indicate that they are suspicious or malicious. For example, if a file has a name or file extension that is similar to a legitimate file (e.g., “iexplore.exe”) but is modified slightly (e.g., “Lexplore.exe”), a ATI in this group reports it. Files with the names of known malware or suspicious extensions are also reported.
 - **Windows Suspicious Based on Parent** – The ATIs in this group detect suspicious activity based on the parent process of an executable.
 - **Windows Suspicious Based on Path** – The ATIs in this group detect file activity in suspicious location, such as file execution in the Recycle Bin or System Volume.
 - **Windows Suspicious Based on Path and File Name** – The ATIs in this group detect suspicious activity based on both file path and file name. For example, one ATI reports System files executing outside the System folder. Another indicator in this group reports execution of rarely used system utilities.
 - **Windows System Configuration** – The ATIs in this group detect suspicious system configuration activity, such as firewall or name resolution tampering, or installation of a language pack.
 - **Mac Application Behavior** – The ATIs in this group detect behavior that is not normally expected from the type of application performing it. For example, one ATI in this group reports an event if an application such as Microsoft Excel creates an unknown executable. Another ATI in this group detects shells being spawned from a browser.
 - **Mac Shell Activity** – The ATIs in this group detect suspicious use of a command shell.
 - **Mac Suspicious Based on Path** – The ATIs in this group detect activities that are suspicious because of where they are attempted, such as execution attempts from the Trash folder.
 - **Mac Suspicious Based on Path and File Name** – The ATIs in this group detect unusual behaviors from a known path. For example, one indicator in this group reports an event if a file is created that is indicative of a known backdoor.
 - **Mac System Configuration** – The ATIs in this group detect suspicious changes to system configuration, such as attempts to escalate privileges.
 - **Linux Possible Backdoor** – The ATIs in this group detect files associated with backdoors to the Linux secure shell.
 - **Linux Startup Configuration** – The ATIs in this group detect suspicious changes to the Linux startup configuration.

To view, enable or disable Indicator Sets:

1. On the console menu, choose **Rules > Indicator Sets**. The Indicator Sets page appears.

<input type="checkbox"/>	Indicator Set Name ▲	Version	Enabled	Platform	Policy	Date Updated
<input type="checkbox"/>	Linux Possible Backdoor	1405	No	Linux	All Policies	Jul 13 2020 06:37:07 AM
<input type="checkbox"/>	Linux Startup Configuration	1405	No	Linux	All Policies	Jul 13 2020 06:37:07 AM
<input type="checkbox"/>	Mac Application Behavior	1407	No	Mac	All Policies	Jul 13 2020 06:37:07 AM
<input type="checkbox"/>	Mac Shell Activity	1407	No	Mac	All Policies	Jul 13 2020 06:37:08 AM
<input type="checkbox"/>	Mac Suspicious Based on Path	1405	No	Mac	All Policies	Jul 13 2020 06:37:08 AM
<input type="checkbox"/>	Mac Suspicious Based on Path and File Name	1407	No	Mac	All Policies	Jul 13 2020 06:37:10 AM
<input type="checkbox"/>	Mac System Configuration	1409	No	Mac	All Policies	Jul 13 2020 06:37:08 AM
<input type="checkbox"/>	Windows Admin Tool Tracking	1408	No	Windows	All Policies	Jul 13 2020 06:37:10 AM
<input type="checkbox"/>	Windows Application Behavior	1408	No	Windows	All Policies	Jul 13 2020 06:36:56 AM
<input type="checkbox"/>	Windows POS Indicators	1405	No	Windows	All Policies	Jul 13 2020 06:37:06 AM
<input type="checkbox"/>	Windows Process Injection	1405	No	Windows	All Policies	Jul 13 2020 06:36:58 AM
<input type="checkbox"/>	Windows Ransomware Indicators	1407	No	Windows	All Policies	Jul 13 2020 06:37:06 AM
<input type="checkbox"/>	Windows Startup Configuration	1405	No	Windows	All Policies	Jul 13 2020 06:36:59 AM
<input type="checkbox"/>	Windows Suspicious Based on File Name	1405	No	Windows	All Policies	Jul 13 2020 06:37:00 AM
<input type="checkbox"/>	Windows Suspicious Based on Parent	1407	No	Windows	All Policies	Jul 13 2020 06:37:11 AM
<input type="checkbox"/>	Windows Suspicious Based on Path	1405	No	Windows	All Policies	Jul 13 2020 06:37:00 AM
<input type="checkbox"/>	Windows Suspicious Based on Path and File Name	1405	No	Windows	All Policies	Jul 13 2020 06:37:00 AM
<input type="checkbox"/>	Windows System Configuration	1408	No	Windows	All Policies	Jul 13 2020 06:37:01 AM

18 Items Page 1/1

2. Check the box next to the name of each Indicator Set you want to enable and then choose **Enable Indicator Sets**.

- or -

Check the box next to the name of each Indicator Set you want to disable and then choose **Disable Indicator Sets**.

3. To see details and exceptions for any one Indicator Set, click the View Details button next to its name in the table.

Initially, all Indicator Sets are disabled. You can enable and disable these rule groups as you choose. For example, if one Indicator Set is generating too many events not of interest in your environment, you can turn it off on the Indicator Sets page. You also can create exceptions to an Indicator Set without disabling all of the indicators in the set. See [“Indicator Set Exceptions”](#) on page 666 for more details.

As with other console tables, you can use the Group By menu, the Show Filters link and the Show Columns link to modify your view of the Indicator Sets table. [Table 99, “Indicator Set Parameters”](#) on page 665 provides a description of all available columns.

Indicator Set Details

In the Indicator Sets table, clicking on the View Details button next to the name of a set opens the Indicator Set Details page for that set. This page includes:

- key details about the Indicator Set, including its name, version, and history

- radio buttons and checkboxes for enabling and disabling the Indicator Set, and for specifying the policies in which the set is active
- an Exceptions panel that shows any exceptions to the Indicator Set and allows them to be enabled, disabled, and deleted
- a Recent Events link in the Related Views menu that opens the Events page filtered to show recent events involving this Indicator Set
- two links in the Related Views menu that show either all computers that have received the rule to enable the Indicator Set or all computers that have not received the rule

?
Indicator Set Details

Indicator Set Name: Windows Suspicious Based on Parent

Version: 1407

CL Version:

Status: Enabled Disabled

Platform: Windows

Rule Applies To: All Current and Future Policies
 Selected policies

Date Created: Jul 13 2020 06:37:11 AM

Date Updated: Jul 13 2020 06:37:11 AM

Date Modified: Jul 13 2020 06:37:11 AM

Last Modified By: System

Save & Exit
 Save
 Cancel

Exceptions

[Show Filters](#) ▾ | [Show Columns](#) ▾ | [Refresh Page](#)

Action ▾

<input type="checkbox"/>	Exception Name ▲	Type	Enabled	Target	Process	User or Group
There are no items to display.						

0 items Page 0/0 25 ▾ rows per page

Table 99 shows the fields available in the Indicator Sets table and the Indicator Set Details page.

Table 99: Indicator Set Parameters

Field	Description
Indicator Set Name	Name of the Indicator Set. Names are assigned when shipped, and include the platform and general purpose of the ATIs in the set.
Version	The version of this Indicator Set. If new versions have been downloaded from the Carbon Black File Reputation, the version number will increment to indicate this.
CL Version	The App Control configlist in which this Indicator Set is enabled. This field is blank if the Indicator Set is not enabled. You can use this field (and the Related Views menu) to determine which endpoints have received this rule.

Field	Description
Status (Details page) Enabled (Table)	On the Details page, Status radio buttons make this Indicator Set Enabled or Disabled. In the Indicator Sets table, the Enabled field shows 'Yes' or 'No'.
Platform	Platform (Windows, Mac, or Linux) for which this Indicator Set is effective.
Rule Applies To (Details page) Policy (Table)	On the Details page, the radio buttons allow you to apply the rule to All policies or Selected policies . If you choose Selected policies , a list of all policies on your App Control Server appears, each with a checkbox. In the Indicator Sets table, the Policy field shows the policies for which the set is activated.
Date Created	Date and time this Indicator Set was first seen on this App Control Server.
Date Updated	Date and time this Indicator Set was last updated to a new version. If there have been no updates to this Indicator Set, this is the same as Date Created.
Date Modified	Date and time of the last <i>user-initiated</i> change to the Indicator Set configuration. This includes enabling or disabling the Indicator Set and changes to the policies it applies to.
Last Modified By	Console user that made the most recent change to editable parameters of this Indicator Set
Exceptions Panel (Details page only)	On the Details page, this panel lists any exceptions made for this Indicator Set. See "Indicator Set Exceptions" on page 666 for more on exceptions and Table 100, "Exception Details (in Indicator Sets)" on page 669 for a description of Indicator Set Exception parameters.

Indicator Set Exceptions

Indicator Set Exceptions are modifications of the Indicator Set that eliminate reports for actions that match the exception. They allow you to reduce or eliminate reporting of events that are not of interest to you while still leaving the rest of the Indicator Set functionality enabled. To create an Indicator Set exception, you identify an ATI-related event on the Events page that you would like to remove from future reporting. You can create an exception specific to that event automatically, or you can modify the exception so that applies to a broader or narrower range of targets, processes, or users.

Indicator Set Exceptions are specific to the Indicator Set that generated the event you use to create them. You can create multiple exceptions at once, but you cannot create an exception using a non-ATI-based event.

To create Indicator Set Exceptions (default method):

1. If the events for which you want to create exceptions are not displayed, choose **Reports > Events** on the console menu and then choose the **Threat Indicators** Saved View. You can also choose an event from another view, but using Threat Indicators ensures that the events shown all have an associated Indicator Set.
Note: You also can choose the **Recent Events** link on an Indicator Set Details page to see all recent events for that set.
2. If necessary, change the Max Age value to view older events.
3. When one or more relevant events are displayed, check the box next to each one and on the Action menu, choose **Create Indicator Set Exceptions**. A status message at the top of the page will indicate if the exceptions have been successfully created, or will show an error if they have not. A common error is selection of an event that does not have an Indicator Set.

Each exception created in this way uses the name of the Indicator Set plus incrementing digits (e.g., the first exception to the Windows System Configuration set is named "Windows System Configuration Exception 1").

You can edit an Indicator Set Exception once it is created (including its name), or you can add special parameters at the time of creation by choosing an *Create an advanced Indicator Set Exception*. However, an advanced Indicator Set Exception may be created for only one event at a time.

To create an advanced Indicator Set Exception:

1. If the event for which you want to create the exception is not displayed, choose **Reports > Events** on the console menu and choose the **Threat Indicators** Saved View. You can also choose an event from another view, but using Threat Indicators ensures that the events shown all have an associated Indicator Set.
Note: You also can choose the **Recent Events** link on an Indicator Set Details page to see all recent events for that set.
2. If necessary, change the Max Age value to view older events.
3. When the event for which you want to create an advanced exception is displayed, check the box next to it and on the Action menu, choose **Create an advanced Indicator Set Exception**. The Add Indicator Set Exception dialog appears with the Indicator Set and Platform entered in read-only form and the other parameters editable. Note that if you check more than one box, an error message appears.
4. In the Add Indicator Set Exception dialog box, enter an Exception Name and optionally a Description.
5. Edit the other parameters to create the rule you want. These parameters are described in [Table 100, "Exception Details \(in Indicator Sets\)"](#) on page 669.
6. When you have finished configuring the exception, click the **Save** button if you want to stay on the page or the **Save & Exit** button to return to the Events page.

The new exception appears in the Exceptions panel of the Indicator Set Details page.

Indicator Set Exception Details

Each Indicator Set Details page includes an Exceptions panel. If exceptions have been created for this set, they appear in a table in that panel.

Indicator Set Details ?

Indicator Set Name: Windows Suspicious Based on Parent

Version: 1407

CL Version: 1169

Status: Enabled Disabled

Platform: Windows

Rule Applies To: All Current and Future Policies
 Selected policies

Date Created: Jul 13 2020 06:37:11 AM

Date Updated: Jul 13 2020 06:37:11 AM

Date Modified: Jul 14 2020 02:53:38 PM

Last Modified By: admin

Save & Exit
 Save
 Cancel

Exceptions

[Show Filters](#) | [Show Columns](#) | [Refresh Page](#)

Action

	Exception Name	Type	Enabled	Target	Process	User or Group
<input type="checkbox"/>	Windows Suspicious Based on Parent Exception 1	Path	Yes	c:\windows\...	c:\program f...	NT AUTHORITY
<input type="checkbox"/>	Windows Suspicious Based on Parent Exception 2	Path	Yes	c:\windows\...	c:\windows\...	NT AUTHORITY

2 items Page 1/1 25 rows

The table shows the Exception name and other details of the exception, and like other console tables, it can be modified using the Show Filters and Show Columns links. The Action menu allows you to Enable, Disable, and Delete exceptions.

When you click on the View Details button for an exception in the table, the Indicator Set Exception Details page appears.

Table 100: Exception Details (in Indicator Sets)

Field	Description
Indicator Set Name	Name of the Indicator Set to which this exception is applied. Names are assigned when shipped, and include the platform and general purpose of the ATIs in the set.
Indicator Name	Name of the specific ATI in the Indicator Set for which an exception is being made
Exception Name	Name of this exception. This is provided automatically if the exception is created using the Create Indicator Set Exceptions command on the Action menu. Automatic naming uses the name of the Indicator Set plus incrementing digits (e.g the first exception to the Windows System Configuration set is named “Windows System Configuration Exception 1”. If the Exception is created using Create an advanced Indicator Set Exception , the name is entered by the console user. In either case, the name may be changed later.
Description	Additional information about the exception. This can be any text you choose to enter. (Optional)
Status	Radio buttons that Enable or Disable this exception.
Platform	Platform (Windows, Mac, or Linux) to which this Exception applies.
Type	The type assigned to this exception when it was created (not editable). The possible values are Path, Process and Registry.

Field	Description
Target	The Target of the action for which the exception was created. There may be multiple values in this field, and the values that are used depend upon the Exception Type: <ul style="list-style-type: none"> • Path – File paths or file names • Process – Processes • Registry – Registry paths Specification of paths and processes in rules is described in the Custom Rules chapter: “Specifying Paths and Processes” on page 407 shows details on specifying a process in rule pages and Table 58 shows process menu options.
Process	This menu allows you to limit the exception so that it is applied only when certain processes attempt to take action matching the target specification. “Specifying Paths and Processes” on page 407 shows details on specifying a process in rule pages.
User or Group	The users or groups to which this exception applies. Specification of users and groups is described in the Custom Rules chapter, in the section “Specifying Users or Groups” on page 424.
Date Created (Table only)	Date and time this Exception was created.
Date Modified (Table only)	Date and time this Exception was last modified.
Created By (Table only)	Console user that created this Exception
Last Modified By (Table only)	Console user that last modified this Exception

Updates to Indicator Sets

App Control provides a mechanism for automatic periodic updates to Indicator Sets. This may involve entirely new Indicator Sets, new indicators added to existing sets, reorganization of Indicators Sets, or changes to existing indicators. These changes are delivered automatically when available if both of the following are true:

- You have Carbon Black File Reputation enabled.
- You have also enabled automatic Indicator Set updates on the System Confirmation/Advanced Options page.

See [“Activating Carbon Black File Reputation”](#) on page 763 for more information about enabling Carbon Black File Reputation and [“Advanced Configuration Options”](#) on page 744 for more information about enabling automatic Indicator Set updates.

You may leave automatic updates enabled all the time or temporarily enable updates periodically if you choose. Updates are scheduled to be delivered within 24 hours and often appear sooner than that.

If you receive automatic updates, the state of Indicator Sets is as follows:

- New Indicator Sets are added in a disabled state.
- Existing Indicator Sets remain enabled or disabled according to the state they were in prior to updating. This is true even if the upgrade adds or modifies threat indicators in the existing Indicator Set.

Alerts for Tracking Indicator Set Updates

There is a built-in **Indicator Set Alert** that, when enabled, will inform you of the following Indicator Set changes:

- Indicator Set updated
- Indicator Set created
- Indicator Set deleted

This alert may be especially useful if you are only enabling the automatic updates temporarily – you will know when to turn off the updates. See [“Using App Control Alerts”](#) on page 608 for more on enabling and configuring alerts.

You can also tell whether a detection Indicator Set has been updated by reviewing Version, Date Created, and Date Updated fields on the Indicator Set Details page or Indicator Sets table.

Monitoring Threat Reports

Suspicious or threatening activity is reported through Saved Views on the console Events page and the Files pages. Check these views periodically as part of your threat monitoring activity. In addition to providing information, monitoring these threat reports also help you take actions to improve reporting and remediate threats:

- **Create Indicator Set Exceptions** – If you see specific threat-related events that you do not want reported, you can create Indicator Set Exceptions to eliminate reporting of those events. See [“Indicator Set Exceptions”](#) on page 666.
- **Disable Indicator Sets** – If you determine that a particular Indicator Set always reports events that are not of interest to you, you can disable the Indicator Set. See [“Indicator Sets for Threat Detection”](#) on page 662.
- **Enable Indicator Sets** – If you have not enabled all Indicator Sets and you think that certain critical activity is not being reported, see whether the disabled Indicator Sets would report that activity. See [“Indicator Sets for Threat Detection”](#) on page 662.
- **Create Alerts** – If you see detection-related events that you consider high priority, consider creating alerts for those events. See [“Threat-Related Alerts”](#) on page 676.
- **Remediate Threats** – As you monitor threats, you may see events that require remediation. This remediation might involve actions done outside of App Control, creation of App Control rules, or some combination of the two. See [“Responding to Threats”](#) on page 677.

Threat Views on the Events Page

On the console Events page, suspicious or threatening activity is reported in several Saved Views, some of which require Indicator Set activation and some of which use other data. The following Saved Views are threat-related:

- **Threat Indicators** – This view shows threats detected by the enabled ATIs in the Indicator Sets on App Control-managed computers. If no Indicator Sets have been enabled, this view will be empty. More details about these reports are shown below in the section [“Reviewing Threat Event Reports”](#) on page 673.
- **Threat Indicators - Legacy** – This view shows threats detected by the ATIs that were installed in releases prior to v7.2.0. If you did not install the Detection Enhancement in a prior release, this view will be empty.
- **Threat Report - Suspicious Executable Created by Shell** – This view shows events in which certain executable files are created by cmd.exe or powershell.exe in locations such as the system directory, RecycleBin, or AppData.
- **Threat Report – Suspicious Files by Location** – This view shows events in which a file is first seen or executed on any computer, or first appears (unapproved) on at least one computer, in an unusual, suspicious location. An example would be unexpected file activity in the Recycle Bin.
- **Threat Report – Suspicious Files by Name** – This view shows events in which a file is first seen or executed on any computer, or first appears (unapproved) on at least one computer, with a suspicious name. This is often a name similar to the name of a legitimate Windows file. For example, discovery of a file named svch0st.exe (using zero in place of the lowercase ‘o’ in svchost.exe) would appear in this event view.
- **Threat Report – Suspicious Files by Parent** – This view shows events in which an unknown, or low prevalence, executable file is written by a program that should not normally create such files. An example of this would be an executable file created by Adobe Reader; this is often indicative of a malformed- or malicious-PDF-style attack.

To view threat reports on the Events page:

1. Choose **Reports > Events** on the console menu.
2. On the Saved Views menu, choose the **Threat** view you want to examine.

Fields in Threat-Related Events Views

Certain fields in the Events table are of particular interest in the Threat views. Some are visible in the table by default and some may be added. They include:

- **Indicator Set** – The name of the Indicator Set containing the indicator that triggered the event.
- **Rule Name** – The name of the rule that triggered the event. For detection events, these are descriptions of the suspicious activity being detected.
- **Indicator Name** – This optional field is the same as the Rule Name for threat events. It is included to make it easier to identify threat events in Syslog output.
- **Process Threat** – The threat level for the process attempting an action in this event, if reported by Carbon Black File Reputation.
- **Process Trust** – The trust level for the process attempting an action in this event, if reported by Carbon Black File Reputation.

- **Process Prevalence** – The prevalence of the file associated with the Process field of the event. Prevalence is the number of computers on which at least one instance of the process file exists.
- **File Threat** – The threat level for the file acted upon in this event, if reported by Carbon Black File Reputation.
- **File Trust** – The trust level for the file acted upon in this event, if reported by Carbon Black File Reputation.
- **File Prevalence** – The prevalence of the file acted upon (the file in the File Name field) in this event. Prevalence is the number of computers on which at least one instance of the file exists.

Note

The initial values and later updates to threat, trust and prevalence data are provided based on access to Carbon Black File Reputation and scheduling of App Control tasks. Updates may have a delay.

Reviewing Threat Event Reports

Different event views provide different types of information, and cover different time windows.

The **Threat Indicator** view is likely to show the most recent or serious potential threats. Because of this, you might choose to concentrate on this view first. However, keep in mind that the Threat Indicators view shows only matching events that occur *after* you enable one or more Indicator Sets.

For an event in the Threat Indicators view, both the Indicator Set and the Rule Name are shown for the ATI that triggered the event. This shows you the type of threat the rule identified. It also provides a way to identify the source of over-reporting or false positives, and so helps you decide whether you want to disable an Indicator Set or create an exception for certain rules within it.

The **Threat Report** views make use of standard App Control events, including those that were present before you added the enhancement. They can report on matching events for whatever time period you choose on the Max Age menu, regardless of whether you have any of the Indicator Sets enabled. Like all events views, the maximum time frame for which threat events can be viewed is delimited by the database trimming choices in effect for your App Control database.

The Description field is also useful for reviewing events. Depending upon the event, it may identify the file that was written, modified or deleted, the process that acted on the file, and other pertinent data. For example, events generated by the following ATIs might have these descriptions:

Rule Name	Sample Description
Suspicious executable based on name	File 'c:\documents and settings\user\temp\explore.exe' was modified or deleted.
Unusual change to startup configuration	Modification of registry '\registry\machine\software\microsoft\windows nt\currentversion\winlogon\shell' was allowed.

Some of the information in the description is also available in specific fields that you can add to the view.

Note

Unlike the Events page threat views, the File Catalog threat view reports on both existing and historical files in the file inventory. If a file matching the view parameters ever existed on an agent-managed computer reporting to your server, it will be included in the view. See [“Threat Views on the Files Pages”](#) on page 675.

Showing and Modifying View Parameters

You can view the filters that are used to create the threat-related event views by clicking on the Show Filters button when the view is being displayed. The example below shows the filters used to build the Threat Indicators view.

The screenshot shows the 'Events' page in the VMware Carbon Black App Control interface. The 'Saved Views' section shows 'Threat Indicators' as the current view, with a 'Cache' button and an 'Add' button. The 'Group By' dropdown is set to '(none)'. Below this, there are links for 'Hide Filters', 'Show Columns', 'Export to CSV', 'Access Event Archives', and 'Refresh Table'. The 'Filters' section is expanded, showing a list of filter conditions. The first filter is 'Indicator Set' with the value 'is not empty'. The second filter is 'Subtype' with the value 'is'. Below these, there is a list of rules, each with a plus icon and a dropdown arrow. The rules are: Report access (Memory Rule), Report execution (Custom Rule), Report execution (removable media), Report read (removable media), Report write (Custom Rule), Report write (Registry Rule), and Report write (removable media). At the bottom of the filter list, there are 'Apply', 'Cancel', and 'Reset' buttons.

You can modify these views to add, remove, or modify filter conditions to further refine the view to either eliminate uninteresting events or broaden the scope of events displayed. Although modifications of default Saved Views cannot be saved, you can save your modified view under a new name.

By default, events from the past day are shown in the threat views. You can choose a different time period on the Max Age menu on the Events page.

Threat Events in Syslog Output

Threat-related events are exported to Syslog with other App Control events. To decide what to filter or search for in Syslog, you can choose one of the threat views and review the Rule Names shown in the table to see the specific rules that generated an event. You also can search for any event in Syslog that contains an Indicator Name field (“indicatorName” in raw output, mapped to different strings depending upon Syslog format), which will identify it as a threat detection event. For these events, Indicator Name is the same as Rule Name.

Events		
Saved Views: (The Current View Has Unsaved Changes - Discard)		
Threat Indicators	Add	Group By: (none) Ascending
Show Filters	Show Columns	Export to CSV Access Event Archives Refresh Page
Action	Search:	Automatically ap
Timestamp	Indicator Set	Rule Name
May 12 2017 02:35:33 PM	Windows Admin Tool Tracking	Processes started by Powershell remoting (WinRM)
May 12 2017 02:34:40 PM	Windows Admin Tool Tracking	Powershell or WinRM remoting activity
May 12 2017 02:34:23 PM	Mac System Configuration	Suspicious OSX persistence
May 12 2017 02:34:06 PM	Mac Shell Activity	Suspicious shell use
May 12 2017 02:33:57 PM	Windows System Configuration	Modification of the powershell execution policy
May 12 2017 02:30:33 PM	Mac Application Behavior	Shell Spawned from a browser
May 12 2017 02:29:48 PM	Linux Startup Configuration	Unusual change to startup configuration
May 12 2017 02:29:41 PM	Windows System Configuration	Possible name resolution tampering

One other potential search approach is to filter Syslog output to show just event subtypes that begin with "Report " (except for "Report execution block") – these are the subtypes for threat-related events. To see the specific list of event subtypes for an Events page view, you can choose the view from the Saved Views menu and then click on **Show Filters**.

See the separate document *VMware Carbon Black App Control Events Guide* for more information on the Syslog output available from the App Control Server.

Exporting Threat Event Data to CSV Files

As with other tables in the console, data in threat-related tables on the Events page may be exported to CSV files, which can be useful for analysis of threats in external tools. If you plan to export this information, consider using the Show Columns feature on the console page to add all columns to the table. This assures that all potentially useful information about a threat event is included in the export.

To export threat events to a CSV file, set up the table with the view, columns, and Max Time value you want and then click **Export to CSV**.

Threat Views on the Files Pages

The Files pages include views that report on the existence of suspicious or threatening files, even if they were created prior to the installation of the App Control Agent on an endpoint.

To view the Files pages, choose **Assets > Files** on the console menu and choose the tab for either **File Catalog** or **Files on Computers**. The following threat-related Saved Views are available on these tabs:

- **Threat Report - Suspicious Files by Extension** (File Catalog only) – This view identifies files that have been analyzed and determined to be executables by App Control but have an extension that is not an executable type. Malware often tries to disguise itself by using normally benign file extensions such as “.gif” or “.jpg”.
- **Threat Report - Suspicious Files by Name** – (Files on Computers only) This view shows files in the inventory that have names similar to the name of a common file (such as an operating system file), zero trust level in Carbon Black File Reputation, and a File State of Unapproved.

As with Event views, you can click on the **Show Filters** button on the Files pages to see the extension and other parameters that create these views. The views have the potential to produce many “false positive” results. To reduce the number of results, additional factors such as file trust, size, and publisher are used in this view. You can further modify and save the view under another name to create your own version of a threat report for files.

Threat-Related Alerts

You can create an alert that is triggered whenever an ATI determines that a potential threat has occurred. Alerts for this purpose are Event Alerts, and can be fine-tuned to include or eliminate certain types of threat events. All Event Alerts must include at least one Subtype in the Select Event Properties panel. In the example below, an alert is being created that used the same properties as the Threat Indicator view on the Events page, and so this event will be triggered any time an event occurs that would be displayed in that view.

Add Alert

General

Alert Name:

Message:

Priority:

Status: Enabled Disabled

Type

Type:

Description: Alerts subscribers when a specified event(s) or event rule(s) triggers it

Mail Template:

Criteria

Threshold:

Time Period:

Trigger On: Event(s) Event Rule

Select Event Properties

Add filter

- Indicator Set
- Subtype
 -
 -
 -
 -
 -
 -

Alerts make it easier to monitor specific events, and can be configured to send email to one or more recipients when triggered. See [“Using App Control Alerts”](#) on page 608 for instructions on creating and configuration alerts.

Responding to Threats

If you see a threat that requires remediation or further attention, there are many ways you can respond. A key step before taking action is to research the files, processes, users, and other information included in the report.

Once you determine that a response is required, you might take actions outside of App Control, such as deleting instances of suspicious files or creating new firewall rules. Within App Control, you can check the box next to events reported in threat views and act on the files reported in the events using the commands on the Action menu, including:

- **View Carbon Black Reputation Data** – If you have enabled Carbon Black File Reputation, you can open the Carbon Black File Reputation site to view additional reputation information about a file (if available), including its first seen date and prevalence on agent-monitored computers.
- **Send Suspicious Files for Analysis** – If you have used the App Control Connector to integrate an external analysis appliance or service, you can send files reported as threats for external analysis. Note that this option sends the target file noted in a threat event for analysis, not the process. See [Appendix C, “App Control Connector,”](#) for more information.
- **Ban Globally/Ban by Policy** – You can ban a suspicious or malicious file directly from the Action menu in one of the threat views, and you can configure policies to terminate running processes for banned files. Bans should be used carefully since it is possible that a file reported in a threat report is used for both acceptable and unacceptable purposes. One way to determine this is to begin with a Report Only ban, an option available on the Ban by Policy page. See [Chapter 8, “Approving and Banning Software,”](#) for more information.
- **Delete Files** – In addition to banning suspicious or malicious files, you have the option of deleting them through the App Control console, either on one computer or everywhere they appear on an agent-managed computer. See [Chapter 9, “Deleting Files,”](#) for more information.

The screenshot displays the 'Events' section of the App Control console. At the top, there are filters for 'Saved Views' (set to 'Threat Indicators') and 'Group By' (set to '(none) Ascending'). Below the filters are links for 'Show Filters', 'Show Columns', 'Export to CSV', 'Access Event Archives', and 'Refresh Table'. The main area shows a table of threat events. An 'Action' menu is open over the first row, listing various actions such as 'Approve', 'Remove Ban or Approval', 'Ban', and 'Upload file to server'. The table columns include 'Rule Name' and 'Details'.

Rule Name	Details
Processes started by Powershell remoting (WinRM)	Admin Tool Tracking
Powershell or WinRM remoting activity	Admin Tool Tracking
Suspicious OSX persistence	System Configuration
Suspicious shell use	Activity
Modification of the powershell execution policy	System Configuration
Shell Spawned from a browser	Application Behavior
Unusual change to startup configuration	Startup Configuration
Possible name resolution tampering	System Configuration
Possible name resolution tampering	System Configuration

In addition to the choices on the Action menu, there may be situations in which creation of a different type of rule, such as a custom or registry rule, could mitigate the threat. These

rules require that you enter their parameters manually. You can copy file, registry, or process information from events in the threat views and then configure the other rule parameters in the way you choose, being careful to restrict the rule to the actions you are certain you want to block or report on to avoid blocking critical files or processes. See [Chapter 14, “Custom Software Rules,”](#) [Chapter 15, “Registry Rules,”](#) and [Chapter 16, “Memory Rules,”](#) for more information about creating these rules.

Responding to Threats with Event Rules

Event Rules allow you to take certain actions when events matching the rule definition occur. This offers an automatic way to respond to threats, even if you haven't reviewed them on the Events page. Although you normally cannot automatically ban files using an Event Rule, you can take other actions that might be useful for reported threats:

- **Remove Approval from Suspicious Files** – You can use event rules to automatically remove local or global approval from files matching the rule parameters.
- **Send Suspicious Files for Analysis** – If you have used the App Control Connector to integrate an external analysis appliance or service, you can create an event rule that sends files reported as threats for external analysis. For example, in the illustration below, files reported in threat events and that are not already banned are sent to WildFire for analysis. This can provide more information that might influence your decision to block or not block a file.
- **Create a Ban** – You can define an event rule that bans any file included in a threat event triggered by an ATI. You can also create a Report Only ban that will generate an event telling you that a file would have been blocked if the Ban was fully activated.
- **Delete Files** – You can define an event rule that deletes any file included in a threat event triggered by an ATI.

Create Event Rule ?

General

Copy Settings From: (none) ▼

Rule Name: Analyze Detected Threat Files

Description: Send suspicious files detected by ATIs to Check Point.

Status: Enabled Simulate only Disabled

Select Event Properties

Add filter ▼

Indicator Name is not empty ▼ +

Subtype is ▼

Report access (memory rule) ▼

or Report execution (custom rule) ▼

or Report execution (removable media) ▼

or Report read (removable media) ▼

or Report write (custom rule) ▼

or Report write (removable media) ▼ +

Select File Properties

Add filter ▼

Global State is not ▼ Banned ▼

or Banned by Policy ▼ +

Select Process Properties

Add filter ▼

Select Action

Action: Analyze file ▼

Priority: Medium ▼

Use Check Point:

Check Point Submit Options: Win81-64bit;Office 2013;Adobe 11 ▼

See [Chapter 19, “Event Rules,”](#) for more about creating and editing event rules.

Chapter 24

Using and Customizing Dashboards

App Control Dashboards are configurable pages containing compact windows called “portlets,” each of which provides access to App Control-related information or controls.

Sections

Topic	Page
Dashboards Overview	681
Using Portlets	683
Changing Dashboard Appearance	688
Creating, Editing and Managing Dashboards	691
Managing the Default Home Page	697
Creating and Customizing Portlets	699

Dashboards Overview

If you have not changed the default start page, the Home Page dashboard is the first page shown when you log in to the console (if not, click **Home** in the console menu).

The screenshot displays the Carbon Black App Control Home Page dashboard. The top navigation bar includes the Carbon Black App Control logo, the instance ID 'APC-12709', and menu items for 'Home', 'Reports', 'Assets', 'Rules', and 'Tools'. Below the navigation bar is a 'Home Page' dropdown and a search bar. The dashboard is composed of several portlets:

- Alerts:** Shows a table with columns for Name, Type, Priority, Date Triggered, and Instances. A single alert is listed: 'New Certificate Alert' (Certificate Alert, Low priority, triggered on Jun 24 2020 02:12:14 PM, 8 instances).
- Top X:** A search interface to find top items, with options for 'Find top' (10), 'Blocks by Computer', and 'Max age' (Last Day).
- Find Computer:** A search interface to find computers, with options to search by 'Computer name or IP' or 'User name'.
- Find Files or Events:** A search interface to find files or events, with filters for 'Computer' (Any Computer), 'User' (Any User), 'Filename' (All Files), and 'Max age' (Last Day).
- Change Policy:** A form to change the policy of a computer, with fields for 'Change policy of computer', 'From existing policy', and 'To new policy'.
- Event Reports:** A summary of events for a specific period (Jun 24 2020 11:25:46 AM to Jun 25 2020 11:25:46 AM), showing counts for 'New installations', 'New unapproved files', 'Blocked files (Banned)', and 'Blocked files (Unapproved)'. A table shows 1 new installation and 1 computer.
- Emergency Lockdown:** A button to lock down computers, with a 'LOCK DOWN' button and a message: 'Click on the button below to move all your connected computers not under High Enforcement Level to High Enforcement Level.'
- Licensing:** A summary of license usage, showing counts for 'Visibility' (0) and 'Control' (100 limit, 1 in use).

A Dashboard consists of a series of *portlets*, each of which provides summary information or controls that can help you manage the security of your computers and the files on them. Some portlets display a specific type of information from your App Control database, such

as events or baseline drift. Others might display news feeds or other information from an outside URL.

Note

This chapter uses the Home Page as an example for explaining dashboard features. For a complete list and description of the Home Page portlets, see [Table 2, “Home Page Quick Access Portlets”](#) on page 58.


- The initial section of this chapter describes basic elements of a dashboard and how to use them. If you intend only to use the dashboards delivered with App Control, this is the only section you need to read.
- The second major section of the chapter describes customizing the *appearance* of a dashboard. If you plan to use only existing dashboards but would like to change some aspects of the way they are displayed, this section will help you accomplish that.
- The third major section of the chapter describes how to create and customize dashboards and the information and controls on them. This includes choosing to share a dashboard with other users.
- The final section of the chapter describes how to create and edit the portlets that make up a dashboard.

What you can do with dashboards depends on the privilege level of your console login account – the descriptions below assume default permissions for each group:

- Administrators and PowerUsers can view, use the features of, create, change, and delete their own dashboards and dashboards shared by other users. They can share dashboards they create, and they can choose a different default Home Page for new users of your console.
- Administrators and PowerUsers can view, use the features of, create, change, and delete portlets.
- ReadOnly users can access and use the features of their own dashboards, built-in dashboards such as the Home Page and System dashboard, and any dashboards other users have created and shared. They can create, change, or delete their own dashboards. They cannot modify or delete other dashboards, share dashboards they create, or choose a different default Home Page for new console users.
- ReadOnly users can view and use the features of portlets except for those that access features they do not have permission to use, such as Emergency Lockdown and Changing Policy for a Computer. They cannot create, modify, or delete portlets.
- You can enable or disable permissions for dashboard access by using the Manage Shared Dashboards checkbox on the Group Details page (see [“Managing Console User Roles”](#) on page 101).

Dashboard Elements

Although the portlets displayed by a dashboard vary, the basic structure of all dashboard pages is standard. The two main areas are the Dashboard toolbar, which shows the name of the current Dashboard and provides buttons and menus to manage it, and the portlets.



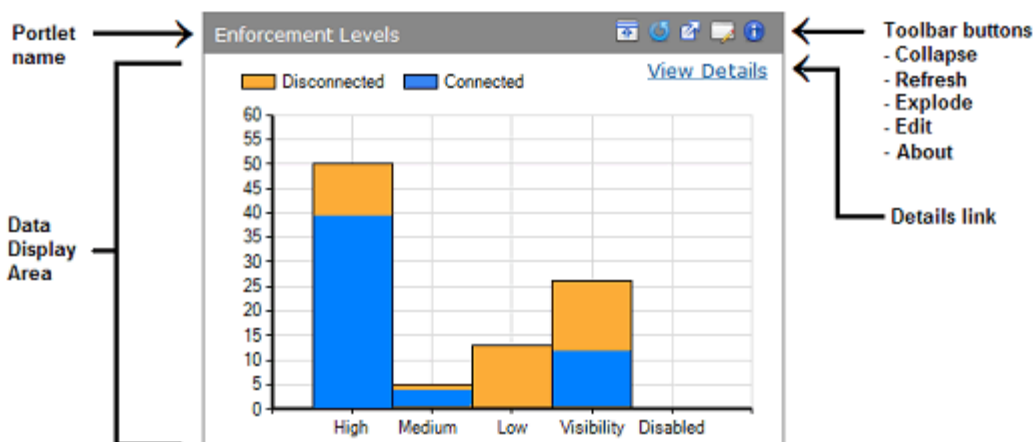
The dashboard toolbar includes:

- Current dashboard name – This appears at the top left of the toolbar.
- Dashboards menu – Clicking on the down-arrow next to the dashboard name opens the dashboards menu, which allows you to choose a different dashboard to display.
- Dashboard Help button – The question mark button in the upper right area of the dashboard page opens general help about dashboards. For each individual portlet, an information button in the upper right corner provides a description of that portlet.
- Dashboard action buttons – The Reload button reloads the current dashboard. The remainder of the buttons are used for more advanced activities described in the section “[Creating, Editing and Managing Dashboards](#)” on page 691.
- Dashboard appearance option menus – These options, on the right half of the toolbar, are described in detail in “[Changing Dashboard Appearance](#)” on page 688.

Using Portlets

The portlets on a dashboard may display file, computer, or event information. They might show the number and types of computers managed by App Control Agents, the number and type of security policies enforced, or the categories of software on your computers. The dashboard might also include portlets that allow you to make inquiries, such as finding an event or file, or portlets that take actions, such as locking down all computers.

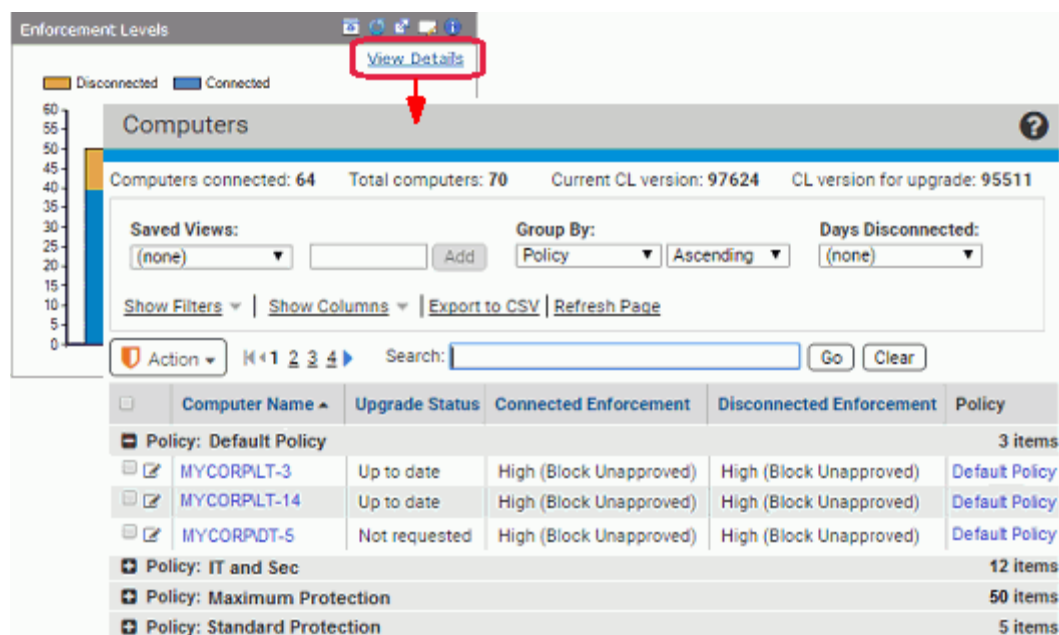
Each portlet has a toolbar with its name in the top left and a series of buttons in the top right. The main content of the portlet is below the toolbar. Data is displayed in this content area in the form of tables, charts, graphs, RSS crawls, or HTML pages. For portlets that take action or allow queries, there are fields to fill in or buttons to click to execute an action. You might also add portlets with other means of conveying data.



In many portlets, moving the mouse cursor over an element of a chart, for example, a bar in a bar chart, provides a description of that element, such as how many computers are represented by a particular bar in the chart.

Getting More Detailed Data

In addition to displaying key information at their top level, many portlets provide a way to “drill down” for more detail. You get more detail by clicking on graphics or data in a portlet (where the mouse cursor changes into a hand shape) and/or clicking on the **View Details** button, if it is available in the portlet. The first level of detail below the dashboard might be a App Control Server page with the additional information about what the portlet shows. Depending upon the portlet, information on the details page might be grouped by the data type shown in the portlet (e.g, computers grouped by Enforcement Level).









To return to a dashboard from a “drilldown” to details, choose the name of the dashboard you were on from the console Home menu. Note that using the back button to return to a dashboard could produce unpredictable results.

Portlet Toolbar Buttons

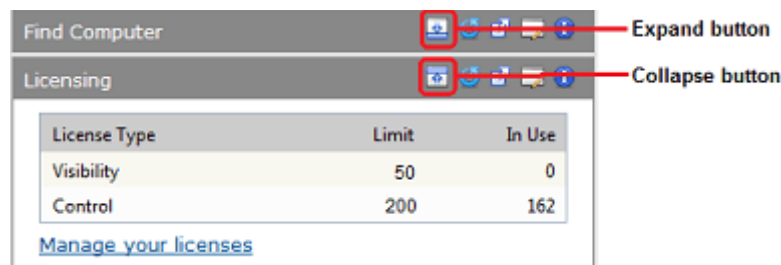
The portlet toolbar offers a variety of options, some of which change the display of a portlet. [Table 101](#) shows the buttons in the toolbar and the actions they take.

Table 101: Portlet Toolbar buttons

Button	Description
 Collapse	Collapse the view of the portlet so that only its toolbar is displayed.
 Expand	Restore a collapsed portlet to its normal display.
 Reload	Reload the portlet with the most current data available.
 Explode	Explode the view of the portlet so that it covers the entire dashboard. Clicking the X in the upper right corner of an exploded portlet restores it to its normal size.
 Edit	Open the Portlet Details page for this portlet, which provides access to editable parameters. What can be edited varies by portlet type and source. For some portlets built-in portlets, the only editable parameters are the name and the description that appears when a user clicks the information button. See “Editing Portlet Details” on page 700.
 Information	Open the information window for this portlet, which provides a brief description of the purpose of the portlet and how to use it. This information may be edited.

Collapsing, Expanding, and Exploding Portlets

There are two features for changing the way portlet windows are displayed on a dashboard. One allows you to “collapse” a portlet to display its name and toolbar only, and then to “expand” the portlet back to its normal state. The Collapse or Expand button (depending upon the current portlet state) is in the toolbar on the right side of each portlet.



Exploding a portlet is a temporary viewing option that allows you to take over the entire dashboard display area with one portlet. When you are finished with the exploded view, click the X button in the top right area of the portlet to return to normal viewing.


The size of an “exploded” portlet depends upon the size of the console browser window at the time the explode button was clicked.

Entering Information into Portlets


App Control is shipped with System portlets, not all of which are on the original Home Page. Some System portlets provide fields for entry of data, such as a computer name, a file name, or a user name, in order to conduct a search for information or to take an action on the item identified in the data. These portlets have several useful features.

Where you type in the name of something stored in the database for your App Control Server, a portlet provides an “auto-complete” feature – as you type, a list of possible

matches to what has been typed so far is displayed in a menu. If the item you are looking for appears in the menu, you can simply point and click it to finish entering the name. As the example below shows, auto-complete matches what you have typed with any object in the category you chose (*User* in the example) that *contains* the string, not just those that begin with it. Note, however, that you can choose an *Exact match* option for Filename rather than the default behavior of finding every file containing the entered string.



When you enter data into a portlet, the data you enter generally stays in the fields (i.e., becomes the default) unless you change it. This can be helpful if you want to do multiple searches (or other actions) with most but not all of the same information you first entered. To start over with no data on the portlet, click the **Clear** button.



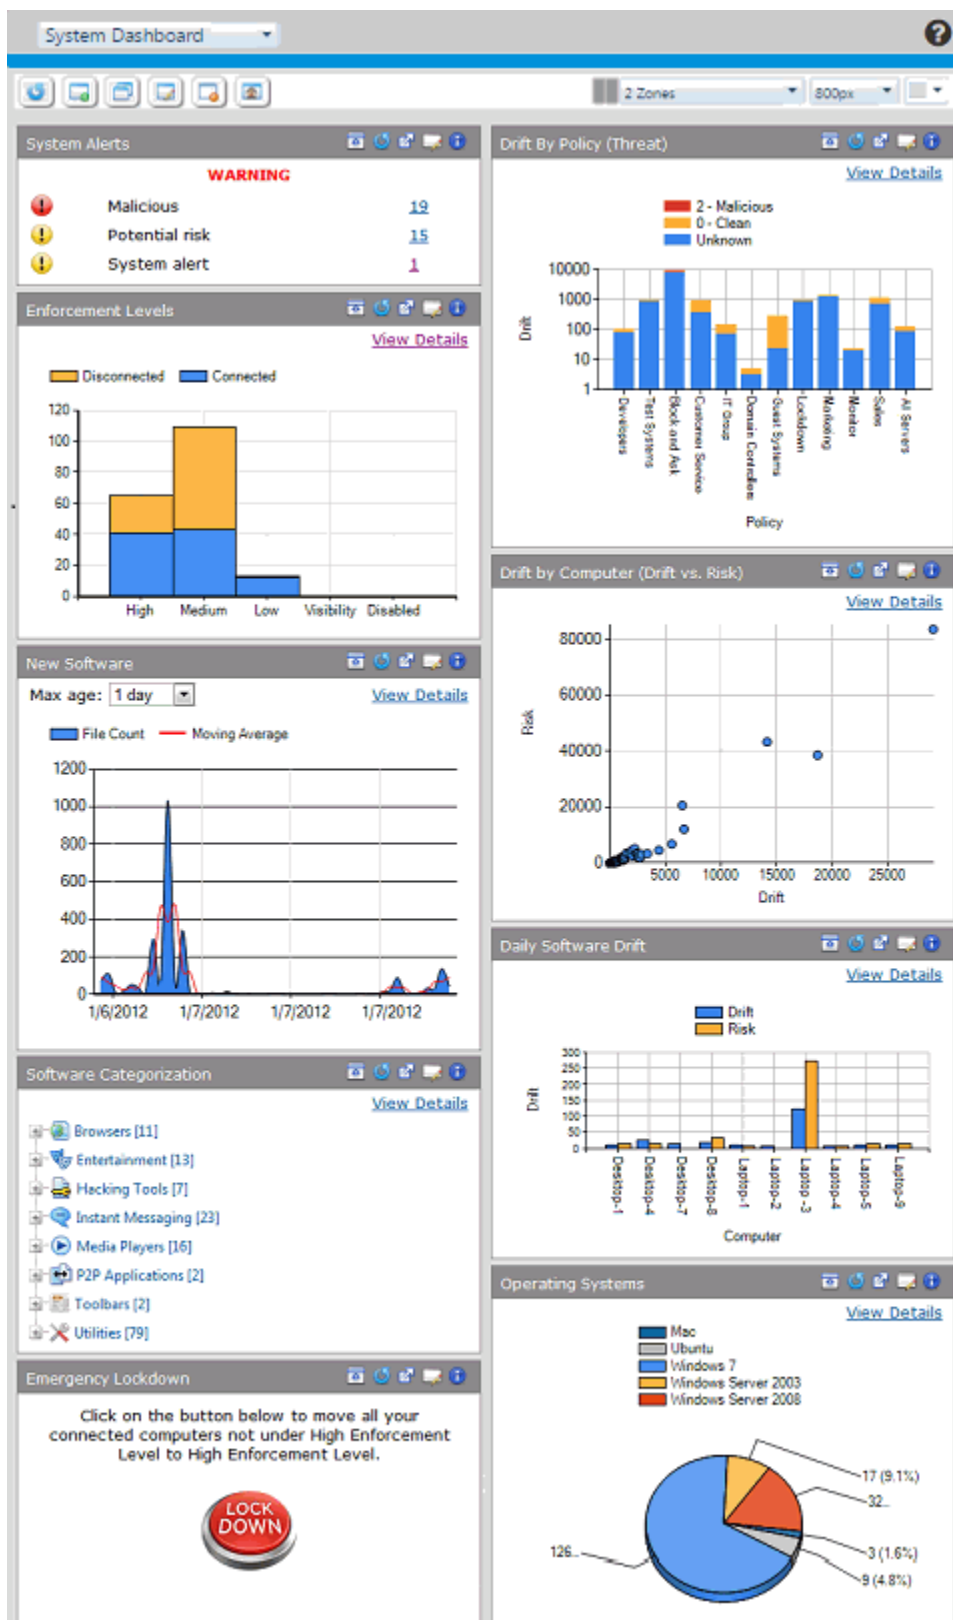
Other Portlet Controls

Portlets can have special controls that provide more information or take an action. For example, the Emergency Lockdown portlet has large buttons for Lockdown and Restore. The Alerts portlet has highlighted text links for resetting some or all links. Where there are special controls, text in the portlet itself describes their purpose.

Viewing Other Dashboards

The Home Page is always available on the console menu. There also is a *System* dashboard with portlets showing a variety of reports on your system, including the number of computers at each Enforcement Level, new software seen on your system, and baseline drift reports. Upgrades from a previous release may include other dashboards available created in the previous version.


The illustration below shows the type of portlets on the System dashboard (your System dashboard might have more, fewer, or different portlets).




There are several ways to choose and open a different dashboard.

To open a dashboard:

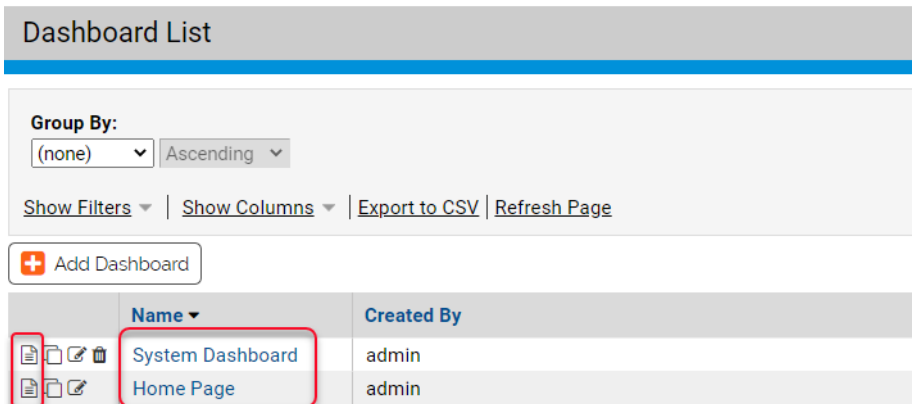
- If you are on a dashboard, choose a different one from the menu in the top left of the toolbar:



- Or, from any console page, move the cursor over **Home** in the console menu to view other dashboard choices. Note that not all dashboards are necessarily added to the menu.



- Or, choose **Reports > Dashboards** on the console menu and on the Dashboard List, either click on the View Dashboard button next to a dashboard name or click on the name itself.




Changing Dashboard Appearance

The following options can be used to change the appearance of a dashboard:

- changing the layout of portlets on the dashboard
- changing the dashboard width
- changing the dashboard background color
- collapsing and expanding portlet windows
- moving portlets on the dashboard

Three of these options are on the menus on the right half of the toolbar:




Note that this section describes what can be done to change the appearance and layout of an existing dashboard with existing portlets. Adding and removing portlets is described in the section [“Editing a Dashboard”](#) on page 695.

These appearance options affect only the current dashboard, and are specific to the currently logged in user.

Changing Dashboard Layout

The Dashboard Layout menu shows the current dashboard layout and allows you to select a different layout from a set of 13 templates. The templates create *zones* in which portlets are placed, and in some layouts, these zones have different widths. Once you choose a layout, you can move portlets from zone to zone so they have width appropriate for their content.



Layouts are labeled with the number of zones and the “style” number if there is more than one style with that number of zones. The default layout is two equal columns, which is the only “2 Zones” layout. The number of zones is not the number of portlets – each zone can and usually will have multiple portlets in it.

Portlet Distribution in Layouts

When you switch between layouts or add portlets, portlets are assigned to zones based on the following rules:

- If you switch to a layout with the same or more zones as your current one, portlets will remain in their assigned zone. For example, if you switch from “2 Zones” to “3 Zones, Style 1,” all of the portlets in zone 1 will remain in zone 1 and all of the portlets in zone 2 will remain in zone 2 until you move them. There is no attempt to map portlets that are in wide zones in one layout to wide zones in a different layout.
- If you switch to a layout that has fewer zones than the current one, portlets will be remapped to new zones. Portlets from even-numbered zones in the former layout will go to even zones in the new, and odd to odd, except when going to the one-zone layout, where all portlets go to the single zone.
- When you add portlets to a dashboard, they are distributed sequentially to each zone, starting with zone one. So if you add three portlets during one editing session, one each goes to zones 1, 2, and 3.
- The console “remembers” the distribution of portlets in layouts you have used. If you change layout and then return to one you used previously, the portlets appear in the same locations they did before as long as you have not added or removed portlets.


In many cases, you will want to rearrange portlets after a layout change.

Changing Dashboard Width

The Dashboard Width menu shows the current dashboard width in pixels and allows you to select a width between 600 and 1700 pixels. When you change dashboard width, the width of portlets is resized proportional to their zone within the current layout. Choose a width appropriate to your screen size and resolution, and to the amount of the screen you want to allocate to the console. The default dashboard width is 800 pixels.


Changing Dashboard Background Color

On the Dashboard Color menu, you can change the background color of a dashboard. Clicking on the menu brings up a palette, and clicking a color on the palette makes the color change. The background color change does not affect portlet color. The default background color is light gray.



Moving Portlets

You move a portlet by clicking in its toolbar and moving the mouse while holding the left mouse button down. When you move a portlet, the portlet you are moving becomes transparent, and only the borders of the other portlets are shown. As you move the portlet, the location in which it would be dropped if you released the mouse button is shown as a dotted-line box, a landing area. If you move from one layout zone into another, the landing area box shows you any change in portlet width due to the move. When you drop the portlet into its new location, all of the portlets return to normal display.



Creating, Editing and Managing Dashboards

This section describes the creation and editing of dashboards as well as other dashboard management tasks. Dashboards are defined by the following basic parameters:

- name
- portlets you want on the dashboard
- whether this dashboard will be shared with other users
- whether this dashboard will be listed on the console menu

You can create a new dashboard from scratch or copy an existing dashboard to a new name, modifying it once copied. Whether you are creating, copying, or editing a dashboard, you enter or edit the basic configuration information on the Edit Dashboard page. The main difference among these cases is what information, if any, is filled in for you on the Edit Dashboard page when you start.

In addition to creating and editing dashboards, you might want to:








- set or reload the default dashboard, which is described in [“Managing the Default Home Page”](#) on page 697
- delete dashboards, described in [“Deleting a Dashboard”](#) on page 697

Note

This section describes how you define and manage a dashboard and its *content*. Ways to customize the *appearance* of a dashboard are described in the section [“Changing Dashboard Appearance”](#) on page 688.

You can access most of the dashboard management tasks described here from either the Dashboards list page or from the toolbar on an individual dashboard. See [“Managing Dashboards from the Dashboards Page”](#) on page 698 for a summary of Dashboards list page features. [Table 102](#) shows the actions taken by the buttons on the dashboard toolbar.

Table 102: Dashboard Toolbar buttons

Button	Description
 Reload	Reloads the dashboard and its portlets with the most current data available.
 New Dashboard	Opens the Edit Dashboard page, where you can enter a name for a new dashboard and choose whether to make it available to other users and whether to show it on the console menu (under <i>Home</i>). You also choose portlets for the dashboard from this page, and can create new portlets using the New Portlet button.
 Copy Dashboard	Opens the Edit Dashboard page for the current dashboard, with all of the current portlets checked for inclusion and a new dashboard name in the form “Copy of <the dashboard you were on>”. You can modify the name as you choose. Saving a copy of a dashboard can be useful if you want to have your own version of a shared dashboard, or if an existing dashboard has some of the portlets you would like to use but you want to add or remove portlets to make it exactly what you need. This also gives you options to add the dashboard to the console menu and share it with all users.
 Edit Dashboard	Opens the Edit Dashboard page so you can modify the current dashboard, including creating new portlets or changing the portlets displayed.
 Delete Dashboard	Deletes the current dashboard (after you choose OK in a confirmation box). See “Deleting a Dashboard” on page 697. Not available on the Home Page.
 Reset to Default	Resets a system-provided dashboard (currently, the Home Page and System dashboard) to its currently saved default settings (see Set as Default below). Not available for user-created dashboards.
 Set as Default	Sets the current dashboard as the default Home Page for users whose accounts are created <i>after</i> this setting is saved. See “Managing the Default Home Page” on page 697.

Shared Dashboards


You can create dashboards strictly for your own use only, or you can share any dashboard you create by checking the *Share with all users* box on the Edit Dashboard page.

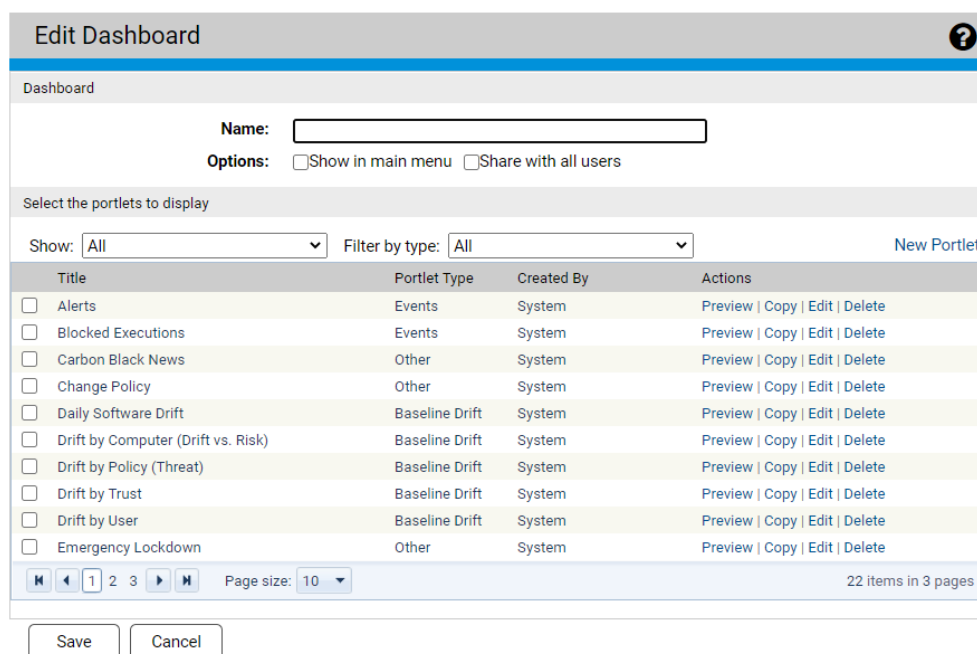
When dashboards are shared, console users in Administrator or PowerUser groups, or in custom groups with *Manage Shared Dashboards* permission, can modify the dashboard, and they also can delete it.

Keep in mind that other users might come to rely on a dashboard you share. If you turn off sharing for a dashboard or delete the dashboard, other users will lose access to it, either immediately, or, if they are on the dashboard, as soon as they navigate away from it.

Creating a New Dashboard

To create a new dashboard:

1. Open the Edit Dashboard page for a new dashboard using one of the following:
 - Choose **Reports > Dashboards** on the console menu, and on the Dashboards page, click the **Add Dashboard** button.
 - **- or -**
 - On any dashboard, click the Create New Dashboard button .



Edit Dashboard ?

Dashboard

Name:

Options: Show in main menu Share with all users

Select the portlets to display

Show: Filter by type: New Portlet

Title	Portlet Type	Created By	Actions
<input type="checkbox"/> Alerts	Events	System	Preview Copy Edit Delete
<input type="checkbox"/> Blocked Executions	Events	System	Preview Copy Edit Delete
<input type="checkbox"/> Carbon Black News	Other	System	Preview Copy Edit Delete
<input type="checkbox"/> Change Policy	Other	System	Preview Copy Edit Delete
<input type="checkbox"/> Daily Software Drift	Baseline Drift	System	Preview Copy Edit Delete
<input type="checkbox"/> Drift by Computer (Drift vs. Risk)	Baseline Drift	System	Preview Copy Edit Delete
<input type="checkbox"/> Drift by Policy (Threat)	Baseline Drift	System	Preview Copy Edit Delete
<input type="checkbox"/> Drift by Trust	Baseline Drift	System	Preview Copy Edit Delete
<input type="checkbox"/> Drift by User	Baseline Drift	System	Preview Copy Edit Delete
<input type="checkbox"/> Emergency Lockdown	Other	System	Preview Copy Edit Delete

Page size: 10 22 items in 3 pages

2. In the Name box, enter the name you want for the new dashboard. This is the name that will appear in the upper left when you display this dashboard, and is also the name that will appear on the list of dashboards on the Dashboards page.
3. If you would like to add this dashboard to the Home section of the console menu:
 - a. In the Options line, check the *Show in main menu* box. Note that even if you do not check this box, the dashboard will be available through the Dashboards page and on the Dashboards menu of any other dashboard.
 - b. If you want a different (usually shorter) name to appear on the menu than the one you chose for the dashboard, enter it in the Menu name field, which appears when you check the Show box.
4. If you want other users to be able to use this dashboard, check *Share with all users*.
5. Check the box to the left of each portlet you want to add to this dashboard. Use the page buttons at the bottom of the portlet list or the filters at the top of the list to view all of the available portlets of interest.

Note: To see what the portlet looks like before adding it to the dashboard, click **Preview** to the right of the portlet name.
6. If you need a portlet not available on the list, see [“Creating and Customizing Portlets”](#) on page 699. Once the new portlet is created, check the box next to its name to add it to this dashboard.



7. Click **Save**. The new dashboard is saved and added to the list on the Dashboards page. If you checked the appropriate box, its name appears on the Home menu on the console menu.

Copying a Dashboard

Copying a dashboard can be useful under a number of circumstances, including:

- if you want your own copy of a shared dashboard created by someone else
- if you find a dashboard that is close to what you want but would like to add or remove portlets or otherwise edit it for your needs

To save an existing dashboard under another name:

1. Open the Edit Dashboard page for a copied dashboard using one of the following:
 - Choose **Reports > Dashboards** on the console menu, and on the Dashboards page, click the  button next to the dashboard you want to copy.
 - **- or -**
 - On the dashboard you want to copy, click the Copy Dashboard button .
2. The Edit Dashboard page opens with all of the same parameters as the dashboard you copied, except for the name, which appears in the form “Copy of <name-of-dashboard-you-copied>”. Replace the default “Copy of” name with the name you want to use for the dashboard.
3. Modify any of the other dashboard parameters you would like to change. See [“Creating a New Dashboard”](#) on page 694 for details.
4. Delete any portlets you do not want to appear on this dashboard by un-checking the box to the left of their names.

Caution

Do not click the Delete link to the right of the portlet name – this deletes it from the App Control Server entirely, not just from the current dashboard.


5. Add any portlets you would like to appear on this dashboard by checking the box to the left of their names.
6. If you need a new type of portlet, see [“Creating and Customizing Portlets”](#) on page 699. Once the new portlet is created, check the box next to its name to add it to this dashboard.
7. Click **Save**.
The copied dashboard appears on the Dashboards page under its new name with whatever modifications you made.

Editing a Dashboard

You can edit a dashboard to add or remove portlets from it, change its name, or change its sharing and menu options.

To edit a dashboard:

1. Display the dashboard you want to edit.
2. Click the *Edit this dashboard* button in the dashboard toolbar. The Edit Dashboard page appears.
3. Modify any of the dashboard parameters you would like to change, including:
 - a. Portlet name
 - b. *Show in main menu* choice
 - c. *Menu name* (if the *Show in main menu* box is checked)
 - d. *Share with all users* choice
4. On the Edit Dashboard page, the portlet list includes all portlets, including those already on the current dashboard. There are several options for filtering the list:
 - a. If you want to see a list of only those portlets *not* currently on this dashboard, on the *Show* menu choose **Portlets not on the dashboard**.





- b. To see only certain *types* of portlets in the list, choose the type on the *Filter by type* menu; for example, you might choose to show only Computer portlets. See [“Portlet Types and Subtypes”](#) on page 699 for a description of portlet types. You can combine choices on the Show menu with choices on the Filter menu. Also, these menu choices affect what appears on the Edit Dashboard page, not what appears on the dashboard.
 - c. Whether the list is complete or filtered, if it includes multiple pages, you can click the page numbers or arrows at the bottom of the list to navigate from page to page. The legend in the bottom right corner of the list tells you how many items and how many pages are in the current list.
5. You can use the **Preview** button next to any portlet in the list to see what it will look like on the dashboard.
6. Check the box to the left of the name of each portlet you want to add to the dashboard. See [“Creating and Customizing Portlets”](#) on page 699 if you need to create a portlet not currently found in the list.
7. Un-check the box next to the name of each portlet you want removed from the dashboard.

Note: Do not click the Delete link to the right of the portlet name – this deletes it from the App Control Server entirely, not only from the current dashboard.
8. When you have checked all the portlets you would like to add, click the **Save** button. The dashboard is redisplayed with the new portlets added.
9. If you need to change the overall dashboard layout to accommodate the new portlets, use the Dashboard Layout menu to make this change. See [“Changing Dashboard Layout”](#) on page 689 for more details.
10. If necessary, move portlets on the dashboard to accommodate the new portlets. If you do not know how to move portlets, see [“Moving Portlets”](#) on page 691.

Managing the Default Home Page

There are two Home Page management buttons on the dashboard:

- Using the Reset to Default button , any user can choose to reset their current, possibly modified, Home Page, to the default Home Page.
- Using the Set as Default button , any user with Administrator or PowerUser privileges (or custom Manage Shared Dashboards permission) can save the current dashboard as the default Home Page for new users.

If you set a different default Home Page, that page becomes the Home Page for anyone using the Reset to Default button. It also is the default Home Page for any new console users who log in for the first time *after* the change to the default. Users who have already logged in before the default Home Page is changed retain their existing Home Page unless they click the Reset to Default button and have permission to make the change.

Note


To be certain you can go back to the original Home Page, before you (or anyone else) make any modifications, you can use the Copy Dashboard command to copy the Home Page, and rename the copy so that you will have a backup. If needed, you can use Set as Default to restore the Home Page from the backup.

Deleting a Dashboard

You can delete any dashboard you created and (unless you are logged in as a ReadOnly user) any shared dashboard made available to you. The only dashboard that cannot be deleted by anyone is the Home Page.

When you choose to delete a shared dashboard, a dialog box warns that the dashboard is shared and allows you to confirm or cancel the deletion. Be careful when deleting a shared dashboard since it is possible that other console users want to continue using it. If another user is using a dashboard *when* you delete it, the dashboard remains displayed until they navigate away from it, at which point it becomes unavailable

To delete a dashboard:

1. Start the deletion process in one of the following ways:
 - On the console menu, choose **Reports > Dashboards** and on the Dashboards page, click the Delete (x) button next to the name of the dashboard to delete.
 - **- or -**
 - On the dashboard you want to delete, click the Delete Dashboard  button.
2. In the confirmation dialog that appears, if you are certain you want to delete this dashboard, click **Yes**. The dashboard is deleted and if you were on the dashboard when you deleted it, it is replaced by the Home Page.

Managing Dashboards from the Dashboards Page

The Dashboards page includes a complete list of available dashboards and controls to manage them. Many of the procedures described in other sections of this chapter reference the Dashboards page for alternative ways to accomplish a task.

Dashboard List

Group By:
(none) ▾ Ascending ▾

[Show Filters ▾](#) | [Show Columns ▾](#) | [Export to CSV](#) | [Refresh Page](#)

+ Add Dashboard

	Name ▾	Created By
	System Dashboard	admin
	Jane Doe Dashboard	admin
	Home Page	admin

Table 103 shows the dashboard-specific actions available on this page – see also Table 102 for similar commands available when you are already on a dashboard:

Table 103: Dashboard List buttons and links

Button/Link	Description
Add Dashboard	Opens the Edit Dashboard page, where you can enter data for creating and configuring a new dashboard. See “Creating a New Dashboard” on page 694 for more details.
View Dashboard	Clicking this button displays the dashboard in this row. See “Dashboards Overview” on page 681 for an overview.
Copy Dashboard	Copies the portlets and other settings for the current dashboard to a new dashboard named “Copy of <current-dashboard>”, and opens the Edit Dashboard page. You can modify the name as you choose. Saving a copy of a dashboard can be useful if you want to have your own version of a shared dashboard, or if an existing dashboard looks like a good template. See “Copying a Dashboard” on page 695 for more details.
Edit Dashboard	Opens the Edit Dashboard page for the dashboard in this row so you can modify the dashboard, including creating new portlets or changing the portlets displayed. “Editing a Dashboard” on page 695 for more details
Delete Dashboard	Deletes the dashboard in this row (after you choose OK in a confirmation box). See “Deleting a Dashboard” on page 697 for more information. Not available on the Home Page.
Dashboard Name link	Clicking a dashboard name in the list displays the dashboard.

Creating and Customizing Portlets

In addition to its dashboard management features, the Edit Dashboard page provides access to portlet management features with which you can:

- edit an existing portlet
- create a new portlet
- copy an existing portlet and modify it
- delete a portlet

Any user with Administrator or PowerUser privileges, or in a custom group with dashboard management permission, can use these features. All changes to portlets, including creation and deletion, affect all console users – there are no “private” portlets.

Portlet Types and Subtypes

Portlets are organized by *types* and *subtypes*. Depending upon the type and subtype, the portlet has different capabilities, and there are different input parameters available when you create or edit it. The types are:

- **Events:** These portlets display event information from the App Control database, such as the number of blocked file executions over a period of time or alerts that have been triggered.
- **Baseline Drift:** These portlets display the results of baseline drift analysis, such as daily drift of software from a baseline or a list of the computers with the greatest deviation from the baseline.
- **Computers:** These portlets display information available from the App Control Server about the computers on your system, such as the number of computers running each operating system or the number of computers at each Enforcement Level.
- **Files:** These portlets show information about the files on agent-managed computers, such as the number of newly seen files over time or the category (browsers, utilities, messaging, etc.) of the files on the system.
- **Other:** These portlets may display an RSS feed or information from another URL, or they may display HTML pages you provide. This category also includes one-of-a-kind system-created “action” portlets such as the emergency lockdown button, or combinations of different types of information from your App Control database.

System Portlets

The console is installed with a large number of pre-configured portlets. Some of these are visible on the Home Page and might also be on other dashboards at your site. They can be identified by the name “System” in the “Created By” column on the Edit Dashboard page.


Some System portlets, such as the Emergency Lockdown portlet or the Change Policy portlet, are designed to be one-of-a-kind, and cannot be copied or deleted (the Copy and Edit links will be grayed out in their rows). The only changes allowed for these portlets are to their names and descriptions.

Editing Portlet Details

You can edit portlets to change their appearance or the data presented. You might, for example, decide that a pie chart better presents the data you want to see than a vertical bar chart. The Portlet Details page, where you edit portlets, can be opened from a currently displayed portlet on a dashboard or from the portlet list on the Edit Dashboard page.

See “[Creating Custom Portlets](#)” on page 701 for more detail on the individual parameters you can edit.

To edit a portlet on the currently displayed dashboard:

1. Click on the Edit button  in the upper right of the portlet you want to edit. The Portlet Details page appears.
2. Make whatever changes you want to the settings on the Portlet Details page. If necessary, click the **Show Advanced Details** button for more editing options.
3. Use the **Preview** link at the bottom of the page to view the effects of your changes. Note that you might need to scroll the browser window down to see the Preview panel. When a preview is showing, you can continue to make changes and click **Refresh** to see the results. Click **Close** when you are finished with the preview.
4. When you are satisfied with the changes you have made, click **Save** at the bottom of the Portlet Details page. The current dashboard appears and shows the portlet with whatever changes you made.

You also can edit portlets via the Portlet Catalog, whether or not the portlet appears on any of your dashboards.

To edit any portlet from the Edit Dashboard table:

1. On the Edit Dashboard page, find the portlet you want to edit.
2. In the list of portlets, click the **Edit** link to the right of the name of the portlet you want to edit. The Portlet Details page appears.
3. Edit as described in the previous procedure.

Deleting Portlets

Caution

Console users in the Administrators group or custom groups with permission to manage dashboards can delete portlets from the Edit Dashboard page (except for certain System portlets). Use this capability with care, since it deletes the portlet from *all dashboards for all users*.

To permanently delete a portlet:

1. From any dashboard or the Dashboards page, click the Edit Dashboard button. The Edit Dashboard page appears.
2. In the list of portlets, click **Delete** next to the portlet you want to delete. A confirmation dialog appears and includes information about how many dashboards use this portlet. Be sure you actually want to delete this portlet from your App Control environment – it will be permanently removed for all users.
3. If you are certain you want to delete this portlet, click **OK** in the confirmation dialog. The portlet is removed from the portlet list on the Edit Dashboard page. It is removed from all dashboards that include it.

If a user is viewing a dashboard containing the portlet, the portlet will remain visible until the user reloads or navigates away from the dashboard.

Creating Custom Portlets

In addition to making available built-in portlets, dashboards provide the means to create and use your own portlets. You can choose from a list of several portlet types that can present data about your App Control-managed assets and rules, and then configure the appearance of data from those reports as you choose.

Regardless of who creates a custom portlet, the portlet is available to all console users through the Edit Dashboard page. Note, however, that ReadOnly users cannot create or modify a portlet.

As you enter details for your portlet, don't hesitate to experiment with different settings on the Portlet Details page and click the **Preview** button. The Preview capability serves as both a debugger, to inform you when you choose incompatible settings for a portlet, and a good way to try different charts or different collections of data before adding a custom portlet to a dashboard.

To create a custom portlet:

1. Click the Edit Dashboard button, either on a currently displayed dashboard or next to the name of any dashboard on the Dashboards list.
2. On the Edit Dashboard page, click **New Portlet**. The New Portlet page appears.
3. On the New Portlet page, choose the type from the *Select portlet type* menu. See [“Portlet Types and Subtypes”](#) on page 699 for a description of the portlet types.
4. If there is more than one choice, choose the subtype from the *Select subtype* menu.
5. Click **Next**. The Portlet Details page appears. This is the same Portlet Details page that appears when you edit a portlet.

Note

The type and subtype of a portlet determine its fundamental structure and many of the available choices on the Portlet Details page. They cannot be edited once chosen. If you want to change type or subtype during the portlet creation process, click **Cancel** and start over.

Adding Portlet Details

6. On the Portlet Details page, enter the General details, which include the following:
 - a. **Title:** Type the title you want to appear on the portlet and in the portlet list on the Edit Dashboard page.
 - b. **Description:** Type the information you want users to see when they click the information button for this portlet, such as a short description of the purpose of the portlet and instructions for how to use it.
7. If the Portlet Details page includes a panel specific to your portlet type, such as *Baseline Drift details* or *RSS details*, fill in the required information there and then click **Next**.

If there is a **Save** link instead of a Next link, click it to save the new portlet and add it to the catalog and current dashboard. For some portlet types, no further configuration is necessary.
8. If a Data Presentation panel appears, you have the option of choosing Table as the Chart type.
 - If you choose **Table**, select the columns and column *order* you want, then continue with step 14. See “[Using Tables in Portlets](#)” on page 705 for details on configuring table portlets.
 - If you choose any other Data Presentation type, continue with step 9.
9. If a Graph Settings panel appears on the page, provide the details for the way in which you want the data for this portlet presented. The available choices vary depending upon the type and subtype of portlet, but generally those shown in [Table 104](#).
10. When you finish choosing Graphic Settings, click **Preview** to see what your portlet will look like. You can try a variety of settings, such as different chart types, to find the one you like best. Use **Refresh** to update the preview as you change settings.
11. Once you have specified the basic appearance of the chart for this portlet, you can do one of two things:
 - a. If you do not want to view and modify advanced graphic details, click **Save** to add the portlet to the Edit Dashboard page.
 - b. If you do want to see additional graphic settings, click the **Show Advanced Settings** button.
12. If you are reviewing advanced graphic settings, you have the choices shown in [Table 105](#). Note that not all advanced settings are appropriate (or available) for all chart types.
13. If you have entered Advanced details, you can click the **Preview** link again to examine your portlet before saving.
14. If the Portlet Details page for the portlet you are creating has a filters panel and you want to filter the data that will be used in the portlet (both graphic and table-only portlets), configure the filter you want. See “[Using Filters in Portlets](#)” on page 708 for more details.
15. When you are satisfied with the appearance and data of your portlet, click **Save** to add the portlet to the Portlet Catalog, add it to the current dashboard, and close the Portlet Editor.

Table 104: Portlet Graphics Settings

Setting	Description
Chart type	This menu lists the ways you can represent data for the portlet type and subtype you chose. The list may include points, bars, and pie charts, among other choices.
X-axis	This lists the types of attributes available for the portlet type and/or subtype you chose. Choose one (for example, Computer name) to distribute along the X axis of the chart. For different types of charts, the choice here might not determine what appears on the X axis but what is the fundamental data in another format, for example, what each slice of a pie represents.
Limit to the 5 10 15 highest lowest values	If you put certain data, such as individual computers, on the X-axis, you can have too many instances to display effectively inside the portlet. The “Limit to” checkbox and menus allow you to show only the instances with the 5, 10, or 15 highest or lowest values of whatever it is you are displaying (drift, for example). Presumably these would present the most interesting information, and the limit allows you to have a usable graphic rather than putting too much information into too little space. This box is not displayed for certain chart types, including scatter charts or columns using the “auto split” feature.
Group by	Appears only if you choose Scatter as the Chart type. If you choose a <i>Group by</i> value, the dots on the scatter chart represent the total value for the group you indicate rather than values for an individual group member. For example, if you choose Policy as the <i>Group by</i> value, instead of dots representing a Y value for individual computers, they would represent the Y value for all the computers in a policy instead.
Exclude “Unknown” X-axis values	If you check this box, data with unknown X-axis values is eliminated from the chart or graph. This is another way to eliminate less useful information from the portlet.
Split by	Specifies the information type whose values split the X-axis data. For example, you might create a portlet that shows raw drift by policy. <i>Split by</i> creates a separate series (bar, column, or segment) for each unique value in selected column, so a bar representing all the computers in a policy can be split (by color) to show how much drift is attributable to each computer.

Setting	Description
Metrics	Lists the choices of attributes you can represent on the Y-axis of your chart. If you can only choose one value for the particular portlet type you are creating, this is a dropdown menu. If you can choose multiple types, this is a multi-select menu that allows you to move more than one item from the <i>Available</i> columns to the <i>Selected</i> column or vice versa. You can add any metrics that are shown as available. For example, for a bar chart of unique files by global state, you could add “Count” to show the number of files in each state and then also add “Prevalence” to show how many computers have files of each type.
Show table below graph	When checked, displays a list of table columns available for this portlet. Move those columns you want displayed into the Selected column. See “Using Tables in Portlets” on page 705 for more details.

Table 105: Portlet Advanced Settings

Setting	Description
Height	Allows you to choose a height, in pixels, for the portlet, or to let the dashboard size it for you (Auto). Note that if you choose a value other than <i>Auto</i> , you may interfere with proper display of the portlet.
Show X axis title/ Show axis titles	When box is checked, includes the X-axis title (that is, the title shown in the X-axis box in Graph Details) on the portlet chart, or if X and Y axes are shown, titles for each.
X-axis labels	For choices other than None, adds labels to the data points on the chart (for example, the bars in a bar chart), in the location and orientation you choose. If you choose Auto, the dashboard specifies label positioning based on the best fit.
Legend	When any button but None is clicked, provides a legend describing the chart elements in the location you specify. For example, if different colors are used for total systems vs. connected systems, the legend identifies which is which.
Include tooltips	(Alternative to Legend) When this box is available and checked, hovering the mouse cursor over a chart element displays a tooltip describing what the element represents.
Show Data Point Values	When box is checked, displays the Y values (or their equivalent) on the portlet chart. For example, if a column represents three computers, the number 3 is displayed above the column.
Draw 3D	When box is checked, displays the chart with 3D effects.
Use logarithmic scale	When box is checked, changes the scale for displayed data from linear to logarithmic.

Using Tables in Portlets

When portlets have content appropriate for display in a table, there are two table options that can appear on the Portlet Details page:

- **Table Only:** The Portlet Details page provides a table option in the *Chart type* menu. This is the option to choose if you do not want any graphic charts on the portlet.
- **Supplemental Table:** If the main chart type choice is something other than *Table*, a *Show table below graph* checkbox appears at the bottom of the Graph Settings panel. When you check this box, you get both a graphic representation and a table.

Table-only Portlets

Table-only portlets can be a good choice when you would like to display App Control data on the dashboard that doesn't lend itself to graphic representation. For example, you might not be interested in *how many* computers or files meet certain criteria but instead in a more complex picture of different kinds of data for each computer, or for each file.

When table-only presentation is possible, a Data Presentation panel appears on the Portlet Details page. In that panel, you can choose **Table** as the Chart type. Choosing this option replaces the Graph Settings panel on the Portlet Details page with a Table Settings panel in which you choose and order the data to include in the table.

The screenshot shows the 'Portlet Details' page for a 'Baseline Drift' portlet. The 'General' section includes a title 'Special Drift Table' and a description: 'Shows the per-computer file drift in the inventory of unapproved files in the past 24 hours, and the risk associated with that drift.' The 'System Portlet' checkbox is unchecked. The 'Data Presentation' section shows 'Drift Report' set to 'Daily drift of all computers' and 'Chart Type' set to 'Table'. The 'Table Settings' section is highlighted with a red box and contains a 'Columns' panel with two columns: 'Available' and 'Selected'. The 'Available' column lists: First Created, File Name, Path Name, Drift, Risk, Relative Drift, Relative Risk, and Company. The 'Selected' column lists: Drift Report Result Id, Weighted Drift, and Drift Type. At the bottom of the 'Table Settings' panel are 'Refresh', 'Close', and 'Preview' buttons. The main content area below the 'Table Settings' panel displays 'No data available.' and 'Save' and 'Cancel' buttons.

You must choose the columns you want to appear in the table. You can double-click on a data element in the Available column to move it to the Selected column, and vice versa. You also can use the arrow buttons to move items back and forth between Available and Selected, and to change the order of data in the table.

Table portlets provide many features for rearranging the data they display:

- You can have multi-page tables and navigate between pages using the page and arrow buttons in the bottom left of the portlet.
- You can determine the number of rows displayed in a table by choosing a different *Page size* (in multiples of 10 rows).
- You can click over a column and drag it to a different location in the table.
- You can click over a column heading and drag it into the labeled zone at the top of the portlet to group the table by the data named in the column heading.
- You can filter the contents of a table by any column head to show data of interest. (You also can pre-filter the data using the Filters on the Portlet Details page.)
- You can click on a column head to sort by the data in that column.

Computer Status

Drag a column header and drop it here to group by that column

Computer Name	Parity Agent Version	Connected	Policy
Mycorp\Desktop-1	8.0.0.2322	True	Domain Controllers
Mycorp\Desktop-4	8.0.0.2322	True	Research Group
Mycorp\Desktop-6	8.0.0.2322	True	Sales
Mycorp\Desktop-7	8.0.0.2322	True	Research Group
Mycorp\Laptop-2	8.0.0.1345	False	Executives
Mycorp\Laptop-3	8.0.0.1345	False	Research Group
Mycorp\Laptop-4	8.0.0.2322	True	Marketing
Mycorp\Laptop-9	8.0.0.2322	True	Customer Service
Mycorp\Laptop-10	8.0.0.1345	False	Sales
Mycorp\Laptop-11	8.0.0.2322	False	Research Group

Page size: 10 167 items in 17 pages

To filter on a column, enter a string in the box below the column – for example, “Laptop” in the Computer Name column, and then click on the filter button to see the operator menu, where you can choose *how* you want to use the string you entered to filter the data.

Computer Status

Drag a column header and drop it here to group by that column

Computer Name	Parity Agent Version	Connected	Policy
Mycorp\Desktop-		True	Domain Controllers

Filter dropdown menu options: DoesNotContain, StartsWith, EndsWith, EqualTo, NotEqualTo

Supplemental Tables in Portlets

You can add a supplemental table within a graphic portlet. Because the space is shared, you probably will not want to create elaborate supplemental tables.

When a supplemental table is possible, a *Show table below graph* checkbox appears at the bottom of the Graph Settings panel. Check this box to display the Table Settings panel.

The screenshot shows the 'Data Presentation' settings panel. It is divided into two main sections: 'Graph Settings' and 'Table Settings'.

Graph Settings:

- Chart Type:** Column
- X-Axis:** Publisher or Company
- Limit to the 5 highest values
- Exclude 'Unknown' X-axis values
- Split By:** None
- Metrics:**
 - Available:** Prevalence, Prevalence (avg)
 - Selected:** Count
- Show table below graph

Table Settings:

- Columns:**
 - Available:** Category, Installed Program, First Seen Computer, Product Name, Company, Description, Global Flags
 - Selected:** Publisher or Company, Count

You must choose the columns you want to appear in the table – your Metrics choices for the Graph Settings are not imported to the table. You can double-click on a data element in the Available column to move it to the Selected column, and vice versa. You also can use the arrow buttons to move items back and forth between Available and Selected, and to change the order of data in the table.




As with table-only portlets, you can drag and drop columns to rearrange them, and can sort data by clicking on column heads. You cannot group by column and cannot filter the data in the table itself.


Using Filters in Portlets

Some portlets allow you to use filters to limit and focus the information a portlet displays. For example, you could create a portlet that shows the connection status of computers but filter out those in Visibility mode policies.

Filters do not make sense for certain portlets – RSS feeds and HTML pages, for example – and are not used on the pre-configured portlets installed with the App Control Server. If the portlet you are creating or editing includes a filtering capability, you will see a Filters panel on the Portlet Details page. The illustration below shows the initial building blocks of a portlet filter.



This initial filter view shows the top-level group operator. To have the filter actually do anything, you need to add at least one *expression*, a set of parameters that can be evaluated as true or false against App Control data. For example, to have the filter include only those computers containing “Laptop” in their name in the portlet data, you would create the following filter.




Each expression consists of a parameter--some kind of data that is available in the App Control database, an expression operator, and a value. You choose the parameter and operator from menus that vary depending upon that type and subtype of portlet. You type in the value you want to match.

Every expression belongs to a group, even if the group includes only one expression. While an expression might evaluate to true on its own, the group operator determines whether the *group* is true, as [Table 106](#) shows.



Table 106: Group operators in portlet filters

Operator	Effect
AND	If <i>all</i> expressions in the group are <i>true</i> , the group is true. For the top-level group, this means that data for which all expressions in the group are true is displayed in the portlet.
OR	If <i>at least one</i> expression in the group is <i>true</i> , the group is true. For the top-level group, this means that data for which at least one expression in the group is true is displayed in the portlet.
NOTAND	If <i>at least one</i> expression in the group is <i>false</i> , the group is true. For the top-level group, this means that data for which at least one expression in the group is false is displayed in the portlet.
NOTOR	If <i>all</i> expressions in the group are <i>false</i> , the group is true. For the top-level group, this means that data for which all expressions in the group are false is displayed in the portlet.



With **AND** as the group operator and a single expression, if the expression is true, the group is true, and the data matching the expression will be included in the portlet. As the table describes, however, adding expressions and using other operators can provide more powerful and complex filters. The illustrations below show some examples:



If you created a “Top 5 Computers with First Seen Files” portlet as shown in the details above, it displays the five computers that have the most first seen files. Note that there is not a filter on this data. Perhaps you would like to eliminate data for files that were on computers when the agent was installed and concentrate on anything that arrived afterward. To accomplish this, you could add an expression and create a filter to eliminate “initialized” files, as shown on the left, below.



Filter removes files present at initialization from the data used by the portlet.



Filter removes files present at initialization and files whose metadata shows Microsoft Corporation as publisher or company from the data used by the portlet.

To further fine-tune your portlet, you might decide to eliminate all files that identify “Microsoft Corporation” as the publisher in addition to initialized files since you know that you installed several Microsoft applications on all computers after initialization and it is not necessary to track these in your portlet. To accomplish this, you could change the group


operator to OR and create a new expression to produce a filter as shown in the right half of the illustration above.

As long as you can use the same group operator to accomplish your goal, you can continue adding expressions to a group.

Nesting Groups of Expressions

You can nest groups of expressions within a filter. Each expression in a filter group shares the same top-level operator (i.e., AND, OR, NOT AND, NOT OR), and the results of the group are treated like an expression for the group above it. Group level can be determined by the indentation of the group and its expressions – those to the left are higher-level groups than those farther to the right.

The filter shown below indicates that files whose data is displayed in the portlet must NOT be both initialized AND either from Dell OR from Microsoft. The OR group is at the same level as the Initialized expression, and the NOT AND group contains everything in the filter.



Note

Because some pre-processing of filters occurs as you choose each building block of an expression or group, you might notice a several second time delay after filter construction actions.

Chapter 25

Locating Files

This chapter explains how to use the Find Files page to locate or verify the existence of specific executable files on computers running the App Control Agent. Find Files locates *instances* of files, not their listings in the File Catalog.

Unified Management

If you have enabled Unified Management, you can use the management server to search multiple App Control Servers at once.

Sections

Topic	Page
Find Files Overview	713
Initiating Find Files from Other Pages	713
Defining a Search on the Find Files Page	714
Using Find Files Results	717
Saved Views for File Searches	720

Find Files Overview


The App Control Server keeps track of all “interesting” files on all connected computers running the App Control Agent, in near-real-time. Because of this “live inventory,” you can quickly locate a file or group of files matching a name, hash, or other criteria available in the database for your App Control Server. For offline computers, the file inventory includes all files from the last time they were connected.

This chapter focuses on the Find Files page, which opens by default with a filter that allows you to search for a file by name. As with the Files on Computers tab, you can add filters to fine-tune the results you get, and for many searches, you can create a Saved View. In addition, some other console pages include a Find Files button or link that displays Find File results for a particular file on the page.

Notes

- You also can search for file instances on the Files on Computers tab of the Files page, although you must add all filters manually, including the file name filter.
- Certain features allow you to exclude files from the file inventory, and excluded files might not appear in search results. See the Overview in [Chapter 7, “File, Publisher, and Application Information,”](#) for details.

Initiating Find Files from Other Pages

In addition to going directly to the Find Files page, you can search for file instances by clicking the Find File button  next to a file name or hash in some tables on other pages. This initiates a search by hash for all instances of that file. You can do this from:

- Files page (both the Files Catalog tab and the Files on Computers tab)
- File Group Details page
- Baseline Drift Report Results page (Files views)
- Snapshot Content page
- Find Files page (for instances of one specific file only)
- Software Rules/Publishers page (for files from one publisher)
- Approval Request Details page (for instances of a file whose approval is requested)

Other console pages have links that initiate a Find Files search pre-configured to find files relevant to the page you are on. These include:

- File name links on the Files page – When you click on a highlighted filename on the Files page, the console displays a Find Files report of all files *associated with* the named file (that is, files installed by or that are copies of the named file).
- File Details page and File Instance Details page – The *All File Instances* link in the Related Views menu initiates a search for the file whose details you are viewing.
- Add/Edit Policy page – The Related Views menu on this page has two file searches: *All Files on computers in this policy* and *Unapproved files on computers in this policy*.
- Computer Details page – The Related Views menu includes *Files on this Computer*, which displays a Find Files report of all files on the computer.

When Find Files results appear for any of these queries, you can further refine, as with any other console table, by showing or hiding columns and applying additional filters – if the Filters panel is not showing, click the Show Filters link.

You can also find files from the Find Files or Events portlets on the Home Page dashboard, which includes a Find Files or Events portlet.

Defining a Search on the Find Files Page

You can create file queries on the Find Files page based on any parameter available on the Filters menu. As with any page, you can combine filters, in some cases including more than one of the same type of filter (for example, *File Name is calc.exe* or *File Name is add.exe*) in the same search. If you are searching for one specific file, you can search by file name or hash identifier.

Tip

Combination searches based on file name and hash are useful for detecting attacks by a malicious program that presents itself with different file names but contains the same data.

Finding Files by Name

Although searching by hash is a better way to be certain you find all instances of a file, searching by name is the easiest type of search to create from scratch. File Name searches allow you to use different operators to expand or narrow the matches you get from the search, as shown in [Table 107](#).

Table 107: Operators for the File Name Filter

Field	Description
contains	Any file whose name <i>contains</i> the text in the box. This operator can cause time-consuming and inefficient searches; use an alternative if possible
does not contain	Any file whose name does <i>not</i> contain the text in the box. This operator can cause time-consuming and inefficient searches; use an alternative if possible
begins with	Any file whose name <i>begins with</i> the text in the box.
ends with	Any file whose name <i>ends with</i> the text in the box.
is	Only files that <i>exactly</i> match the text you enter. When you choose is , be sure to include the full file name, including extension, in the File Name text box.
is not	Any file whose name does not exactly match the text you enter. Note that if you enter “calc” as the File Name, for example, the results from is not will include “calc.exe”, “mycalc”, etc.
is empty	Any file whose name is missing or blank.
is not empty	Any file whose name is <i>not</i> missing or blank.

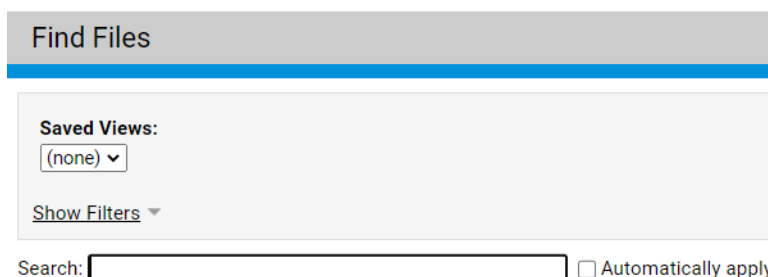
By default, the Find Files page opens with the File Name filter and the operator “is”, meaning file instances exactly matching the text you enter in the box will be in the results.

When searching for a file, consider the following best practices:

- **No Wildcards** – Do not use wildcards (*, ?, etc.) in your search string for a file name. The App Control Server will attempt to match them literally, and the results will not likely be what you want. Instead, use the operator menu, which provides choices that accomplish the same thing, without requiring you to type in special symbols.
- **Case Sensitivity and Platforms** – Although case-sensitivity varies among operating systems, file searches in App Control are not case sensitive; for example, searching for “Myfile.exe”, myFiLE.exe”, or “myfile.exe” will return the same results
- **Limit Results** – Try to define your search parameters so that the results are limited to a reasonable number of files. The console does limit the number of matching files it will return, and you will see a message instructing you to try a narrower search if the number of results exceeds what can reliably be inserted into one table.
- **Choose the Most Efficient Search Criteria** – Some search criteria are more efficient than others. In general, a filter that allows searching for an exact match rather than requiring a string analysis will be much faster and have less likelihood of database timeouts. For example, if you want to find all files with a particular extension (such as .exe) using the File Name filter and choosing “ending with .exe” is very inefficient. In this case, use the Extension filter. Searching for a file using the “containing” operator (such as, “File Name contains setup”) is particularly inefficient.
- **Auto-Completion** – Many fields on the Find Files page, including File Name, provide automatic matching of the string as you type it, showing matching choices in a menu.

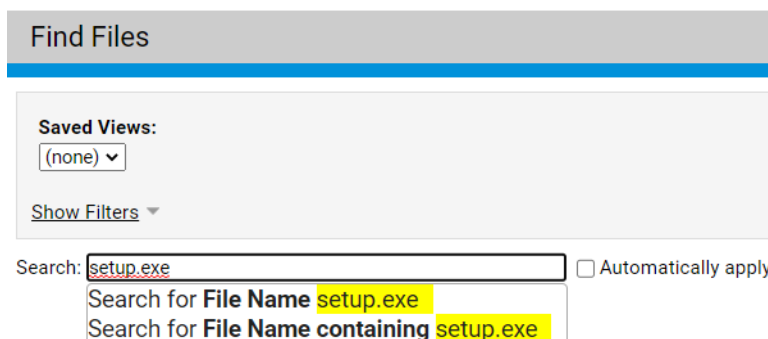
To locate instances of a file by name:

1. In the console menu, choose **Tools > Find Files**. The Find Files page appears.



The screenshot shows the 'Find Files' page header. Below the header is a 'Saved Views' section with a dropdown menu set to '(none)'. Below that is a 'Show Filters' link with a dropdown arrow. At the bottom is a search bar with the text 'Search:' followed by an empty input field and a checkbox labeled 'Automatically apply'.

2. Specify a File Name, or a portion of a filename, that you would like to use in the search. As you type, the search bar gives different file search options.



The screenshot shows the 'Find Files' page header. Below the header is a 'Saved Views' section with a dropdown menu set to '(none)'. Below that is a 'Show Filters' link with a dropdown arrow. At the bottom is a search bar with the text 'Search:' followed by an input field containing 'setup.exe' and a checkbox labeled 'Automatically apply'. A dropdown menu is open below the search bar, showing two options: 'Search for File Name setup.exe' and 'Search for File Name containing setup.exe', both with yellow highlights.

3. Choose the option that matches what you are looking for. Choose the *File Name* option if you want only files exactly matching the File Name you entered so far (including extension). In general, avoid the “containing” operator unless it is absolutely necessary since it results in a less efficient search.
4. If you check the **Automatically Apply** box, as soon as you press Return on your keyboard or click on an option in the search box, all files (on all computers) matching the File Name-option combination you entered are displayed in the Find Files table.
5. If you do not check the **Automatically Apply** box, choosing an option in the search box opens the Show Filters panel so that you can add additional search parameters. In this case, you click **Apply** in the Filters panel to see the search results when you have finished adding filters.

Adding a Pathname to a File Search

File Path is one possible addition to a search for files by name. It may also be useful in other searches, for example, if you want to find all files from a specific publisher in a specific directory and its subdirectories.

You specify a pathname *without* the name of the file you want to find. For example, if you wanted to find *calc.exe* in *c:\windows\system32*, you would specify the following filters:

The screenshot shows the 'Find Files' search interface. At the top, there is a 'Saved Views' dropdown menu set to '(none)'. Below it is a 'Hide Filters' dropdown. The 'Filters' section contains an 'Add filter' dropdown. Two filters are currently active:

Filter Type	Operator	Value	Action
File Name	is	calc.exe	+
File Path	is	c:\windows\system32	+

At the bottom of the filters section are three buttons: 'Apply', 'Cancel', and 'Reset'.

Specifying that the File Path **is** *c:\windows\system32* indicates that you want to find files only in the named folder, not in subfolders. If you want to search for all files in a named folder and its subfolders, you use the operator **contains**. For example, if you specified File Name is *calc.exe* and File Path **contains** *c:\windows\system32*, you would find all instances of *calc.exe* in *system32* and at any level underneath it.

Platform Note: Using a pathname in a file search will limit your search to computers that support the platform-specific delimiters (i.e., ‘\’ or ‘/’) and other special path characters you use.

Finding Files by Hash

App Control supports three hash types: SHA-256, SHA-1, and MD5. If you have a hash from some source other than App Control and want to search for it, you can search for that file on your computers from the Find Files page by choosing the hash type from the Filters menu entering the hash into the filter field.

Find Files

Saved Views:
(none) ▾

Hide Filters ▾

Filters

Add filter ▾

SHA-1 is 43a4b8fe96c5bfe5c861cdecbf19ee2c5da7bb7e +

Apply Cancel Reset



On some files, App Control does special processing to create SHA-256 hashes that will be identical for identical files. Because of this, searching by using *externally created* SHA-256 hashes is not recommended.

The best way to search by hash is to locate the file of interest in one of the Files tabs and then click on the Find File button next to the file. The console will run the Find File search without you needing to type or cut and paste the hash string.

As with file names, the console shows a list of matching hashes as you type in digits, and if there is only one item on the list, you can pick it without entering the entire hash string.

Using Find Files Results

The Find Files results page provides all of the tools available on the Files page, both for getting further information and taking action on one or more files in the table:

- When your initial search is broad enough to include different files (not just different instances) in the results, you can initiate a new search for all instances of one specific file by clicking the Find File button  next to that file.
- You can click the View Details button  next to any found instance of the file to get more information about that instance.
- You can select files from the results and operate on them with the approval or ban commands on the Action menu. For example, you can **Approve Locally** or **Remove Local Approval** for any file in the results by checking the box to the left of the file listing and clicking the appropriate button.
- If you have Carbon Black File Reputation enabled, you can view additional information (if available) for any file in the results by checking the box to the left of the file name and choosing **View Carbon Black Reputation Data** from the Action menu.
- If you have enabled third-party analysis tool integrations via the App Control Connector, **Analyze with ...** commands appear on the Action menu. You can use the available commands to send one or more of the found files for analysis.
- The Action menu also includes commands that allow you to find computers that have, or are missing, one or more files in the Find Files results.

Notes

- View Carbon Black Reputation Data results open in a new tab for each file. For multi-file requests in Internet Explorer, the popup blocker may block the results for each file after the first one.
- In Find Files results, you can rearrange columns, download results in comma-separated-value format, and add the Find File results to a Snapshot. For more information, see [“Console Tables”](#) on page 66.

Special Cases in Results

Files on Offline Computers

If a computer is offline, a Find Files search will include the matching files from that computer's most recent synchronization with the App Control Server in the results. The next time the computer connects to the App Control Server, its file information is updated within a short time (depending upon the network traffic and how many computers are being updated), and the updated information becomes available to Find File.

Find File results tables that include the Computer column have an indicator to the left of the computer name showing whether the computer is connected and up-to-date. A blue circle indicates that the computer is connected and up-to-date. A yellow circle indicates a computer that is connected but awaiting an action (agent out of date, requires reboot, or other reasons). A gray circle indicates a disconnected computer. When you move the mouse cursor over a status circle, more information for that computer's status appears below the name, including how long a computer has been offline.

	Date Created	Computer	File Name	Trust	Threat	Local State	Global State
	Feb 17 2016 10:55:32 AM	MYCORPLaptop-21	notepad.exe	10		Approved	Approved
	Feb 17 2016 10:41:27 AM	MYCORPLaptop-6	notepad.exe	10		Approved	Approved
	Feb 17 2016 10:41:27 AM	MYCORPLaptop-3 Disconnected for less than a day	notepad.exe	10		Approved	Approved
	Feb 17 2016 10:55:32 AM	MYCORPLaptop-1	notepad.exe	10		Approved	Approved
	Feb 16 2016 01:53:55 PM	MYCORPLaptop-2	notepad.exe	10		Approved	Approved
	Feb 17 2016 11:17:54 AM	MYCORPDesktop-33	notepad.exe	0		Approved	Unapproved

Deleted Files

If a file matching a Find Files search has been recently deleted from a computer, it can be included in Find File results if you choose, although this is not done by default. To include deleted files, check the *Show deleted files* box in the bottom right of the Find Files page;

the table is immediately updated to show any deleted files matching your search parameters. Deleted files are labeled as such in the Find Files results.

Find Files

Saved Views: (The Current View Has Unsaved Changes - Discard) (none) Add Group By: (none) Show Deleted Files

Hide Filters Show Columns Show Snapshot Export to CSV Refresh Table

Filters
Add filter

File Name is notepad.exe

Apply Cancel Reset

Action Search: Automatically apply Showing 6 out of 6 item(s)

	Date Created	Computer	File Name	Threat	Local State	Global State
<input type="checkbox"/>	Feb 16 2016 01:53:55 PM	MYCORPLaptop-2	notepad.exe	<input checked="" type="checkbox"/>	Approved	Approved
<input type="checkbox"/>	Feb 17 2016 11:17:54 AM	MYCORPDesktop-33	notepad.exe	<input checked="" type="checkbox"/>	Approved	Unapproved
<input type="checkbox"/>	Feb 17 2016 11:05:24 AM	MYCORPDesktop-12	notepad.exe (Deleted)	<input checked="" type="checkbox"/>	Approved	Approved
<input type="checkbox"/>	Feb 17 2016 11:05:24 AM	MYCORPLaptop-5	notepad.exe	<input checked="" type="checkbox"/>	Approved	Unapproved

Deleted files are removed from the database on the same schedule as old events. See [“Advanced Configuration Options”](#) on page 744 for information about configuring this time period.

Notes

- If you are searching for deleted files using the Deleted filter, you must check the *Show deleted files* box in the bottom, right corner of the page before any matching results will appear.
- Including deleted files in a search will slow down the search and consume more resources, so use this feature only when necessary.

Files on Deleted Computers

If a computer has been deleted from the Computers list, its files remain in the database of Files on Computers for one day. This means that a Find Files search could include results from deleted computers. Deleted computers are labeled as such in the Find Files results.

	Date Created	Computer	File Name	Threat	Local State	Global State
<input type="checkbox"/>	Feb 16 2016 01:53:55 PM	MYCORPLaptop-2	notepad.exe	<input checked="" type="checkbox"/>	Approved	Approved
<input type="checkbox"/>	Feb 17 2016 11:17:54 AM	MYCORPDesktop-33	notepad.exe	<input checked="" type="checkbox"/>	Approved	Unapproved
<input type="checkbox"/>	Feb 17 2016 11:05:24 AM	MYCORPLaptop-6 (Deleted)	notepad.exe	<input checked="" type="checkbox"/>	Approved	Approved
<input type="checkbox"/>	Feb 17 2016 11:05:24 AM	MYCORPLaptop-5	notepad.exe	<input checked="" type="checkbox"/>	Approved	Unapproved

Files on Computers Still Initializing or Synchronizing

If a computer has just had the agent installed and is still initializing, some of its files are available to Find Files, but its full file inventory is not available until initialization is complete. To determine whether a computer is still initializing, go to the Computers page and search for the computer.

Similarly, if an agent is re-synchronizing with the server, changes in its file information are not complete until the synchronization is finished. You can view synchronization progress on the Computer Details page or, if you add the Synchronization column to the table, on the Computers page.

Saved Views for File Searches

If you have a complex search that you think you will use often, you may be able to save it as a Saved View.

Notes

- Certain Find File reports, including those initiated from the Find File button on other pages, cannot be saved because they were run in a specific context that might not be in effect if executed again from the Find Files page – the Saved Views panel does not appear in these cases. As an alternative, you might be able to duplicate and save the search you want by using filters on the Files on Computers tab of the Files page.
- ReadOnly users cannot save views. Also, some custom login account groups might not have permission to save views.

To create a Saved View on the Find Files page:

1. In the console menu, choose **Tools > Find Files**. The Find Files page appears with the default filter, *File Name*, and the default operator, *is*.
2. Choose each filter you would like to add to the search criteria, provide any text required to configure the filter, and click **Apply**.
3. When you have finished adding filters, enter a name in the Saved Views box above the table and then click **Add**. You now will be able to choose the Saved View you created from the Saved Views menu and get results for this same search whenever you choose.

Chapter 26

System Configuration

This chapter explains settings that enable you to configure and maintain your App Control Server installation. Access to the System Configuration page is available only to login accounts in the Administrators group or in customized groups with View System Configuration and Manage System Configuration boxes checked.

Sections

Topics	Page
Overview	722
Viewing Server Status and Options	724
Configuring Active Directory Integration	726
Configuring Agent Management Privileges	727
Managing the App Control Event Database	733
Securing Agent-Server Communications	739
Advanced Configuration Options	744
Adding a Login Banner to the Console	749
Backing Up the App Control Server	750
Restoring the App Control Server	753
Configuring Alert and Approval Request Mail	754
Managing App Control Licenses	759
Activating Carbon Black File Reputation	763
Activating Carbon Black EDR Integration	768
Configuring Unified Management	769
Configuring SAML Logins	769

Overview

The System Configuration pages present both read-only status information and configurable settings for use by App Control Administrators. The configuration information is organized on a series of tabbed views, some of which have several panels:

- **General** tab – Server status information, options for integrating the App Control Server with Active Directory or LDAP, App Control Agent Management options, and user password configuration requirements.
- **Events** tab – Configuration settings for managing the App Control Server's own database and options for setting up supplemental external event logging, including Syslog.
- **Security** tab – Shows current status of secure communications between App Control Agents and the App Control Server, and provides options for enabling certificate verification for these communications if not already enabled.
- **Advanced Options** tab – Options for database backup, automatic agent upgrades, App Control Console login timeout, files for App Control to ignore, API access, deleting offline computers, allowing use of expired publisher certificates, and letting Carbon Black File Reputation update detection indicators, system health indicators, and definitions of updaters.
- **Mail** tab – Configuration settings for sending email when a App Control alert is triggered or an approval request is resolved.
- **Licensing** tab – Shows the number and type of App Control Agent licensed for your server, and allows you to update your license key; also allows you to enable and configure Carbon Black File Reputation.
- **External Analytics** tab – Configuration settings for External Analytics, which enables the App Control Server to export data it collects from endpoints to external analysis tools. See [Appendix F, "Exporting Data for External Analysis,"](#) for information about the settings on this tab.
- **Connectors** tab – Configuration settings for integrating the App Control Server with one or more network security devices or services. [Appendix C, "App Control Connector,"](#) provides additional information about some of the connections enabled by the settings on this tab.
- **Unified Management** tab – Configuration settings required to enable Unified Management. See ["Enabling Unified Management and Adding Servers"](#) for information about the settings on this tab.
- **SAML Login** tab – Configuration settings required to integrate App Control with an identity provider that uses Security Assertion Markup Language (SAML). See ["Configuring SAML Logins"](#) for more information regarding the settings on this tab.

To display the System Configuration page:

1. On the console menu, click the configuration (gear) icon and choose **System Configuration**. The System Configuration page appears.
2. By default, the App Control Console displays the General tab of the System Configuration page. Select another tab if you want to view or change something not on this tab.

The General Configuration Tab

The General tab of the System Configuration page has three sets of configuration fields:

- The [Server Status](#) panel shows information about your App Control and database servers, including their addresses.
- The [Active Directory/LDAP Integration](#) panel allows you to configure AD or LDAP integration with the App Control Server.
- The [Agent Management](#) panel allows you to set up access to special agent management commands by user, group, or password.
- The [User Passwords](#) panel allows you specify the password criteria that each user must adhere to.

The screenshot displays the 'General' configuration tab with the following settings:

General Settings

Server Status

- Carbon Black App Control Version: 8.5.0.12
- Server Address: APC-58610
- Server Port: 41002
- Server Timezone: -Automatic-
- Database Schema Version: 8.5.0.12
- Database Address: local
- Database Auth.Type: NT
- Database Size: 224.25 MB
- Free Local Disk Space: 41.5 GB / 79.7 GB
- CL Version: 1169

Active Directory / LDAP integration

- AD-Based Logins: Disabled
- AD Security Domain: (none)
- AD-Based Policy: Disabled
- Windows 2000 DCs:
- Test AD Connectivity:

Agent Management

- Windows User/Group To Manage Agents: None User or group Pre-defined group
Authenticated Users
- Mac User/Group To Manage Agents: None User Group
- Linux User/Group To Manage Agents: None User Group
- Enable Global Password:
- Enter Password:
- Confirm Password:

User Passwords

- Minimum Length: 8
- Minimum Special Character Count: 0
- Expiration: 0 Days
- Expiration Warning Period: 14 Days
- Expiration Grace Period: 0 Days

Buttons: Edit, Update, Cancel

Viewing Server Status and Options

The top panel on the General tab of the System Configuration page is *Server Status*, which displays App Control Server parameters and allows editing of some of them (see [Table 108](#) for details).

Important

Parameters on the Server Status panel tell you about the size of the App Control database and the amount of free space on the computer running App Control Server. These do not, however, report on whether an *external* SQL database is running out of space. Regardless of which database option you choose, monitor your App Control database regularly to be sure it does not overflow and prevent the App Control Server from operating. See the *VMware Carbon Black App Control Installation Guide* manual for more information on database configuration. Also, see “[Creating Alerts](#)” on page 612 for information on database-related alerts.

To display server status information:

1. On the console menu, click the configuration (gear) icon and choose **System Configuration**.
2. If it is not already showing, click on the **General** tab. The General configuration options appear, with the Server Status panel showing at the top.

The screenshot shows the 'General' tab selected in the System Configuration page. The 'Server Status' panel is visible, displaying the following information:

Carbon Black App Control Version:	8.5.0.12
Server Address:	APC-58610
Server Port:	41002
Server Timezone:	-Automatic- (dropdown menu)
Database Schema Version:	8.5.0.12
Database Address:	local
Database Auth.Type:	NT
Database Size:	224.25 MB
Free Local Disk Space:	41.5 GB / 79.7 GB
CL Version:	1169

3. To change timezone, click the **Edit** button, make the changes, and click **Update**, and then click **Yes** on the confirmation dialog. See [Table 108](#) for details about the other settings.

Table 108: Server Status Information and Configuration Options

Field	Description
App Control Server Version	Version number of the installed App Control Server software. (Read Only)
Server Address	<p>IP address or qualified DNS name for the App Control Server.</p> <p>If you change the server address, you must reinstall the App Control Agent on all computers (although not if you change from an IP address to an equivalent DNS name, or vice versa). As soon as the agent is installed, computers reinitialize and all files except those explicitly banned on the server become locally approved and permitted to run. So that you can use the same policies, App Control automatically updates existing agent installation packages with the new IP address so that they direct computers to report to the correct server when they come back online.</p> <p>Note: IPv6 may be used for communications with the App Control Server, but a numeric IPv6 address may not be accepted in certain versions of the Firefox browser. To avoid this problem, use one of the other supported browsers or a fully qualified DNS name.</p>
Server Port	App Control Server port that is dedicated to communications with computers running the App Control Agent. This cannot be changed after server installation. (Read Only)
Server Timezone	The timezone used by the App Control Server. Normally this will be set to Automatic, which uses the same time zone as the operating system. However, to account for non-standard handling of daylight saving time in certain zones, you can set the server timezone explicitly, using the dropdown menu here.
Database Schema Version	Normally the database schema version is the same as the App Control Server version. You can, however, use existing databases when you upgrade or reinstall the server. In this case the database schema version can be different. For Carbon Black Support use. (Read Only)

Field	Description
Database Address	Shows whether your database is Local or on a separate server, in which case it provides the address. (Read Only)
Database Auth. Type	This indicates the type of database authorization you chose when you installed App Control Server. It is either NT , indicating that you are controlling database access by Windows NT account or group, or SQL , indicating that you are using a login and password specific to your SQL Server. (Read Only)
Database Size	Amount of disk space currently used by the App Control Server database. (Read Only)
Free Local Disk Space	Amount of available <i>local</i> disk space on the App Control Server. If the App Control Server database is on the same system as the App Control Server, you can periodically monitor this value to determine how quickly events are accumulating and whether you need to adjust the event log deletion period. (Read Only) Important: This field reports free space on the App Control Server system only. If you are using a remote database, you must check available space directly on that system.
CL Version	This is a configuration list version number reflecting the current set of policy rules. As App Control Console users create bans, changes policies, and take other actions that change the configuration of your App Control Server, this number increments. Carbon Black Support can use CL version in certain troubleshooting situations. (Read Only)

Configuring Active Directory Integration

The App Control Server can take advantage of your Active Directory (AD) environment to set access privileges for users of the App Control Console, assign security policies to computers, provide user and computer metadata, and designate certain groups or users to be able to install software (and have it automatically approved) on App Control-managed computers. You configure AD integration on the General tab.

To display AD integration configuration options:

1. On the console menu, click the configuration (gear) icon and choose **System Configuration**.
2. If it is not already showing, click on the **General** tab. The General configuration options appear, with the AD/LDAP integration options showing in the middle panel.

Active Directory / LDAP integration

AD-Based Logins:

AD Security Domain:

AD-Based Policy:

Windows 2000 DCs:

Test AD Connectivity:

3. To configure AD or LDAP integration, click the **Edit** button at the bottom of the page.

4. Enter the AD Security Domain for the server, and check the Windows 2000 DCs box if you are using a Windows 2000 domain controller.
5. In the Active Directory/LDAP integration panel, click the **Test AD Connectivity** button to determine whether there is an AD server accessible to the App Control.

Note: This test does not confirm that the domain you provided is legitimate.
6. Click the **Update** button, and then click **Yes** on the confirmation dialog. See [Table 109](#) for more details about these settings.

Table 109: Active Directory/LDAP Integration Options

Field	Description
AD-based logins	Choosing Enabled in this field allows users to log in to the App Control Console using AD accounts and passwords. See “Enabling Console Access via AD Accounts” on page 92 in Chapter 3 , “Managing Console Login Accounts,” for more detail.
AD security domain	Specifying an AD security domain in this field directs the App Control Server to look in that domain for the security groups to use for App Control Console user login validation. If you do not specify a security domain, the login domain for each console user is used, and so the relevant security groups must be in each user’s domain for that user to be able to log in.
AD-based policy	Choosing Enabled in this field allows you to automatically assign App Control policies to computers based on AD or LDAP. See Chapter 4 , “Managing Computers,” for more detail.
Windows 2000 DCs	Checking this box indicates that your network is using Windows 2000 domain controllers. This disables the AD security domain value you provided, if any, since it relies on cross-domain membership tests that are only available with Windows 2003 SP2 domain controllers.
Test AD Connectivity	Clicking the Test button tests connectivity between the App Control Server and Active Directory. If it reports Success, you can use App Control’s Active Directory integration features. If it reports Error, your App Control Server cannot access Active Directory, and you need to resolve this problem before the integration features can be used. <p>Note: This test only discovers whether an AD server – any AD server – is accessible to your App Control Server. It does not confirm that the integration of the two servers will be successful with the domain and Windows 2000 choices you provide.</p>

Configuring Agent Management Privileges

You may, in conjunction with your Carbon Black Support representative, use special Agent Management commands for App Control Agent management. Because App Control Agent plays a critical role in managing and protecting your computers, you can and should limit access to these commands. In the Agent Management section of the General tab, you can choose one or both of the following methods for controlling agent command access:

- For each client platform, you can specify a user or group allowed to run the commands

- You can specify a password that will be required to run the commands. The password must be between 1 and 64 characters long, be in the ASCII character set, and must not contain the following special characters: | > < & % () @ . [] { } ; : ^ = ! ' " ` ~ ,

If you define both a user/group and a password, *either* access method is sufficient on its own. The current agent management configuration when agent installation packages are created is built into the agent. If you change the password, the App Control Server updates online agents with the new password, but agents not online must continue using the old password. Likewise, changes in the user or group access definition are not effective on an offline agent unless the old agent is uninstalled and a new one is installed by some method.

Note

Configuring the Agent Management options *before* generating any agent installation packages is the most efficient way to set a global agent password or user/group access choice.

For new installations of App Control Server, you are prompted to provide an Agent Management access method during the installation process – this is the best time to choose an option.

To display agent management configuration options:

- On the console menu, click the configuration (gear) icon and choose **System Configuration**.
- If it is not already showing, click on the **General** tab. The General configuration options appear, with the Agent Management options showing in the bottom panel.

- To configure agent management, click the **Edit** button at the bottom of the page, make the needed changes, click the **Update** button, and then click **Yes** on the confirmation dialog. See [Table 110](#) and [“Connection Status and Agent Management Choices”](#) on page 729 for more details about these settings and guidance on choosing options.

Table 110: Agent Management Configuration Options

Field	Description
Windows User/ Group to Manage Agents	<p>If defined, the specified Windows user or group is allowed to run special commands for App Control Agent management on computers that recognize that user or group.</p> <ul style="list-style-type: none"> Choose the <i>User or group</i> radio button to enter a user or group name manually; you also can enter a user or group SID in this box. Choose the <i>Predefined group</i> button to choose a Windows group (e.g., Local Administrators), from a menu.
Mac User/ Group to Manage Agents	<p>If defined, the specified Mac user or group is allowed to run special commands for App Control Agent management on computers that recognize the user or group. Choose the <i>User</i> radio button or the <i>Group</i> button and enter a name in the box.</p>
Linux User/ Group to Manage Agents	<p>If defined, the specified Linux user or group is allowed to run special commands for App Control Agent management on computers that recognize the user or group. Choose the <i>User</i> radio button or the <i>Group</i> button and enter a name in the box.</p>
Enable Global Password	<p>If defined, the specified password may be used by <i>any</i> user to run special commands for App Control Agent management from the client computer. Check the box to enter the password.</p> <p>If you define both a password and a user or group for agent management, you only need one or the other for access.</p> <p>The password must be between 1 and 64 characters long, be in the ASCII character set, and must not contain the following special characters: > < & % () @ . [] { } ; : ^ = ! ' " ` ~ ,</p>

Connection Status and Agent Management Choices

Your Agent Management access choice may be dictated by whether or how often your client systems running the App Control Agent are connected to the App Control Server.

If a computer is *never connected to the server*, you can provide access by choosing an Agent Management password before generating installation packages. This password is built into the agent, and can be changed only by one of the following means:

- Installing a new agent package generated after the password change.
- Importing a new configuration list from the App Control Server after you have changed the global password; see your Carbon Black Support representative for instructions on importing a configuration list.

Another option for systems never connected to App Control Server is specification of a group that can be guaranteed to exist on all machines, such as Local Administrators for Windows computers. The suitability of this method depends on how your organization manages administrative accounts, but it lets you control access to agent management commands by adding or removing users from the named group, independent of changes in the App Control.

If a computer will be *connected to the App Control Server occasionally*, you have more flexibility in choosing and changing client management access methods. Changes to a

password, or to user or group definition, propagate to the agents the next time they connect.

If all of your computers will *always be connected to the App Control Server* (or can be if needed), you have the most flexibility in configuring Agent Management access since changes you make will go to your connected agents as soon as the agent and server are in contact. In this case, you might find it more convenient to choose a well-known group, or define a new group, such as "App Control Local Administrators", and give its members access to the management commands. Groups also allow the use of such tools as *runas*, *psexec*, or *sudo*, to run commands using alternate credentials. You also can use a password if you choose.

Note

When running on Windows Vista and later, membership in pre-defined security groups like Administrators requires that the application run as an administrator. If you are not certain that a user has this elevated privilege, using a built-in group for Agent Management access may not be a good choice if you will be using computers running Vista or Windows 7.

Specifying User Password Requirements

You can increase security to the product by specifying password complexity and ensuring periodic password changes. You can specify:

- The minimum length of the password.
- The minimum amount of special characters required, if any.
- How long the password is valid for before it expires.
- How many days of warning the user will receive before the password expires.
- The number of days after expiration where users can still log in; however, they are immediately redirected to change their password and access to other content is prohibited.

NOTE: If a user's password expires after the grace period or if there is no grace period, the user must request a password reset from their App Control Administrator. The administrator would have two options to reset the password:

- Reset the password to an agreed upon password and the user could reset it at a later date.
- Create a temporary password that redirects the user to immediately change it upon login.




See: ["To change a console password and other login account details:"](#)

To specify password requirements for console users:

1. On the console menu, click the configuration (gear) icon and choose **System Configuration**.

2. If it is not already showing, click on the **General** tab. The General configuration options appear, with the User Passwords panel at the bottom of the page.

User Passwords

Minimum Length:	<input type="text" value="8"/>
Minimum Special Character Count:	<input type="text" value="0"/>
Expiration:	<input type="text" value="0"/> Days 
Expiration Warning Period:	<input type="text" value="14"/> Days 
Expiration Grace Period:	<input type="text" value="0"/> Days 

3. To configure user password requirements, click the **Edit** button at the bottom of the page, make the needed changes, click the **Update** button, and then click **Yes** on the confirmation dialog. See for more details about these settings and guidance on choosing options.

Table 111: User Password Configuration Options

Field	Description
Minimum Length	This is the minimum number of characters that a password must be. Default is 8 characters. NOTE: This does not apply to the admin username on initial install. However if you update the admin password after the initial install, it will apply.
Minimum Special Character Count	This is the minimum required number of non-alphanumeric characters, such as: > & % () @ . [] { } : ; ^ = ! ' " ` ~ , Default is 0 special characters required. NOTE: This does not apply to the admin username on initial install. However if you update the admin password after the initial install, it will apply.
Expiration	This is the number of days after a password is set that it will expire. Default is zero (0) days; this means there is no expiration. NOTE: This does not apply to the admin username on initial install. However if you update the admin password after the initial install, it will apply.
Expiration Warning Period	This is the number of days before password expiration that the user will see a console message warning them about the impending password expiration. Default is 14 days.
Expiration Grace Period	This is the number of days after the password has expired that the user can still successfully log in using the expired password but will be forced to change their password immediately after logging in. When logging in under the grace period, the user can do no other actions in the console until the password change is completed.

Event Management Options

App Control event data is stored in a SQL Server database, and grows over time at a rate that corresponds to file activity on your network. The Events tab provides two sets of options for managing events data generated by App Control:

- The Event Log Management panel provides options for managing the size of the *primary* App Control Server database and for archiving events.
- The External Event Logging panel provides options for enabling supplemental, external logging of App Control events to another SQL Server or a Syslog management server. Use of supplemental external logging may allow you to reduce the amount of data you keep in the primary database.

System Configuration

General
Events
Security
Advanced Options
Mail
Licensing
External Analytics
Connectors
Unified Management
SAML Login

Event Logging

Event Log Management

Delete Events Older Than: 4 Week(s) ▾

Delete If More Than: 10000000 Events

On Limit, Delete Oldest: 10 % of Events

Archive Events Enabled:

External Event Logging

Syslog Enabled:

Syslog Address: 127.0.0.1

Syslog Port: 514

Syslog Format: Enhanced (RFC5424) ▾

Syslog Export Process Command Lines:

Use External Database:

DSN String: (none)

Edit
 Update
 Cancel

Important

Your choices for event log management may be determined by your disk capacity and the availability of an external SQL Server database for storing App Control data. Please consider this before making any changes to your logging configuration.

Managing the App Control Event Database

The Event Log Management tab includes options for limiting the growth of the App Control Server database and setting up event archiving.

Setting Limits for Event Deletion

You can set limits that delete data to keep the App Control Server database at a reasonable size. The App Control Server provides several mechanisms for handling this volume of data. App Control provides for automatic deletion of event data based on two different parameters:


- **Delete Events Older Than** – By default, events older than 4 weeks are automatically deleted, which means that event data is purged on the system and is not available for

display in reports generated by the App Control Server. You can modify the time period in the Management Configuration table.

- **Delete if More Than** – This threshold defaults to 1 million for SQL Server Express and 10 million for other SQL Server editions. This works with a second parameter, **On Limit Delete Oldest**, which allows you to define the percentage of the events deleted when you reach the limit you set. The default percentage is 10%.

Event data is deleted when *either* condition is met. You can configure these automatic deletion parameters based the available disk space on the SQL Server and your need for historical information. To determine the right values for your network, monitor disk space use on the server and adjust the event database deletion parameters accordingly.

Enabling Daily Event Archiving

If *Archive Events Enabled* is checked on the System Configuration Events tab, a separate compressed CSV file is produced for each day's event data. Daily event files are stored for one year and accessible through the Event Log Archives, which lists the date-stamped files in chronological order. Through the log, you can click on and open (or save to another location) any listed event log file. You open the Event Log Archives by clicking the Archives button  in the header of an Event log.

If *Archive Events Enabled* is not checked, no event archives are generated from that point forward.

Moving the Database to an External Server

When you installed the App Control Server, one of your choices was whether to put the App Control Server database on the same computer as the App Control Server. You might find that the volume of App Control data requires a transition from a shared to a dedicated database server.

Moving the primary App Control Server database requires steps outside of the App Control Console, including running the App Control Server installation program to reconnect to the new server. Contact Carbon Black Support if you need to make this transition.

Note

The External Event Logging options on the System Configuration Events tab are for enabling *supplemental* event logging, not for moving the primary database.

Setting up External Event Logging

App Control allows you to copy event data to an additional, external SQL Server. You also can configure event output to a Syslog server using several different output formats. The full set of settings for external event logging are shown in [Table 112](#) on page [738](#).

Logging Events to a Syslog Server

The App Control Server supports integration of its event information with Syslog servers using several formats. You configure Syslog integration in the External Event Logging panel of the Events tab.

The supported formats are:

- **Basic (RFC3164)** – the default for upgrades to v7.2.2 from pre-6.0.1 (Parity) versions
- **Enhanced (RFC5424)** – a newer standard and the default for new installations of v6.0.1 (Bit9, Parity, or App Control) and later.
- **CEF (ArcSight)** – the format to use to integrate App Control event logs with HP ArcSight ESM or HP ArcSight Logger
- **LEEF (Q1 Labs)** – the format to user to integrate App Control event logs with QRadar Log Manager or QRadar SIEM

Notes

- See the separate document *VMware Carbon Black App Control Events Guide* for more information on syslog formats supported by App Control and how to map events to them.
- If you used HP ArcSight or Q1Labs products with previous versions, you will need to see the Integration guide for information about upgrading your integration to this release.
- If you worked with Carbon Black Support to manually enable special Syslog formatting in pre-6.0.2 releases, your changes will be overwritten on upgrade to this release. Use the Syslog format menu to choose formatting.

To enable event logging to a Syslog server:

1. Prepare the Syslog server to which you want to log App Control events. See the separate *VMware Carbon Black App Control Events Guide* for more details about preparing the server.
2. On the App Control Console menu, click the configuration (gear) icon and choose **System Configuration**, and on the System Configuration page, click on the **Events** tab.
3. On the Events tab, click the **Edit** button at the bottom of the page.
4. In the External Event Logging panel, check the **Syslog Enabled** box.

External Event Logging

Syslog Enabled:

Syslog Address:

Syslog Port:

Syslog Format:

Syslog Export Process

Command Lines:

5. Provide the address (IP address or FQDN) and port number of your Syslog server in the Syslog Address and Syslog Port boxes, respectively.
6. Choose the output format from the Syslog Format menu.
7. Click **Update** and choose **Yes** on the confirmation dialog to save your configuration.

Logging Events to a Supplemental SQL Server

External logging gives you the option of creating custom report implementations directly through SQL. Using an external server can also allow you to meet forensic or compliance requirements for long-term event storage while maintaining events for a shorter period in the App Control Server database. You might also choose to implement external event logging for performance reasons.

Note several key points about what happens when external logging is activated:

- External logging does not eliminate local logging in the primary SQL Server database. Event logging continues, and saves events for whatever time period (or total number of events) you specify.
- To facilitate better system performance, event data is copied from the primary SQL Server database to the external event SQL Server database approximately every 30 seconds rather than continuously.
- Events that happened prior to your activation of external logging are not copied to the external log, so if you intend to set up external logging and want it to be comprehensive, it is best to do so at the same time you are setting up the App Control Server.
- If the external server becomes inaccessible, an error is logged, but there will be no change in App Control Server behavior. Once the external server is available again, events that were missed will be copied.

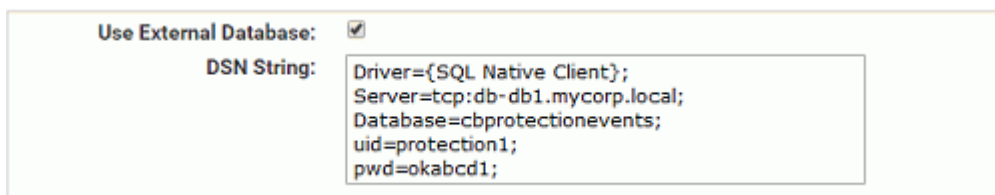
[Table 112, “External Event Logging Options,”](#) on page 738 describes each of the parameters on the External Event Logging panel of the Events tab. See the Carbon Black User Exchange or contact Carbon Black Support for additional details.

The following describes the high-level procedure for setting up external event logging to a supplemental SQL database. If you want to use NT Authentication for your external database, use the special DSN shown in the following procedure.

To enable external event logging to an additional SQL server:

1. Install SQL Server on a machine with sufficient capacity for App Control event logging. Be sure to note the information for the DSN (Data Source Name) string – this will be necessary for use in the App Control Console.
2. Run the external-events script **external_events.sql** to configure the SQL database so that it can properly store events. This script is located in the **\sql** folder under the server folder in the App Control installation directory, by default Bit9\Parity Server. It must be run on the newly installed SQL Server before you can use external events logging.
3. On the console menu, click the configuration (gear) icon and choose **System Configuration**. The System Configuration page appears.
4. Click on the **Events** tab. The External Event Logging panel appears.
5. Click the **Edit** button and then check the *Use External Database* box. This activates the **Test** button as well as the data fields on the panel.

6. In the DSN String field, enter the DSN for this database.
 - a. For manual authentication, this will include the following, each on its own line and separated by semicolons (the illustration following shows an example):
 - Driver={SQL Native Client};
 - Server=**tcp**:*yourfullyqualifiedservername\instancename*;
 - Database=**cbprotectionevents**;
 - Uid=*usernameforSQLadmin*;
 - Pwd=*password*;



Use External Database:

DSN String: `Driver={SQL Native Client};
Server=tcp:db-db1.mycorp.local;
Database=cbprotectionevents;
uid=protection1;
pwd=okabcd1;`

- b. You can use NT authentication, using the Domain credentials you supplied during App Control Server installation, for access to the external event logging server. To do this, replace the “Uid” and “Pwd” lines shown above with a “Trusted_Connection” line in the following format:
 - Driver={SQL Native Client};
 - Server=**tcp**:*yourfullyqualifiedservername\instancename*;
 - Database=**cbprotectionevents**;
 - Trusted_Connection=**Yes**;

Note

If you have difficulties with the DSN string, see the file **shepherd.dsn** in the App Control Server home directory.

7. To make sure your DSN works, click the **Test** button. If your DSN was configured appropriately, a “Testing: Success” message appears below the DSN String box. Otherwise, you will see an error message.
8. Once your DSN Test has succeeded, click the **Update** button (this replaces the “Test” button when the test is successful and you check the checkbox) and choose **Yes** on the confirmation dialog. This activates external logging.

Important

If you upgrade the server to a new version, external databases are not automatically upgraded. You must run the external-events script **external_events.sql** to configure the SQL database so that it can properly store events. This script is located in the **sql** folder under the server folder in the App Control installation directory.

Table 112: External Event Logging Options

Field	Description
Syslog Enabled	<p>A checkbox determining whether event information is output to another server for further analysis with a Syslog management tool. If checked, you also must specify a Syslog server address and listening port. This option is off by default.</p> <p>Note: See the <i>VMware Carbon Black App Control Events Guide</i> for this release for guidance on using event output with your Syslog management tools.</p>
Syslog Address	<p>IP address for a Syslog server (optional). If you specify a Syslog address, you must also enter a port for the server.</p> <p>Note: No error is reported if you set the Syslog address and/or port incorrectly. To verify that Syslog address is correctly set, confirm the receipt of events on the Syslog server after you have completed this configuration.</p>
Syslog Port	<p>Port number for a Syslog server.</p> <p>Events directed to the listening port include activity messages such as blocked files, new files on the system, and changes to login accounts.</p> <p>If you export event data, events continue to be written to the Events page, which is accessible from the App Control Console. If you specify a Syslog port, you must also enter an address for the Syslog server.</p>
Syslog Format	<p>One of the following:</p> <ul style="list-style-type: none"> • Basic (RFC3164) – this is the default for upgrades from pre-6.0.2 (Parity) versions • Enhanced (RFC5424) – this is a newer standard and the default for new installations beginning with Bit9 v7.0.1. • CEF (ArcSight) – format to use if you want to integrate App Control event logs with HP ArcSight ESM or HP ArcSight Logger • LEEF (Q1Labs) – format to use if you want to integrate App Control event logs with QRadar SIEM or QRadar Log Manager <p>See the separate document <i>VMware Carbon Black App Control Events Guide</i> for more information on syslog formats supported by App Control and how to map events to them.</p> <p>Note: If you worked with our Technical Support team to manually enable special Syslog formatting in pre-6.0.2 releases, your changes will be overwritten on upgrade to v7.2.2. Use the Syslog format menu to choose formatting.</p>
Syslog Export Process Command Lines	<p>A checkbox determining whether process command lines are included in syslog output. Not checked by default. Passwords may be specified on the command line and so sending command line output to an external server may be inappropriate.</p>

Field	Description
Use External Database	Check the box to enable use of an external SQL database. Un-check to disable reporting of events to the external database.
DSN String	The DSN string that identifies the external database you will be using. This will vary depending upon whether you use manual or NT authentication. The procedure “To enable external event logging to an additional SQL server:” on page 736 describes how to configure these choices.

To disable external event logging:

1. On the console menu, click the configuration (gear) icon and choose **System Configuration**. The System Configuration page appears.
2. Click on the **Events** tab. The External Event Logging panel appears.
3. Click the **Edit** button. This activates the data fields on the panel.
4. Click the **Use External Database** box to remove the check. This turns the “Test” button into an “Update” button.
5. Click **Update** and choose **Yes** on the confirmation dialog. External event logging is disabled.

Securing Agent-Server Communications

App Control uses SSL security for communication between its server and its agents. By default, this is based on a self-signed Carbon Black security certificate generated when the App Control Server is installed, although a different certificate can be supplied as part of the installation process.

The System Configuration **Security** tab displays the Agent Server Communications configuration page. There, you can make one or more of the following changes:

- If the current certificate for agent-server communications is self-signed, you can edit its details.
- You can import another certificate from a PKCS#12 file, either your own self-signed certificate or from a certificate authority.
- You can increase security by enabling certificate verification so that computers running the agent always verify that the correct certificate exists on the App Control Server. This is a one-time change with no reversal. You can do this for certificates from known certificate authorities, and possibly with your own self-signed certificates imported into App Control. Do not enable verification for the self-signed certificates created during the App Control server installation process.

The screenshot displays the 'System Configuration' interface with the 'Security' tab selected. The 'Agent Server Communications Security' section is active, showing the following details:

- Security Status:**
 - Certificate Source:** Self-signed, no certificate authority
 - Certificate Issuer:** APC-58610
 - Certificate Verification:** Disabled
- Current Server Certificate Details:**
 - Common Name:** APC-58610
 - Expiration Date:** Jul 13 2022 06:34:31 AM
 - Country Code:** US
 - State:** Massachusetts
 - City:** Waltham
 - Company:** Carbon Black, Inc.
 - Department:** Support
 - Email Address:** support@carbonblack.com
 - Subject Alternative Name:** [Redacted]
- Import Server Certificate From PKCS12 File:**
 - File Name:** Choose File No file chosen
 - Password:** [Empty field]
 - Import:** [Button]

Security Status

The top panel of the page shows the security status of agent-server communications. Specifically, it reports on the source of the certificate (self-signed or imported), whether there is a certificate issuer associated with the certificate, and whether App Control is configured to require that agents check the server to verify the legitimacy of the certificate. For self-signed certificates, the Certificate Issuer is the name of the App Control Server and the certificate has no known certificate authority. This panel also contains the button that enables certificate verification.

Current Certificate Details

The Current Server Certificate Details panel shows the standard details available from a security certificate. If the certificate is self-signed, you may edit the details and re-generate the certificate.

To edit the details of a self-signed communications security certificate:

1. On the console menu, click the configuration (gear) icon and choose **System Configuration**. The System Configuration page appears.
2. Click on the **Security** tab. The Agent Server Communications Security page appears.
3. In the Current Server Certificate Details panel, click **Edit**. The fields in the details panel are activated for editing, and the Edit button is replaced by Generate and Cancel buttons.
4. Change certificate details changes as you choose, then click **Generate** to generate a certificate with the new details. To cancel the changes, click **Cancel** instead.

Table 113: Agent-Server Communications Certificate Details

Field/Button	Description
Common Name	This must be the fully qualified domain name of the App Control Server to which your agents are connected.
Expiration Date/ Valid For	Shows the date and time when the certificate will expire. When you are editing the certificate details, this field changes to Valid For and provides box in which you can enter the number of days or years you want the certificate to be valid. Note: You cannot enter a Valid For period longer than 20 years or 7300 days for a self-signed certificate.
Country Code	Standard two-letter country code for organization responsible for the certificate.
State	State (if applicable)
City	City
Company	Company responsible for the certificate
Department	Department (if any) within the company
Email Address	Contact information for anyone needing more information about the certificate.
Subject Alternative Name	Subject Alternative Name (SAN) is an alternative means of verifying the certificate against the server hostname. SAN allows the use of multiple DNS names and/or IP addresses, separated by commas, for a single server so that the certificate can be verified even when there is access from different network routes or the same certificate can be used on multiple servers. The Subject Alternative Name field is empty by default. A tooltip shows the required format. The following is an example of the format for a SAN entry: DNS=cbprotection.mycorp.com, DNS=cbprotection.mycorp.local,IP=10.0.8.123 You can use wildcards in a DNS name (e.g., *.mycorp.com).

Verifying that the Server Name and Certificate Match

How the agent verifies that the server name matches the certificate depends upon the server information provided by the server certificate:

- If there are Subject Alternative Name (SAN) DNS entries in the certificate, these are compared to the server address used by the agent, and the two must match.
- If there are no SAN DNS entries, the server address used by the agent is verified against the Common Name (CN) in the server certificate and the two must match.

Mismatches in address/name format between the agent and the server certificate will fail, even if the name resolves to the IP address. For example, where the agent is using an IPv6 address and the SAN is not, verification will fail. You can correct this problem by adding an additional address (the IPv6 address) to the SAN, in the format DNS=[IPv6].

Importing a Certificate

You can import a new SSL certificate if you choose. Keep the following in mind when planning to import a certificate:

- You cannot import an expired certificate.
- Only PKCS#12 certificates are supported. You cannot use another PKCS version. To use a certificate in another format, you must convert it to a PKCS#12 file format first.
- When you import a certificate, the Edit button is removed from the Current Certificate Details panel since the imported certificate cannot be edited.
- App Control supports use of multi-level certificates. The actual certificate must be specified *last* in the PKCS#12 container file.
- Only a certificate matching the App Control Server hostname or IP address may be imported.

Note

During App Control Server installation, you must either generate a self-signed certificate or import a real certificate for the App Control Console. If you import a real certificate, you may use the same certificate for the Agent-Server communications. If you choose this option, you do not need to complete the following procedure.

To import a new certificate for agent-server communications security:

1. On the console menu, click the configuration (gear) icon and choose **System Configuration**. The System Configuration page appears.
2. Click the **Security** tab. The Agent Server Communications Security page appears.
3. In the Import Server Certificate panel, click **Browse** to navigate to the location of your new certificate file, and when you locate the file in the Chooser dialog, click **Open**.
4. Enter the Password for the certificate file.
5. When you have provided the necessary information, click **Import**. A dialog box appears describing the impact of the change.

6. To complete the certificate import, click **OK** on the confirmation dialog. A status message reports on the success or failure of the import. If successful, the new certificate is installed in the certificate repository and all fields in the Current Server Certificate Details panel are updated.

Enabling Certificate Verification

Enabling certificate verification instructs all App Control Agents to verify the authenticity of the App Control Server certificate against a Certificate Authority or their Root certificates. This adds a level of security to communications because communications between agent and server cannot be spoofed.

There are three classes of certificate that might be used for App Control communications, and you should be aware of the differences before deciding whether to enable verification:

- **Third-party certificates** – You should be able to enable and use certificate verification successfully with certificates from a known certificate authority assuming they are valid and up to date.
- **Imported self-signed certificates** – You also should be able to enable and use certificate verification successfully with your own imported, self-signed certificates assuming they are valid and up to date.
- **Self-signed certificates created during installation** – The self-signed certificates generated by the App Control Server installation program are not from a known certificate authority, so certificate verification should never be used in that case. The **Enable Certificate Verification** button is not exposed when App Control detects a self-signed certificate that it created. In order to expose the "Enable Certificate Verification" button, you would need to contact Carbon Black Support to understand how to override this safety measure.

Caution

Once certificate verification is enabled, it cannot be revoked, so be certain you have the certificate you want in place and you are sure you want to implement the feature *before* you click the button.

To enable agents to verify the server communication certificate:

1. On the console menu, click the configuration (gear) icon and choose **System Configuration**. The System Configuration page appears.
2. Click the **Security** tab. The Agent Server Communications Security page appears.
3. Make any changes you intend to make to the certificate, whether it is editing the details of a self-signed certificate or importing a new one from a file.
4. In the Security Status panel, click **Enable Certificate Verification**. If you are sure you want to make this change, click **OK** in the confirmation dialog; this cannot be undone in the App Control Console. When you click OK, the Enable Certificate Verification button disappears, and the Certificate Verification field changes to *Enabled*.

Advanced Configuration Options

The Advanced Options tab on the System Configuration page includes options related to database backup, computer and agent management, certificate and updater rules, and general console management. It may also include settings for optional features.

For information about Database Backup options, including backup and restore instructions, see [“Backing Up the App Control Server”](#) on page 750 and [“Restoring the App Control Server”](#) on page 753.

This section provides a basic description of the other Advanced Options. [Table 114](#) describes the parameters on this page, except for the Database Backup parameters, which are described in the sections referenced above.

To view and edit Advanced configuration options:

1. On the console menu, click the configuration (gear) icon and choose **System Configuration**. The System Configuration page appears.
2. Click the **Advanced Options** tab. The Advanced Options configuration page appears (see next page).
3. If you need to change any of the configuration information, click **Edit** and make any changes necessary.
4. To submit changes, click the **Update** button and click **Yes** on the confirmation dialog.

General Events Security **Advanced Options** Mail Licensing External Analytics Connectors Unified Management SAML Login

Advanced Options

Database Backup

Backup Type: Network ▾

Backup Path: (none)

Username: (none)

Password:

Windows Domain: (none)

Enabled:

Status: Idle

Carbon Black App Control Agent

Automatic Agent Upgrades: Disabled ▾

Full OS Inventory Tracking:

- Track inventory for locally approved support files signed by "Microsoft Windows" or "Microsoft Corporation" publishers
- Discard information about locally approved support files signed by "Microsoft Windows" or "Microsoft Corporation" publishers at server
- Discard information about locally approved support files signed by "Microsoft Windows" or "Microsoft Corporation" publishers at agent

Resource Download Location: https://**ServerIP**/hostpkg/pkg.php?pkg=

Carbon Black App Control Console

Log Users Out After: 120 Minutes

Files To Ignore: (none)

API

API Access Enabled: ●

File Uploads

Delete Uploaded Files After: after 4 Week(s)

Default Upload Location: files\ Test

Old Computer Cleanup

All Computers: delete after 30 Day(s) offline

Table 114: Advanced (Configuration) Options

Section:Field	Description
Database Backup	See “Backing Up the App Control Server” on page 750 for a description of these options.
App Control Agent: Automatic Agent Upgrades	When Enabled , App Control Agents are notified when a new agent version is available, <i>if</i> the Policy the agent is a member of also has agent upgrades activated. It normally is Disabled and is for use during a App Control Server upgrade. It has no effect on a new App Control Server installation. See the <i>VMware Carbon Black App Control Installation Guide</i> guide for full instructions on agent upgrades.
App Control Agent: Full OS Inventory Tracking	If the <i>Track inventory for locally approved support files signed by “Microsoft Windows” or “Microsoft Corporation” publishers</i> radio button is selected, all files from Microsoft are tracked in the file inventory for this server. If either of the other radio buttons is selected, locally approved support files whose publisher is “Microsoft Windows” or Microsoft Corporation” are excluded from tracking in the App Control database, which can significantly reduce the load on the server. See “Excluding Tracking of Microsoft Support Files” on page 230 for details of the exclusion options.
App Control Agent: Resource Download Location	This field allows you to change the location from which agents and configuration files are downloaded. Note: If you change this setting, you must restart the server to make it take effect.
App Control Console: Log Users Out After	Time period of no activity after which a user is automatically logged out the App Control Console.
App Control Console: Files to ignore	Files that you want to exclude from the Files page lists, separated by commas with optional wildcard character (*). Events associated with ignored files still appear in the Events table and can trigger alerts. Ignored files can be located as Find Files results. Not normally used in normal App Control Server operation.
API	If <i>API Access Enabled</i> is checked, the App Control APIs are made available on this server. APIs allow access to the App Control Server and its database via automation and scripting using a variety of languages. See Appendix B, “App Control API,” for details.
File Uploads	(Optional) Settings for the separately licensed feature for uploading files from agent computers. Determines the location to which files are uploaded and the length of time they remain on the server before deletion. See “Uploading Files from Agents” on page 895 for more details.

Section:Field	Description
Old Computer Cleanup: All Computers	<p>Period of time offline after which <i>any disconnected computer</i> is deleted from the list of computers managed by this App Control Server. Check the box to activate cleanup, and enter the number of days offline after which a computer will be deleted.</p> <p>If you reconnect a deleted computer and the computer is still running App Control Agent, the computer will resync its file list and return to its last configured policy (if available) or the Default Policy. See “Deleting Computers” on page 173 for more details.</p>
Old Computer Cleanup: Computers Matching Filter	<p>A filtered version of automatic deletion of computers from the list of App Control-managed computers after a certain period of time. Check the box to activate cleanup, and enter the number of days offline after which a computer will be deleted.</p> <p>You also add one or more filters to limit deleted computers to those matching criteria you specify. For example, you can delete only virtual computers when they reach the time limit. Or you can delete all computers matching a particular tag (e.g., “Visitor”). The filter options are:</p> <ul style="list-style-type: none"> • Computer name • Computer tag • IP Address • Identifier (MAC address) • Parent Template • Platform • Policy • Virtualized • Virtual Platform <p>Computers must match all filter criteria to be deleted.</p>
Software Rule Options: Updaters	<p>If <i>Automatically update application updaters from Carbon Black File Reputation</i> is checked, Carbon Black File Reputation keeps the Updaters list in the Software Rules section on your App Control Server up-to-date with any new versions it confirms.</p> <p>If not checked, the updaters listed continue to be those provided at server installation time, supplemented by any updaters you have manually defined.</p>
Software Rule Options: Event Rules	<p>If <i>Process event rules</i> is checked (the default), events matching rules defined and activated on the Event Rules page can trigger actions such as file analysis or file banning. See “Event Rules” on page 520 for more details.</p>
Software Rule Options: Indicator Sets	<p>If <i>Automatically update from Carbon Black File Reputation</i> is checked (the default), Carbon Black File Reputation keeps the Indicator Sets used for threat detection up-to-date. See Chapter 23, “Advanced Threat Detection,” for more on Indicator Sets.</p>
Software Rule Options: Health Indicators	<p>If <i>Automatically update from Carbon Black File Reputation</i> is checked (the default), Carbon Black File Reputation downloads Health Indicators used to monitor and report on system health and updates them when necessary. If not checked, the System Health feature is not available. See Chapter 28, “Monitoring System Health,” for more on health indicators.</p>

Section:Field	Description
Certificate Options: Expired Certificates	<p>If <i>Allow approval of software with expired certificates</i> is checked, an expired certificate may be used for publisher-based approval of a file, if the certificate was valid and the certificate timestamp is within the period during which it was valid. See “Approval with Expired Certificates” on page 288 for more details.</p> <p>If not checked, software with expired certificates cannot be approved by publisher.</p>
Certificate Options: Exclude Publisher Approvals With These Certificate Algorithms	<p>This option determines which certificates are <i>excluded</i> from use for publisher approvals. If the box for a certificate algorithm is checked, files signed by a publisher whose certificate uses that algorithm cannot be approved by publisher. See “Excluding Certificate Algorithms” on page 288 for more details.</p> <p>The options are:</p> <ul style="list-style-type: none"> • MD2RSA • MD5RSA • SHA1RSA • SHA256RSA
Certificate Options: Minimum Certificate Key Size For Approval	<p>This option specifies a minimum key length for a certificate to be used for file approval by publisher. Certificates whose key size is greater than or equal to the chosen value may be used for approval by publisher. Certificates whose key size is smaller than the chosen value may not be used. The default value is 512. See “Minimum Key Size” on page 288 for more details.</p>
Certificate Options: Digital Signatures	<p>If <i>Require countersignature</i> is checked, a countersignature is required for the digital signature of each certificate used to identify a publisher. See “Countersignature Options” on page 288 for information that may assist you in configuring this option.</p>
Certificate Options: Initial Revocation Check	<p>Determines whether and how a certificate revocation check is done at initial file discovery on an agent. There are three possible values:</p> <ul style="list-style-type: none"> • Network – If revocation information is not locally available then use the network to retrieve a certificates revocation status. • Cache – Use locally available revocation status information when performing certificate revocation (the network will not be used). • None – Do not perform certificate revocation checking. <p>Consider your agent deployment scenario when setting these values since they can impact agent performance. See “Revocation Checks” on page 289 for more details.</p>

Section:Field	Description
Certificate Options: Background Revocation Check	<p>Determines whether and how certificate revocation checks are done for existing files on an agent every 24 hours. If activated, these checks are done in the background. The possible values are the same as those for Initial Revocation Check (above). See “Revocation Checks” on page 289 for more details.</p> <p>Note: Certificate revocation checks are also done by the <i>server</i>, generally on a weekly basis. Server-based revocation checks are not affected by the initial or background revocation check settings. If you are monitoring network connections, be aware that some traffic you see will be for these revocation checks, and could involve sites in a variety of countries.</p>
Login Banner: Display Login Banner	This checkbox enables and disables display of a custom banner on the login page when users log in to App Control. You must provide the text for this banner in the field below.
Login Banner: Banner Text	<p>The text displayed in the custom login banner when enabled. In addition to Unicode text characters, you can use the following markup symbols:</p> <p>
 for a line break</p> <p><p> and </p> to begin and end paragraphs</p> <p> and to bold and unbold text</p>
Login Banner: Font Color	The font color for the banner on the login page. This can be in any valid CSS format or six hexadecimal digits. Default is 333333 (a dark gray). See https://www.w3.org/TR/2018/REC-css-color-3-20180619/ for CSS color specifications.
Login Banner: Background Color	The background color for the banner on the login page. This can be in any valid CSS format or six hexadecimal digits. Default is D3D3D3 (a light gray).
Login Banner: Border Color	The border color for the banner on the login page. This can be in any valid CSS format or six hexadecimal digits. Default is 000000 (black).

Adding a Login Banner to the Console

If regulatory requirements or your organization’s policies require that users see special text when they access your information resources, this can be configured for the App Control console login page. You can define the text that appears as well as the colors of the text, background, and outline of the text box that appears.

To configure and enable a special login banner:

1. On the console menu, click the configuration (gear) icon and choose **System Configuration**. The System Configuration page appears.
2. Click the **Advanced Options** tab and click Edit at the bottom of the page.
3. In the Login Banner panel, enter the text you would like to have appear in the Banner Text box. You can use bold, linebreak, and paragraph markup in this box.
4. Choose the colors you would like to use for each element in the banner:

- Font Color
 - Background Color
 - Border Color
5. When you have configured the login banner, check the **Display Login Banner** checkbox and then click the **Update** button at the bottom of the page.
 6. Log out to view your login banner. If changes are necessary, you can log in again and return to the Advanced Options page to edit the text or colors.

You might want to disable the login banner, either temporarily or because your requirements have changed and it is no longer necessary.

To disable a login banner:

1. On the console menu, click the configuration (gear) icon, choose **System Configuration**, and then click the Advanced Options tab on the System Configuration page.
2. At the bottom of the page, click the **Edit** button.
3. In the Login Banner panel un-check the **Display Login Banner** checkbox and then click the **Update** button at the bottom of the page.

The banner will no longer appear when users log in to the App Control console, but it can be re-enabled simply by following these same steps and checking the Display Login Banner checkbox.

Backing Up the App Control Server

If your SQL Server administrator has a standard backup plan and mechanism, Carbon Black recommends that you use that mechanism to backup your App Control database.

If you do not have or choose not to use a separate database backup mechanism, App Control Server provides a mechanism to fully back up and restore the system as currently configured, including computer configuration, system settings, file database, and event log. The built-in backup mechanism backs up all database changes within 6 hours of a critical change, such as a change in policy. Full backups occur once a day. Continuous automated backups ensure that the server and connected computers remain synchronized after you restore your backup configuration.

The free space available to the backup folder should be at least twice the size of the App Control Server database. For both your backup folder and your main SQL database, monitor disk space regularly to prevent overruns.

The App Control Server Backup function requires that **xp_cmdshell** support be enabled on the SQL Server instance where the database is hosted. See your SQL Server documentation for instructions on enabling xp_cmdshell. The following links provide some information about this task:

- SQL Server 2008: <http://www.mssqltips.com/sqlservertip/1673/where-is-the-surface-area-configuration-tool-in-sql-server-2008/>
- SQL Server 2012: <http://msdn.microsoft.com/en-us/library/ms190693.aspx>

Important

Because enabling `xp_cmdshell` has security implications, the SQL Server administrator at your site should follow all best practices to limit any exposure it creates. This includes, but is not limited to, the following:

- Never grant access to non-sysadmin principals.
- Ensure that the sysadmin SQL Server right is granted only to trusted administrators of the SQL Server system.

If you stop using the built-in backup mechanism, disable `xp_cmdshell`.

To use the App Control Server database backup mechanism:

1. Make sure `xp_cmdshell` is enabled on your SQL Server.
2. On the console menu, click the configuration (gear) icon and choose **System Configuration**.
3. Click the **Advanced Options** tab. The Advanced Options page appears, with the Database Backup panel at the top.
4. Click the **Edit** button at the bottom of the page, and specify backup location and configuration options (see [Table 115](#)):

The screenshot shows the 'Advanced Options' configuration page in the App Control Server console. The 'Database Backup' section is visible, with the following settings:

Backup Type:	Network
Backup Path:	(none)
Username:	(none)
Password:	*****
Windows Domain:	(none)
Enabled:	<input type="checkbox"/>
Status:	Idle

5. Click the **Update** button and then click **Yes** on the confirmation dialog. Each time you save the backup configuration with backup enabled, the App Control Server tests backup settings and displays an error message if the configuration fails. The server also writes messages to the Events page to inform you about backup success, problems, or failure.

Table 115: Database Backup Options

Field	Description
Backup Type	Network or Local. Local backups should only be used on a different physical drive than the App Control Server drive.
Backup Path	<p>The full path to the computer or storage media that will store the database/configuration backup. Secure the backup directory and ensure that only App Control Server administrators have access to it. For best performance, avoid creating unnecessary subdirectories and keep the backup directory as close as possible to the server root directory. For example:</p> <pre>\\server_name\cbprotection_backup</pre> <p>Notes:</p> <ul style="list-style-type: none"> Local paths are recommended for local backups. You may use a UNC path (shown above) for a Local drive, but the local option does not include username, password, or Windows domain information and no privileges are used to access this path. If the App Control Server is connected to a remote database, the backup path you provide is relative to the database server, and the Username, Password, and Windows domain fields will not appear.
Username (Network backups)	User name with write permission to the network backup directory.
Password (Network backups)	Domain password for the user account that writes to the network backup directory. This password is encrypted in the database.
Windows domain (Network backups)	Windows domain to which the user account for the network backup location belongs.
Enabled	<p>Check the box to begin backups at two-minute intervals to the specified storage location.</p> <p>Clear the checkbox to discontinue automatic backups.</p>
Status (read only)	Time of the next scheduled backup, or status of the most recent backup (including any errors).

Important

After you configure the backup directory, do not add, delete, or edit any of its files. Because updating is continuous, such changes adversely affect file synchronization and the integrity of your backup.

Restoring the App Control Server

You can restore the App Control Server to its most recent state. App Control database and settings restoration is a manual procedure that requires reinstalling the App Control Server. As a precaution, the restoration procedure disables automatic backups to ensure that your only backup copy is not overwritten before you can copy it to a safe location.

The App Control Agent runs independently of the App Control Server. While you reinstall App Control Server and restore the backup configuration, computers remain protected according to the configuration settings received from the App Control Server during their last polling instance.

To restore the App Control Server to its most recent configuration:

1. If your Windows installation is corrupted, reinstall the operating system on the App Control Server hardware. See the *VMware Carbon Black App Control Installation Guide* guide for installation guidelines.
2. Reinstall the App Control Server:
 - a. Insert the App Control CD (or an executable image of it) in a drive connected to the designated server.
 - b. To run the installer, follow the installation prompts. See the *VMware Carbon Black App Control Installation Guide* guide for information about installation options, including changing the server IP address, installing via terminal services, or using a DNS name.
 - c. On the Install Type Option screen, select the **Restore from backup** option.
 - d. Navigate to the backup directory.
 - e. Follow the remaining standard installation prompts, and after completing the installation, exit the procedure.

Important

When you reinstall, the IP address of the installation computer is detected. If you installed the App Control Server using a DNS name, you can sometimes reinstall on a computer with the same name but a different IP address. Otherwise, if you are reinstalling on a computer with a different IP address, you must also reinstall the agent on all computers. After installation, computers reinitialize their files and locally approve previously Unapproved files. The restore procedure automatically updates existing agent installation packages to use the new server IP address.

3. During the restoration procedure, continuous backups are automatically disabled. Resume automatic backups as follows:
 - a. Copy all files in the backup folder to a new location so they are not overwritten (or specify a new backup folder and leave existing backup files in place).
 - b. Verify that the currently specified backup directory is now empty so that the fresh backup completes without potential corruption by old files.
 - c. On the App Control Console menu, click the configuration (gear) icon and choose **System Configuration** and then click the **Advanced Options** tab. The Database Backup panel is at the top of the Advanced Options page:

The screenshot shows the 'Advanced Options' tab in the System Configuration interface. Under the 'Database Backup' section, the following settings are visible:

- Backup Type:** Network (dropdown menu)
- Backup Path:** (none)
- Username:** (none)
- Password:** (masked with asterisks)
- Windows Domain:** (none)
- Enabled:**
- Status:** Idle

- d. Select Edit and then check the **Enabled** check box.
- e. To commence backups in the specified location, click the **Update** button at the bottom of the page and then click **Yes** on the confirmation dialog.

Configuring Alert and Approval Request Mail

Some App Control features require configuration of a mail server so that messages can be sent to administrators or endpoint users under certain conditions. The current features that require this are:

- **Alerts** – email notification of administrators when a App Control alert is triggered. See [“Creating Alerts”](#) on page 612 for more information about alerts.
- **Approval Requests** – email notification of a user when their Approval Request is closed. See [“Reviewing and Resolving Requests and Justifications”](#) on page 572 for more information about Approval Request responses.

To enable email notifications, you must give the App Control Server access to an SMTP (Simple Mail Transport Protocol) server to send messages when notification conditions are met. You configure this on the Mail tab of the System Configuration page. There, you can:

- Specify the mail server for notifications.
- Choose standard or secure mail for notifications.
- Enable or disable sending of alert mail to subscribers of specific alerts.
- Specify an optional global subscriber to receive all alert emails.
- Enable or disable automatic delivery of approval request response email.

[Table 116](#) describes all fields for these options.

Table 116: Mail Configuration settings

Panel:Field	Description
Alert Settings: Mail Notification Enabled	Checkbox determining whether email subscribers to App Control alerts receive email when the alerts are triggered. You might choose to disable this if you are monitoring alerts closely on the App Control Console, or are generating a large number of alerts during testing or monitoring activities. Enabled by default.
Alert Settings: Global Subscriber Enabled	Checkbox determining whether a <i>global</i> subscriber to email alerts is enabled. If this is enabled and a subscriber is entered in the Global subscriber field, the subscriber receives email every time any App Control alert is triggered. You can enable or disable this as needed.
Alert Settings: Global Subscriber	Email address of the global alert subscriber. Appears only if Global Subscriber Enabled is checked.
Approval Request Settings: Mail Notification Enabled	Checkbox determining whether the user making an Approval Request receives automatic email when the request is closed. Disabled by default.
Server Settings: Mail Server	Mail server address. This can be an IP address or a fully qualified domain name.
Server Settings: Mail Server Port	Port for the mail server. Specify the port in use for your server. Default value of 25 is used for standard SMTP mail; default value of 587 is used for Secure Mail. Make sure the port you use is available for outbound traffic.
Server Settings: Mail "From" Address	Address used as the <i>from</i> address in notification emails. The <i>from</i> address need not be an actual, functioning email address, but it must be in the proper syntax for an email address (e.g., info@mycorp.com) or it will generate event log errors. Also, some mail servers automatically discard email without a proper <i>from</i> address as spam.
Server Settings: Secure Mail (TLS)	Checkbox determining whether email is sent by secure mail. Secure mail requires a username and password to authenticate communication with the email server. Note: You cannot use Secure Mail with an account that requires multi-factor authentication. Use of such an account will cause App Control notifications to fail.
Server Settings: Secure Mail Username	Username for authenticating access to the mail server. Appears only if Secure Mail (TLS) is checked.
Server Settings: Secure Mail Password/ Confirm Password	The password for authenticating access to the mail server. Must be entered in both password fields. Appears only if Secure Mail (TLS) is checked.

Panel:Field	Description
Validate Server: Test Address	Email address used to test your email server configuration. For example, you can use your own email address so that you can click the Send Mail button and immediately know whether the mail server configuration works. Do the test before the settings on this page are updated so that any issues are exposed and can be remedied.

Configuring Standard Email for Notifications

To configure email using standard (unsecure) mail:

1. On the console menu, click the configuration (gear) icon and choose **System Configuration**. The System Configuration page appears.
2. Click the **Mail** tab. The Mail Notification Configuration table appears:

The screenshot shows the 'Mail Notification Configuration' page. At the top, there are tabs for 'General', 'Events', 'Security', 'Advanced Options', 'Mail', 'Licensing', 'External Analytics', 'Connectors', 'Unified Management', and 'SAML Logi'. The 'Mail' tab is selected. The page is divided into several sections:

- Alert Settings:** 'Mail Notification Enabled' is checked (checkbox with blue checkmark). 'Global Subscriber Enabled' is unchecked (checkbox).
- Approval Request Settings:** 'Mail Notification Enabled' is unchecked (checkbox).
- Server Settings:** 'Mail Server' is an empty text field. 'Mail Server Port' is a text field containing '25'. 'Mail "From" Address' is an empty text field. 'Secure Mail (TLS)' is unchecked (checkbox).

Below the Server Settings, there is a note: 'Please validate settings by sending a test mail before updating the Carbon Black App Control Server.' Below this note is a 'Validate Server' section with a 'Test Address' text field and a 'Send Mail' button with a play icon. At the bottom of the page, there are three buttons: 'Edit' (with a pencil icon), 'Update' (with a save icon), and 'Cancel' (with a red X icon).

3. Click the **Edit** button to activate the email configuration fields for editing. Fields are added or removed depending upon the options you enable or disable. When you enable an option, required fields for that option appear in red if not filled in.
4. The Alerts Settings Mail Notification Enabled box is checked by default. Leave it checked if you want alert notification emails to be sent.
Note: See [“Specifying a Global Alert Subscriber”](#) on page 758 before deciding whether to enable a global subscriber.

5. Check the Mail Notification Enabled box in the Approval Request Settings panel if you want automatic email to be sent a requestor when an approval request is resolved.
6. In the Server Settings panel, enter the Mail Server address, either as a fully qualified domain name or IP address.
7. By default, the Mail Server Port defaults to 25 when you use standard mail. If you are using a different port, change the field.
8. Enter a Mail “From” Address. This is the address that recipients will see as the sender of notification email.
9. If you want to use Secure Mail for notifications, provide the information described in [“Configuring Secure Email for Notifications”](#) on page 757.
10. To test the mail server configuration, enter a Test email address at which you can receive mail and click **Send Mail**. The App Control Console sends a test email to that address.
11. If the mail server configuration test reported an error in the Validate Server section, correct the problem. Wait for the Validate Server test to be successful before you proceed.
12. Click the **Update** button and then click **Yes** on the confirmation dialog. The updated mail configuration is displayed on the Mail Notification Configuration page.

Configuring Secure Email for Notifications

App Control provides the option of using secure mail instead of standard mail for notifications from the server. The secure mail requires a username and password for access to the mail server. Secure mail uses Transport Layer Security, which is an explicit method of securing communication to the mail server. By default, it uses port 587 and initiates the communication with **–BEGINTLS** sent in plain text.

Important

You cannot use Secure Mail with an account that requires multi-factor authentication. Use of such an account will cause App Control notifications to fail.

To configure the App Control Server to use SMTP/TLS for notifications:

1. In the console menu, click the configuration (gear) icon and choose **System Configuration** and then click on the **Mail** tab. The Mail Notification Configuration page opens.

2. Click **Edit** and check the Secure Mail (TLS) box. Secure mail options appear.

The screenshot shows the 'Server Settings' panel with the following fields and options:

- Mail Server:** [Empty text box]
- Mail Server Port:** [587]
- Mail "From" Address:** [Empty text box]
- Secure Mail (TLS):**
- Secure Mail Username:** [Empty text box]
- Secure Mail Password:** [Empty text box]
- Confirm Password:** [Empty text box]

Please validate settings by sending a test mail before updating the Carbon Black App Control Server.

3. If you have not already done so, provide the Mail Server and Mail "From" Address.
4. By default, the Mail Server Port defaults to **587** when you choose Secure Mail. If you are using a different port, change the value in this field.
5. In the Security Mail Username field, provide a username for authentication on the secure mail server.

Notes

- For an Exchange Server, the Username should be in the format DOMAIN\username, and the From address field must contain a user email return address.
- For Gmail, the Username should contain the Gmail username without any domain. The value in the From address is ignored.

6. In the Secure Email Password field, enter the password for the mail server username, and enter it again in the Confirm Password field.
7. In the Validate Server panel, enter a **Test Address** and test your mail server settings by clicking on **Send Mail**. If the configuration is valid, a message appears that confirms that the test mail was sent. Check that the mail was received at the address specified.

The screenshot shows the 'Validate Server' panel with the following elements:

- Test Address:** [Empty text box]
- Send Mail:** [Button with a play icon]

8. When you have confirmed that the email was received as specified, click **Update** to save the configuration, review the changes on the confirmation dialog, and click **Yes** if you are satisfied with the changes.

Specifying a Global Alert Subscriber

You can designate one user as the global alert subscriber. Because this has the potential to generate a large amount of mail for that user, think carefully before enabling this feature, and consider a special address dedicated to alert tracking. You enable the global subscriber in the Mail Notification Configuration panel of the System Configuration page.

To enable one subscriber to receive all alert emails:

1. On the console menu, click the configuration (gear) icon and choose **System Configuration**.
2. Click the **Mail** tab. The Mail Notification Configuration page appears.
3. In the Settings panel, click **Edit**.
4. Check the **Global Subscriber Enabled** box. The Global Subscriber text box appears.
5. In the Global Subscriber text box, enter the name of the subscriber.
6. Click the **Update** button and then click **Yes** on the confirmation dialog.

Note

To disable the global subscriber, *un-check* the Global Subscriber Enabled box and then **Update**.

Managing App Control Licenses

The Licensing panel of the System Configuration page provides the ability to manage App Control licenses and to activate, deactivate, and configure Carbon Black File Reputation. The Carbon Black File Reputation options are described in the section [“Activating Carbon Black File Reputation”](#) on page 763.

App Control can be licensed at two feature levels:

- **Visibility** – Enables all of App Control’s file and event tracking and reporting capabilities, but does not include control features such as file bans and device blocking.
- **Suite** – Enables both Visibility and Control features.

License keys determine the number of agents allowed to run in each mode. You can mix licenses on the same server, having, for example, 20 Visibility licenses and 20 Suite licenses. In addition, you can purchase the Control upgrade at any time to bring the Visibility licenses up to Suite level. Some optional features are also controlled by the license key.

Viewing Your App Control License Limits and Use

The Licensing panel of System Configuration shows the licenses you have at each level, allows you to add new licenses, and shows how many licenses of each type are in use. It also might show that optional or custom features are activated. For example, if you have licensed the App Control Connector, it is shown here.

To view the App Control Licensing configuration page:

1. On the console menu, choose **Administration > System Configuration**. The System Configuration page appears.
2. Click the **Licensing** tab. The Licensing options appear:

System Configuration

General Events Security Advanced Options Mail **Licensing** External Analytics Connectors Unified Management SAML Login

Licensing

Summary

Carbon Black App Control Suite license Limit: 100 In use: 1
 There is 1 computer currently in Control policies.
 The following additional features have been enabled: File Uploads, Connectors

Licenses

Paste license key Specify license file

In the Licensing window, the Summary panel shows the following information:

- **App Control Suite license** shows the **Limit** for the number of agents (if any) you are licensed to run under full Control mode and the number of these licenses currently **In use**.
- **App Control Visibility license** shows the **Limit** for the number of agents (if any) you are licensed to run under Visibility mode only and the number of these licenses currently **In use**.
- **There are x computer(s) currently in Visibility policies** and **There are y computer(s) currently in Control policies** not only show the number of systems you currently have in each mode but also provide access to a list of each. When you click the highlighted number in each line, the Computers Page opens showing only the computers in the category you clicked. For example, in the illustration above, clicking on **40** shows a list of computers in Control policies. This line also shows how many computers managed by this App Control Server are *servers*.
- If your current license includes optional features, these will also be shown in the Summary panel.

Notes

- App Control licenses specify the allowable number of agents (computers) in each category; licenses are not locked to particular agents. The number of agents actually operating at each level is controlled by the Mode setting on the Add/Edit Policy page for the policy controlling the agent. You can move a computer or group of computers from Visibility mode to Control mode, or vice versa, as long as you have a sufficient number of Suite licenses for the systems in Control.
- For agents in Visibility mode policies, Visibility Only licenses are used first, up to the number you purchased (if any), and then, if necessary, Suite licenses are used.

App Control server administrators can also see licensing information on the App Control Console Home Page if the Licensing portlet is displayed. This portlet provides a **Manage your licenses** link that takes you to the Licensing configuration page.

License Warnings

When you create or edit a policy, or add computers to it, you may change the number of licenses of each type you are using. If the number of agents in Control mode exceeds the number of Suite licenses you have, the console displays a warning message. A warning also appears if the total number of agents exceeds the total number of licenses. If you see one of these warnings, take one of the following actions:

- Contact your Carbon Black Sales representative to purchase additional licenses.
- Move enough agents out of Control policies to comply with your Suite license limit. You can accomplish this by either moving some of your computers to a different policy or by changing one or more policies to Visibility mode.
- Move enough agents to Agent Disabled mode (and uninstall the agent if you do not plan to acquire more licenses) to comply with your license limits.

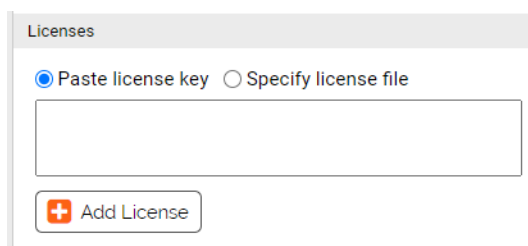
Adding Licenses

If you acquire a license key for additional agents at either licensing level, you activate the new license on the Licensing page. There are two ways to add a new license:

- by entering a string of characters in a text box
- by identifying the location of a file containing the license key

To add new App Control licenses by entering the key:

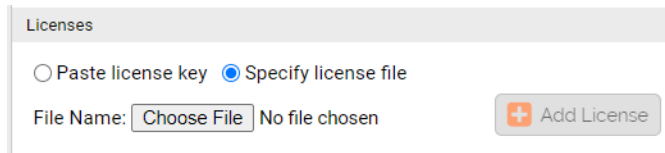
1. On the console menu, click the configuration (gear) icon and choose **System Configuration**. The System Configuration page appears.
2. Click the **Licensing** tab. The Licensing options appear.
3. In the Licenses panel, click the **Paste license key** radio button.



4. Paste or type the license key you received from Carbon Black in the text box.
5. Click the **Add License** button.

To add new App Control licenses by filename:

1. On the console menu, click the configuration (gear) icon and choose **System Configuration**. The System Configuration page appears.
2. Click the **Licensing** tab. The Licensing options appear.
3. In the Licenses panel, click the **Specify license file** radio button.



4. Click the **Choose File** button to open the file chooser, locate the license file, and click **Open** in the file chooser.
5. Click the **Add License** button.

Confirming License Addition

If your license addition is successful, the following message will display within the Add License panel: "App Control License has been successfully added."

If your license addition is unsuccessful, the following message will display: "App Control License has not been added:" along with information about why the addition was unsuccessful. Correct the problem if possible; otherwise, contact your Carbon Black Support representative.

Activating Carbon Black File Reputation

Carbon Black File Reputation is a web service that provides features to enhance the value of the App Control Server. Enabling Carbon Black File Reputation can provide:

- The Carbon Black File Reputation service, which helps identify and classify software discovered on your computers by comparing it to an extensive database of known files. It provides a threat level and a trust rating to files in its database.
- Carbon Black access to your server for remote diagnostics and troubleshooting.
- Cloud-based updates to Trusted Updaters and Advanced Threat Indicators.

While most features are enabled by default when you activate Carbon Black File Reputation, you can opt in or out of feature groups.

Important

- The Carbon Black File Reputation (CDC), which provides file trust and threat information and allows automatic updates of certain rules, requires a TLS 1.2 connection from the App Control Server. If you intend to connect to the CDC, use of .NET 4.6 (or later) is recommended. Earlier versions of .NET will default to pre-TLS-1.2 protocols, and this will prevent a CDC connection unless you disable those older protocols. Disabling older TLS/SSL protocols may be a security issue for connections to other services from your App Control Server.
- If your App Control license key included a Carbon Black File Reputation subscription, the key for Carbon Black File Reputation will already appear on the Licensing page. You will still need to follow the procedure below to accept the terms and conditions of Carbon Black File Reputation use and activate the service.

To enable and configure Carbon Black File Reputation:

1. On the console menu, click the configuration (gear) icon and choose **System Configuration**. The System Configuration page appears.
2. Click the **Licensing** tab. The Licensing configuration options appear, with the Carbon Black File Reputation Activation and Proxy Settings panels at the bottom of the page.

Carbon Black Collective Defense Cloud Activation

Your subscription to Carbon Black Collective Defense Cloud is currently activated.

Carbon Black Collective Defense Cloud Key:

Activate

Carbon Black Collective Defense Cloud Proxy Settings

Enabled:

URL: (none)

Example: `http://hostname_or_ip[:port]`

- If you want to use a Proxy Server to communicate with Carbon Black File Reputation, go to the Carbon Black File Reputation Proxy Settings panel, click **Edit**, and configure the settings as described in the table below: See [“Using a Proxy Server for Carbon Black File Reputation”](#) on page 766 if the proxy server requires authentication.

Table 117: Carbon Black File Reputation Proxy Settings

Field/Button	Description
Proxy Settings: Enabled	If checked, use of a proxy server for communication with Carbon Black File Reputation is enabled. You must provide its URL in the URL box.
Proxy Settings: URL	The URL to use as proxy for Carbon Black File Reputation communications. You can use a hostname or an IP address, and optionally add a port specification.

- Click **Update** and then click **Yes** in the confirmation dialog.
- If there is already a Carbon Black File Reputation key showing in the Carbon Black File Reputation Activation box, skip to the next step.

- or -

If the Carbon Black File Reputation key field is empty, enter the key you have or contact your Carbon Black Support representative to get an activation key.

Carbon Black Collective Defense Cloud Activation

Your subscription to Carbon Black Collective Defense Cloud is currently activated.

Carbon Black Collective Defense Cloud Key:

Activate

Note: Connectivity between the browser and the Carbon Black File Reputation site is required for the remainder of the steps in this procedure.

- When a Carbon Black File Reputation key is showing, click **Activate**. The Activation panel of the page is updated with new buttons.
- Click the **Accept Terms and Activate** button. The Carbon Black File Reputation Terms and Conditions page appears in a new browser window.
- Review the Carbon Black File Reputation terms and conditions. If you agree, check the box to confirm that you have read the terms and click the **Submit** button. This activates your subscription and enables you to connect to Carbon Black File Reputation.
- Close the Carbon Black File Reputation Activation browser window and return to System Configuration in the App Control Console.
- Click the **Verify Activation** button to determine whether Carbon Black File Reputation was successfully configured for communication with your App Control Server.
- Click the **Options** button, which appears after you complete the activation, to open a web page that allows modification of certain Carbon Black File Reputation parameters. The options include the following checkboxes (note which are enabled by default):

- **Enable file metadata sharing for Reputation and Threat results from Carbon Black File Reputation** – This enables transmission of file metadata (but not file content) collected from your agents to Carbon Black File Reputation for analysis. This option is enabled by default, and keeping it enabled is required for you to have access to the reputation services provided by Carbon Black File Reputation.
 - **Enable remote diagnostic analysis by Carbon Black Support** – This enables transmission of diagnostic data and aggregate usage information from your App Control Server to be sent to Carbon Black on an ongoing basis to ensure optimal performance. This is enabled by default.
 - **Enable direct file transfer to Carbon Black Support for troubleshooting** – This allows any files placed in the App Control Server support directories to be sent to Carbon Black Support, including log and agent cache files. This helps the support team respond to questions and issues you report about your installation. This option is *not* enabled by default.
 - **Enable automatic updates of Trusted Updaters and Advanced Threat Indicators** – This allows Carbon Black to remotely update or add trusted Updaters and advanced threat indicators (for detection) on your App Control Server. This option is enabled by default.
 - **Enable Health Indicators** – This allows Carbon Black to remotely deliver Health Indicators, which monitor and report on the health of your App Control environment. It also allows updates to existing health indicators. This option is enabled by default. See [Chapter 28, “Monitoring System Health,”](#) for more details on this feature.
- 12.** Examine all of the Carbon Black File Reputation Options – there may be added or changed options when you are enabling the connection to your App Control. If you are unclear on what any option does, contact Carbon Black Support for more information. When you know which options you want to enable or disable, click the **Edit Settings** button and check or uncheck the boxes next to each option you want to change. When you are finished, click the **Save Settings** button.
- Note:** You may be prompted to provide your Customer Portal login credentials to gain access to the Carbon Black File Reputation Options page. Have these credentials available when you want to view and edit the options.
- 13.** To see the history of Carbon Black File Reputation configuration changes for your server, click the **View Log** link. When you are finished with this configuration, close the browser window.

Once integration with Carbon Black File Reputation is activated, synchronization of files on your server with the Carbon Black File Reputation server begins. To initiate a look up of specific files by hash in Carbon Black File Reputation, you can click the file **View Carbon Black Reputation Data** button on the File Details pages or choose the same command from the Action menu on the Files table page. The analysis results for each file are displayed in a new browser tab. Note that for multi-file requests in Internet Explorer, the popup blocker may block the results for each file after the first one.

Carbon Black File Reputation Availability Status

The App Control Server verifies its connection to Carbon Black File Reputation continuously. If Carbon Black File Reputation is not available, an error is displayed on the Licensing tab indicating the reason for the service interruption.

In addition, there is a built-in *Carbon Black File Reputation Unavailable Alert* that is triggered when expected Carbon Black File Reputation tasks are not performed during a

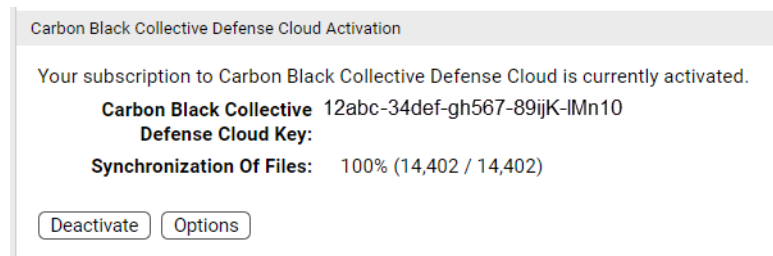
period of time specified in the alert (by default, three hours). When triggered, the alert may also send an email notification to a list of alert subscribers.

The three-hour default setting for the Carbon Black File Reputation Unavailable Alert helps eliminate unnecessary alerts for temporary network issues that would be resolved before they would have significant impact on Carbon Black File Reputation users. However, you can change the length of time Carbon Black File Reputation must be unavailable before the alert is triggered. See [“Using App Control Alerts”](#) on page 608 for more on alerts, including where they are displayed.

Another connection relevant to Carbon Black File Reputation is the connection between the console user’s browser and Carbon Black File Reputation. This connection is required for activation of Carbon Black File Reputation, and also, when you choose Analyze on a App Control Console file details page, for redirection to the Carbon Black File Reputation file assessment page. When a user navigates to the Licensing tab, the App Control Server checks whether that user can access the Carbon Black File Reputation site and displays the following error if there is a problem with that connection: *Carbon Black File Reputation is currently not accessible. Please check back later.*

Deactivating the Carbon Black File Reputation Connection

If you need to deactivate your connection to Carbon Black File Reputation for some reason, you use the same panel on the System Configuration page Licensing tab that was used for activation.



Carbon Black Collective Defense Cloud Activation

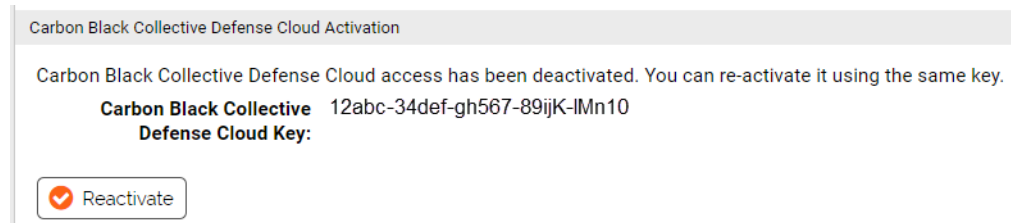
Your subscription to Carbon Black Collective Defense Cloud is currently activated.

Carbon Black Collective 12abc-34def-gh567-89ijk-lMn10
Defense Cloud Key:

Synchronization Of Files: 100% (14,402 / 14,402)

When you click **Deactivate**, a dialog appears warning that trust and threat information will no longer be provided. You confirm deactivation on that dialog.

The key you previously provided to activate the service is stored so that you can reactivate your Carbon Black File Reputation connection simply by clicking the **Reactivate** button.



Carbon Black Collective Defense Cloud Activation

Carbon Black Collective Defense Cloud access has been deactivated. You can re-activate it using the same key.

Carbon Black Collective 12abc-34def-gh567-89ijk-lMn10
Defense Cloud Key:

Using a Proxy Server for Carbon Black File Reputation

You can use a proxy server to handle your communications with the Carbon Black File Reputation. If the proxy server you use does not require authentication, simply provide the URL in the field provided and check the box that activates use of a proxy.

If the proxy server you use requires authentication, you must allow access for the App Control service user account that was configured during App Control Server installation.

You can determine the name of this account by opening the Windows Task Manager and clicking the **Services** button in the bottom right corner. The name in the Log On As field next to App Control Reporter must be allowed to access the proxy server.

The screenshot shows the Windows Services console for 'Services (Local)'. A table lists several services, with a red box highlighting the 'Carbon Black App Control Reporter' and 'Carbon Black App Control Server' services. The 'Log On As' column for these services is 'Local System'.

Name	Description	Status	Startup Type	Log On As
Carbon Black App Control Reporter	Reporting e...	Running	Automatic	Local System
Carbon Black App Control Server	Provides ap...	Running	Automatic	Local System
Carbon Black Sensor	Collects sys...	Running	Automatic	Local System
Certificate Propagation	Copies user ...	Running	Manual	Local System
Client for NFS	Enables this...	Running	Automatic	Network Service

Carbon Black File Reputation Synchronization

When Carbon Black File Reputation is activated, it begins synchronizing file information with the App Control Server. This synchronization allows Carbon Black File Reputation to provide any trust and threat levels it has for files on the server. The amount of time this takes depends upon the number of files to be synchronized.

After the initial synchronization, Carbon Black File Reputation and App Control Server continue to communicate. New files discovered on the server are synchronized with Carbon Black File Reputation, trust and threat levels are updated when they change, and other file metadata, such as publisher and certificate data, may be updated.

The Carbon Black File Reputation Activation panel displays file synchronization status. It includes the total number of unique files in the file inventory, the number and percent synchronized so far, and the estimated time until synchronization is complete. This is especially useful during the initial synchronization, but can also help track the availability of trust and threat information on the server when a large number of new files appears on the server.

The screenshot shows the 'Carbon Black Collective Defense Cloud Activation' panel. It displays the following information:

- Your subscription to Carbon Black Collective Defense Cloud is currently activated.
- Carbon Black Collective** 12abc-34def-gh567-89ijk-IMn10
- Defense Cloud Key:**
- Synchronization Of Files:** 50% (7,201/14,402)

At the bottom, there are two buttons: 'Deactivate' and 'Options'.

Note

The estimated time to complete synchronization might not be accurate if there are technical difficulties with the database or an interruption in network connectivity to Carbon Black File Reputation. If an error occurs during synchronization, the process is paused temporarily to allow normal operations to be restored, and an error message indicates the length of the pause.

Activating Carbon Black EDR Integration

If you are managing your endpoints with both a App Control Server and a Carbon Black EDR Server, you can configure the App Control Server to connect to the Carbon Black EDR Server to receive and display information about files and Carbon Black EDR watchlist events. [Table 118](#) shows the configuration settings for this integration. Carbon Black EDR Server configuration is located on the Licensing tab of the System Configuration page.

Carbon Black EDR

URL: Example: https://hostname_or_ip

Validate SSL Certificate:

API Token:

Receive Watchlist Events: Force Strong SSL:

Table 118: Carbon Black EDR Integration Configuration settings

Field/Button	Description
URL	The URL of the Carbon Black EDR server you want to link to the App Control Server. Port is only necessary if you do not use standard ports on the Carbon Black EDR server (80 for HTTP, 443 for HTTPS).
Validate SSL Certificate	Checking this box causes a validity check on the Carbon Black EDR server certificate. Check this option only if the Carbon Black EDR server certificate is not signed.
API Token	You enter the API Token here for a Carbon Black EDR server user that will be used for the App Control integration. Click the Test button to confirm that the server is accessible and the key works. The test returns one of the following values: <ul style="list-style-type: none"> • Success, version: <Carbon Black EDR product version> • Invalid API Token • Server not accessible Important: See “Creating a Carbon Black EDR User for the Integration” for more on this field.
Receive Watchlist Events	Checking this box activates delivery of Carbon Black EDR watch list events from the configured server to the App Control Server.
Force Strong SSL	Checking this box causes the Carbon Black EDR server to check the App Control Server certificate before sending events. <i>Do not</i> check this if the server uses a self-signed Carbon Black certificate on IIS.

Important

These settings also appear in the Carbon Black EDR console. Although these settings in the Carbon Black EDR console allow editing, changes made there will not be applied. App Control-Carbon Black EDR integration settings should be edited only in the App Control Console.

Creating a Carbon Black EDR User for the Integration

Because you may enter only one API token when configuring the integration between a Carbon Black EDR Server and a App Control Server, you should create a new Carbon Black EDR user for this purpose and use the API Token for that user. The Carbon Black EDR user whose token is used must be in the Administrators group and also must be a Global administrator. The summary of basic steps is shown below:

To create a Carbon Black EDR integration user and API Token (summary steps):

1. Login to the Carbon Black EDR server as a user who can create other administrative users.
2. In the Carbon Black EDR console, go to **Administration > Users**, click **Add User**, and create the new App Control integration user. Be sure to assign this user to the Administrators team and also check the *Global administrator* box.
3. Logout and login to the Carbon Black EDR console again as the new user you created.
4. In the Carbon Black EDR console menu, choose *username* > **My Profile**.
5. On the My Account page, choose **API Token** on the left menu and copy the string in the **Your API Token** box. This is the string you will use to configure Carbon Black EDR integration in the App Control Console, on the Licensing tab of the System Configuration page.

See the *VMware Carbon Black EDR User Guide* for more detailed instructions.

Configuring Unified Management

See [Chapter 27, “Unified Management of Multiple Servers,”](#) for instructions on configuring this feature.

Configuring SAML Logins

The App Control console can be integrated with identity providers (IdPs) that use the Security Assertion Markup Language (SAML). This allows you to require two-factor authentication (2FA) for logging in to the App Control console, for compliance purposes or to meet your own best practice standards.

Integrating a SAML identity provider with App Control requires the following:

- An account with an IdP whose sign-on and logout locations have a binding of type HTTP-redirect.

- For each IdP identity, mapping requires specification of an email address from the IdP account using *one of* the following attributes:
 - A **NameID** of type `urn:oasis:names:tc:SAML:1.1:nameid-format:emailAddress`
 - An attribute with the name **EmailAddress** (capitalized as shown here)
- An App Control login account matching the value of NameID or EmailAddress for each IdP user you want to give access to the console. If you are unfamiliar with creation of login accounts, see [“Creating Login Accounts in the Console”](#) on page 95 for instructions.

Note: This integration allows you to use SAML to authenticate *existing App Control accounts*. It does not import accounts from an IdP to App Control. Also, if both NameID and EmailAddress are found, the EmailAddress attribute is always used, and it must match the email address in a App Control account. NameID is not used as a backup if EmailAddress exists.

- Completion of the configuration procedures below.

Integrating App Control with an IdP

In the terminology of SAML, App Control is a Service Provider. Identity providers and Service Providers must create a trust relationship to be able to work together. The key step required for this trust is to exchange XML metadata with each other. The following procedure requires that you log in to both your IdP and your App Control Console.

To add App Control as a service provider for an identity provider:

1. Log in to your identity provider’s website, or if you have not yet activated an IdP, create an account with one.
2. Go to the page where your provider allows you to add a new service provider.
3. Where prompted, enter the name (FQDN) for your App Control Server. For example: *cbpserver1.myorg.local*
4. Choose one of the following to map IdP accounts to existing App Control Console accounts – in either case the data used for mapping must include an email address that matches an existing console account:
 - Use **NameID** in the following format :
`urn:oasis:names:tc:SAML:1.1:nameid-format:emailAddress`
and provide the NameID Attribute that identifies email addresses on your IdP
- or -
 - Use an attribute with name **EmailAddress** (capitalized as shown). If EmailAddress is provided, it is always used for mapping, even when there is no matching App Control Console account.
5. Log in to the App Control console, click on the configuration (gear) icon in the console menu, and choose **System Configuration**.

6. Click on the **SAML Login** tab.

System Configuration

General Events Security Advanced Options Mail Licensing External Analytics Connectors Unified Management **SAML Login**

SAML Configuration

Identity Provider

Current Identity Provider: None

+ Add Identity Provider

Service Provider

XML Manual

Service Provider XML: [Click here](#) to download the XML for this service provider or copy and paste from below. i

Select all

```
<?xml version="1.0"?>
<md:EntityDescriptor xmlns:md="urn:oasis:names:tc:SAML:2.0:metadata"
entityID="https://apc-58610.cbenglab.com/simplesaml/module.php/saml/sp/metadata.php/default-sp">
  <md:SPSSODescriptor
protocolSupportEnumeration="urn:oasis:names:tc:SAML:2.0:protocol
urn:oasis:names:tc:SAML:1.1:protocol">
    <md:SingleLogoutService
Binding="urn:oasis:names:tc:SAML:2.0:bindings:HTTP-Redirect"
```

7. In the Service Provider section, do one of the following:
 - In the Service Provider XML field, click the **Click here** link to download the Service Provider XML.
 - or-
 - Click the **Select all** button and copy the XML from the window.
8. Go to your IdP's page for configuring a service provider, follow the instructions for importing or pasting the App Control XML you copied.
9. Enter any other information required by the IdP site, and when finished, submit or save your service provider information. Keep both the App Control console and the IdP website page open and continue with the next procedure.

To add an identity provider to App Control:

1. Make sure you have completed the steps for adding App Control as a Service Provider for your IdP.
2. On the App Control SAML Login page, click on **Add Identity Provider**.

System Configuration

General Events Security Advanced Options Mail Licensing External Analytics Connectors Unified Management **SAML Login**

SAML Configuration

Identity Provider

Current Identity Provider: None

+ Add Identity Provider

3. Leave the Configuration Type as **XML** and enter an Identity Provider Name. This is the name that will appear on the App Control login page. It is for easy identification only and has no programmatic impact.

The screenshot shows the 'Add Identity Provider' dialog box. It features a title bar at the top with the text 'Add Identity Provider'. Below the title bar, there are three main sections:

- Configuration Type:** Two radio buttons are present, with 'XML' selected and 'Manual' unselected.
- Identity Provider Name:** A text input field is provided, followed by a blue information icon.
- Identity Provider XML:** This section includes the text 'Choose file or paste XML' and another blue information icon. Below this is a 'Choose File' button and the text 'No file chosen'. A large, empty text area is provided for pasting XML content.

 At the bottom of the dialog, there are two buttons: 'Save' (with a floppy disk icon) and 'Cancel' (with a red 'X' icon).

4. On the IdP site, locate the page that contains that provides public IdP XML metadata for this provider and either copy it or download it to a file.
5. Go back to the App Control SAML Login page, and in the Add Identity Provider dialog, do one of the following:
 - In the Identity Provider XML field, click the **Choose File** button to download the identity provider XML.
 - or-
 - Paste the XML you copied from the identity provider into the window below Identity Provider XML.
6. Click the **Save** button at the bottom of the Add Identity Provider dialog. The IdP is now configured in App Control. You can exit the IdP website, but keep the App Control Console window open to configure User Roles that will be able to log in locally.

Allowing Non-SAML Logins for Specified User Roles

You must set up at least one User Role in App Control for local (non-SAML) login to the console so that you have a way to log in if there are problems with the IdP. This can also be useful for troubleshooting purposes, although you should be sure you limit the accounts for which you allow local login permissions in a way that still meets your compliance or standards requirements.

To allow users with selected User Roles to log in locally (without SAML):

1. On the App Control SAML Login page, go to the Local Login Override Permissions section. All currently defined User Roles for your App Control Server are listed.

2. Check the box for each User Role whose users you want to allow local (non-SAML) login to the console. You must choose at least one role.

The screenshot shows the 'System Configuration' page with the 'SAML Login' tab selected. Under 'SAML Configuration', the 'Identity Provider' is set to 'Sign-On Central'. Below this, there are 'Edit' and 'Delete' buttons. The 'Local Login Override Permissions' section contains a table with columns 'Allow Local Login' and 'User Role'. The 'Administrator' and 'Administrator (Unified Management)' roles have their 'Allow Local Login' checkboxes checked. At the bottom, there are 'Save User Roles' and 'Cancel' buttons.

Allow Local Login	User Role
<input type="checkbox"/>	ReadOnly
<input type="checkbox"/>	PowerUser
<input checked="" type="checkbox"/>	Administrator
<input checked="" type="checkbox"/>	Administrator (Unified Management)
<input type="checkbox"/>	User (Unified Management)

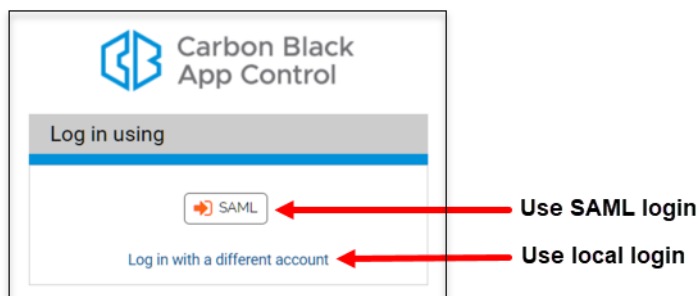
3. When you have checked all appropriate User Roles, click the **Save User Roles** button. User accounts that have these roles will be allowed to login either with SAML or with their own local credentials.
4. If you have completed all SAML Login configuration tasks, you can navigate to a different page or log out of the console.

Note

You can also enable and disable local logins on the Login Accounts Edit User Role page.

Logging In Using SAML

Once you save a properly configured IdP in App Control, you can use SAML to log in. A new button with your identity provider name appears on the App Control login page.



When a user clicks the button, they are directed to the login page of the IdP. If they provide their correct credentials in the IdP, they are logged into the App Control Console as the user whose email address matches the one for the IdP account used.

Users with one of the roles configured to allow local logins can click **Log in with a different account** and enter their App Control credentials to access the console.

A user who attempts a local login without having a user role that allows this will see a login error that instructs them to use a valid user name and password.

Important

The IdP account you log in with must have the same email address (specified through the EmailAddress or NameID value of the IdP) as a user in App Control. The identities are not matched by name. If the email address for the IdP account a user logs does not match the email address of any App Control user, the login will fail.

Deleting or Editing an Identity Provider

You can delete or modify an identity provider in App Control. You might decide not to use an identity provider at all, or to change to a different provider. Only one identity provider per server is allowed, so if you wanted to change providers, you would first delete the old one and then add the new one.

To delete the identity provider for a App Control Server:

1. In the console menu, click on the configuration (gear) icon and choose **System Configuration**.
2. Click on the **SAML Login** tab.

3. In the Identity Provider panel, click **Delete**, and if you are certain you want to delete the provider, click **OK** in the confirmation dialog.

The identity provider is removed and SAML logins are disabled. Logins will be handled locally via the usernames and passwords configured in App Control itself.

To edit the configuration of an identity provider for a App Control Server:

1. In the console menu, click on the configuration (gear) icon, choose **System Configuration** and click on the **SAML Login** tab.
2. In the Identity Provider panel, click **Edit**. The Edit Identity Provider dialog appears.

3. [Table 119](#) describes the fields available in the dialog. Edit the settings as needed and then click the **Save** button at the bottom of the dialog.

Table 119: Identify Provider settings

Field/Button	Description
Configuration Type	This field has two radio buttons that determine editing mode: XML and Manual . When you edit an IdP, the default choice is Manual.
Identity Provider Name	This field shows the current IdP name that appears on the login button in App Control. Changing this changes the button name but does not affect anything else about the IdP configuration.

Field/Button	Description
Identity Provider Entity ID	This is the base URL for the IdP. For example: https://idp.socentral.com
Sign-on Location	This is the URL for signing on to your IdP. For example: https://idp.socentral.com:443/sso/SSORedirect/metaAlias/publicidp
Logout Location	This is the URL for logging out of the IdP. For example: https://idp.socentral.com:443/sso/IdPSloRedirect/metaAlias/publicidp
Signing Certificate	The identity provider's signing certificate.
Encryption Certificate	The identity provider's encryption certificate

Chapter 27

Unified Management of Multiple Servers

Sections

Topic	Page
Overview	778
Unified Management Features	778
Configuring Unified Management	780
Enabling Unified Management and Adding Servers	780
Creating Unified Management Console Accounts	785
Editing Client Server Configuration	788
Disabling Unified Management	791
Unified Management of Rules	791
Managing Unified Rules from the Software Rules Page	792
Copying Existing Rules to Other Servers	798
Editing a Unified Rule	801
Changing a Unified Rule to a Local Rule	803
Managing Unified File Rules from File Table and Details Pages	803

Overview

If you have multiple App Control Servers, you can centralize the management of those servers. Unified Management allows you to specify that one server can control many common management functions not only for itself but for any of your other connected App Control Servers. You might choose this option for one or more reasons, including:

- You need to host local App Control Servers in several different regional locations but want to manage some of their functions from a central server.
- You would like to have different types of endpoints (for example, servers, desktops, POS systems) reporting to different servers while still having the capability of managing certain rules or functions from a central server.
- You would like to spread your endpoint management load among smaller App Control Servers and databases rather than doing all of the processing and storage through one very large server. For example, you might have regional offices with lower capacity network connections that can't support pushing all endpoint data to a central location.

In any of these situations, Unified Management provides the ability to execute some of the primary App Control Console functions from a centralized location and have it affect all of your App Control Servers or just selected servers. All App Control licenses that include Control features can use Unified Management.

Both the management server and the client servers it manages must be at v8.0.0 or greater to use this feature.

Unified Management Features

In this release, Unified Management makes the following features available:

- **Configuration of Unified Management** – You can enable or disable Unified Management for your entire environment, designate a central management server, add or remove other servers from unified management, and determine which users can access unified management as users or as administrators.
- **File Information** – The file inventory for all servers under unified management is available through the central management server. This includes the File Catalog, the Files on Computers list, and the File Details pages. In addition, you can search for a file on all managed servers in one operation.
- **Rules** – For many rule types, you can create and manage a rule on the central management server and have it apply to some or all of the managed servers. This includes File Rules (approvals and bans), Custom Rules, Memory Rules, and Registry Rules.
- **Single Sign-On from Management Server** – Once a user with unified management credentials has been authenticated on a client server, that user can login to the client server directly from the management server, without providing username and password again.
- **Events** – New or modified events have been added to track Unified Management activities. There is a new Unified events saved view on the Events page, and there are is a new (optional) Unified Server Source column that can be added to event tables.
- **User Interface Changes** – On menus, tabs, and page banners, a symbol showing three cubes together indicates that a feature has Unified Management enabled, and on pages that show results from multiple servers, the word “Unified” appears next to

the page heading. For example, the following example shows the top of a Files on Computers page on a server with Unified Management enabled.

Table 120: Unified Management Features

App Control Page	Unified Management Features
All (console menu bar)	Log in to client servers from the management server without providing additional login credentials. Requires initial authentication for user.
File Catalog	View file inventory and take actions on files reported by the management server and all servers reporting to it.
Files in Computers	View inventory and take actions on files instances reported by all unified servers.
File Details	View file details and take actions on a file reported by any unified server.
Find Files	Find instances of a file on all unified servers.
File Rules	Apply an approval or ban, or remove a ban or approval, to file instances on any unified servers. May be further defined by server and/or policy.
File Rules Details	Apply an approval or ban, or remove a ban or approval, to file instances on any unified servers. May be further defined by server and/or policy.
Custom Rules	Apply, enable, or disable a rule on any unified servers. May be further defined by server and/or policy.
Custom Rule Details	Apply, enable, or disable a rule on any unified servers. May be further defined by server and/or policy.
Memory Rules	Apply, enable, or disable a rule on any unified servers. May be further defined by server and/or policy.
Memory Rule details	Apply, enable, or disable a rule on any unified servers. May be further defined by server and/or policy.
Registry	Apply, enable, or disable a rule on any unified servers. May be further defined by server and/or policy.
Registry Rule Details	Apply, enable, or disable a rule on any unified servers. May be further defined by server and/or policy.

Configuring Unified Management

Configuring Unified Management on App Control servers requires that you:


- Have App Control v8.0.0 or later installed on the management server and all systems that it will manage, and have connectivity between the servers
- Make sure security protocols are compatible among servers; if one server is using the TLS 1.2 protocol only, all servers must use it
- Have a console account on the management server that has the *Administrator (Unified Management)* role enabled
- Choose a management server and enable Unified Management on that server on the System Configuration page
- Add each managed client server on the System Configuration page of the management server; this requires providing its URL and authenticating access by entering a username and password for a user on the client server, and also requires that the client server be at the minimum version number shown on the configuration page

Enabling Unified Management and Adding Servers

You enable a server to be the management server and specify the client servers that you want centrally managed on the System Configuration page of the console.

To enable Unified Management on a server:

1. Login to the console for this server with an account that has the *Administrator (Unified Management)* user role enabled. The default “admin” account has this role enabled.
2. On the console menu bar, click the configuration (gear) icon, choose **System Configuration**, and click the **Unified Management** tab.
3. Click the **Edit** button at the bottom of the page
4. In the Unified Management field, choose **Enabled** and then click the **Update** button at the bottom of the page.

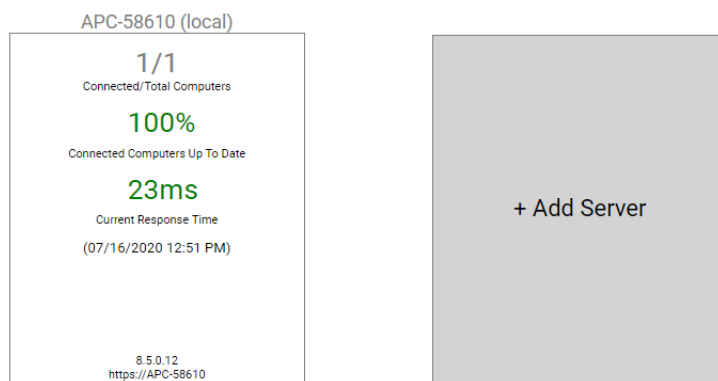


If you know which servers you plan to add as clients for Unified Management, continue with the next procedure.

Adding Client Servers to Unified Management


To add a client server to the Unified Management server:


1. If you are not already logged in, log in to the management server with an account that has the *Administrator (Unified Management)* user role enabled. The default “admin” account has this role enabled. If you would rather use a different account, see [“Creating Unified Management Console Accounts”](#) on page 785.
2. If Unified Management is already enabled on this server, you can choose the **Configure Unified Management** on the servername dropdown – this is a shortcut to the Unified Management tab on the System Configuration page. The shortcut appears only for users with the *Configure Unified Management User Role* enabled.
3. On the Unified Management tab, click the **Edit** button at the bottom of the page and then click the **Add Server** box. The Configure Server panel opens.



Configure Server

Please specify a valid server URL to be used with Unified Management. Make sure to authenticate before clicking on Update button.

Server URL: 

Enable Certificate Verification: 

4. In the Server URL box, provide the URL for the server you want to be managed and click the **Authenticate** button. This opens a new dialog box.
5. In the Authenticate dialog box, provide a console user account on the *client server* that has *Configure Unified Management* permission enabled in one of its user roles, then click **Submit**.

The dialog box is titled 'For server: APC-91166'. It contains the text: 'Credentials are required to get authentication token from the remote server. Your password won't be stored.' Below this are two input fields: 'User Name:' and 'Password:'. At the bottom is a blue 'Submit' button.


If authentication is successful, the dialog closes and you can continue with the next step. There are several authentication failure conditions, including:

- **Incompatible Server Version** – If you try to authenticate with a server running a version of App Control that does not support Unified Management, a message such as *“The remote server needs to be at least version 8”* or one indicating a specific build number appears. You must upgrade the client server if you want it to be under Unified Management.
 - **Server not Reachable** – If you try to authenticate using an incorrect URL, or if the system is not connected to the network, or if a server is not installed at that address, the dialog closes and the message *“Server is not reachable. Authentication could not be tested,”* appears on the Unified Management page. Close the dialog box and check the name, connection, and server status of the client server.
 - **Non-existent Account** – If you try to use an account that doesn’t exist on the client server, the dialog box and Unified Management page show the message *“Server is reachable but authentication failed.”* Close the dialog box and use an existing account with the necessary permissions.
 - **Existing Account without UM Permissions** – If you use an account that exists on the client server but does not have the proper permission, the following error message appears: *“Server is reachable but authentication failed or required permissions are not assigned.”* If this happens, click **Cancel** on the Unified Management configuration page and use a different account.
 - **Incompatible Security Protocols** – If one server is using the TLS 1.2 protocol only, all servers must use it. Otherwise, you will see the message *“Remote server does not support TLS 1.2, please upgrade it to latest version.”*
6. If the client server is using a trusted SSL certificate and you want to use this to verify the connection, you can check the Enable Certificate Verification box.
 7. When Authentication succeeds, click the **Add Server** button to continue adding another server to be managed by this server and follow the configuration steps.
 8. When you have finished adding servers that you will manage from this server, click **Update**.

Notes

- Once the connection to the client server is authenticated, it remains authenticated unless the server URL is changed. Users with either Use Unified Management or Configure Unified Management permission can view details and take actions on the client servers while logged in to the management server.
- A user accessing a client server from the management server has the permissions of the account used to authenticate the connection, not their own permissions.
- When a user accesses a client server from the management server, actions the user takes appear in events as having been performed by the authentication account, not the logged in user.

Server Information on the Unified Management Page



The Unified Management tab on the System Configuration page displays information about the management server and its client servers. The panel for the management server and each successfully connected client server includes the following information:

- **Server Name**
- **Connected/Total Computers** – The number of agent computers currently connected to the server and the total number of connected and disconnected agents registered with this server
- **Connected Computers Up To Date** – The percentage of currently connected computers whose Policy Status is *Up to date*; if there are no connected computers, the value shown is *N/A*.
- **Current Response Time** – The current latency (in milliseconds or seconds) of the communications between the management server and its client servers, measured at the time of the Unified Management page is loaded. In addition to numeric values, color coding provides a quick status of the connection quality:
 - **500 Milliseconds or Less:** This is shown in green. It indicates that the remote server is responding well, with latency similar to a local server.
 - **501 Milliseconds to 1 Second:** This is shown in yellow. It indicates that the remote server is responding slowly.
 - **Greater than 1 Second:** This is shown in red. It indicates remote server is responding very slowly.
 - **N/A:** This is shown in red. It indicates that the server has not responded during a 10-15 second waiting period.
- **Timestamp** – This is shown in parentheses below the Current Response Time. It shows the time of the most recent latency check.
- **Server Version** – The App Control Server version number for this server
- **Server URL** – The URL for the server. For client servers, this is the URL provided when it was added to Unified Management.

Disconnected and Unreachable Servers

When client servers are active, they appear on the Unified Management configuration page as shown in the previous illustration, with data about number of connected agents, response time, etc. If a server is disconnected, its panel on the Unified Management configuration shows one of two indications:

The screenshot shows a 'Servers' section with three panels:

- APC-58610 (local)**: Shows 1/1 Connected/Total Computers, 100% Connected Computers Up To Date, and 21ms Current Response Time. URL: https://APC-58610.
- APC-20359**: Gray background. Text: 'This server is temporarily disconnected. Please reconnect it before using it.' URL: https://APC-20359. Label: **Disconnected by User**.
- APC-91166**: Red background. Text: 'This server cannot be contacted. Please validate its connectivity and SSL certificate. Also make sure that you have authenticated with valid credentials and required permissions. For further assistance, please contact Carbon Black support (error code 404)'. URL: https://APC-91166. Label: **Server Unreachable due to error. ie. powered off**.

- **Disconnected (gray)** – If the server was intentionally disconnected using the Temporary Disconnect button, the panel has a gray background.
- **Error (red)** – If a server has not been intentionally disconnected but cannot be contacted when the configuration page loads, its panel shows a red background to indicate an error. If the error type can be identified, an error code may be displayed. You can also do the following to search for the cause of the error:
 - On the Events page, use the **Unified events** Saved View.
 - In the error log, look for errors involving the client server or Unified Management. The error log is in the following location:

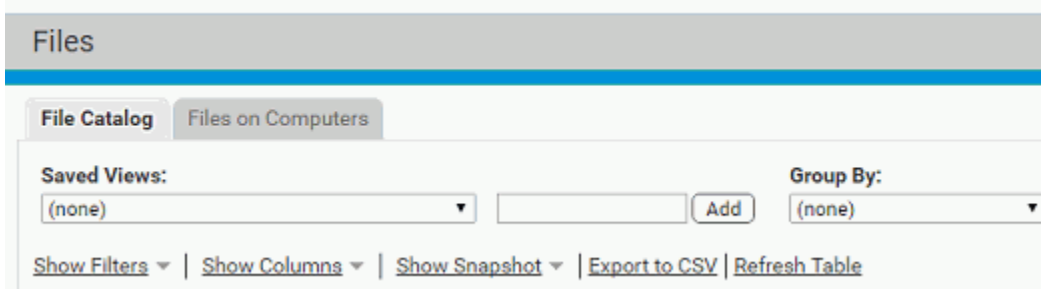
```
<ServerInstallDir>\Parity Console\WebUI\Logs\php_errors.log
```

Authentication Errors

Lack of authentication is one of the most common Unified Management errors you will see. This error occurs when a user logs in to the management server and that user has never authenticated the connection with one or more client servers. It might also happen if an account used for authentication has changed and no longer has unified management permissions.

If a user is not authenticated for a client server, the user is denied access through unified management to information about that server. The error message for this condition appears as soon as the user attempts to access a page that has unified management features. It lists all client servers that are not authenticated for the user.

This unified server is not authenticated: protect-west_mycorp.com. Please see the [User Settings](#) page to authenticate.

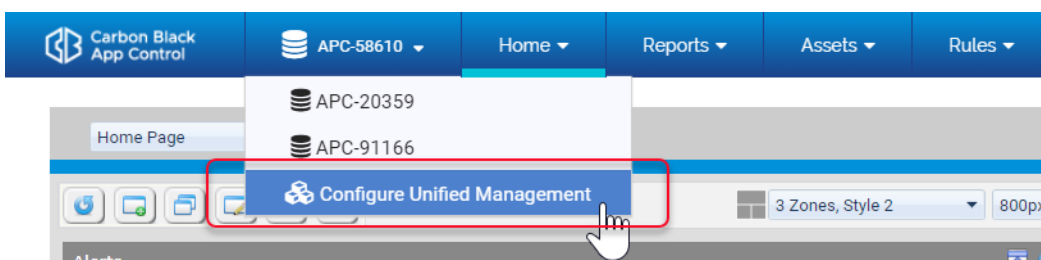


If you see this error, click on the [User Settings](#) link in the message and configure authentication for each listed server on the [User Settings](#) page, as described in [“Authenticating a User on Client Servers”](#) on page 786.


Configuration Information on Managed Servers

You can determine whether a server is managing other servers by choosing **Configure Unified Management** on the server name menu in the console menu bar. This opens the Unified Management tab on the System Configuration page.

Note: This option displays on the server managing the other servers. The child servers do not have a drop-down option.



If a server has been enabled as the management server, the full management interface is shown, including the details on all servers and the Add Server box. For client servers that are *being* managed, the page shows only a message that the current server is “under management” from another server.



Creating Unified Management Console Accounts

You can add Unified Management privileges to an existing account or create users specifically for this purpose. There are two user roles for Unified Management:

- **Administrator (Unified Management)** – This role has all permissions for all App Control features, including Configure Unified Management, enabled. The account you create for the initial configuration of the management server, and the account used for authentication when you add new client servers, must have this role. The the default “admin” account includes this role.
- **User (Unified Management)** – This role allows the user to view information in the App Control console, such as computers, events, and files, including information provided by other computers that are under Unified Management. It also allows creation, distribution, and modification of unified rules and single-sign-on access to client servers from the management server. It does not allow modification of the Unified Management configuration.

Note

In previous releases, each console user belonged to one *group*, and that group defined the user’s permissions. In v8.0.0, groups have been replaced by *user roles*. Although roles and groups define permissions in a similar way, you can assign more than one role to a user. Since the Administrator (Unified Management) role has all permissions, it is not necessary to give it any other roles. See [Chapter 3, “Managing Console Login Accounts,”](#) for more details on role-based access control.

Each user account that will be using Unified Management features must authenticate their connection to each client server when they log in for the first time.

To create a Unified Management login account:

1. Log in to the App Control console with an account that has permission to manage login accounts. The default ‘admin’ account has that permission.
2. On the console menu bar, on the configuration (gear) menu, choose **Login Accounts**.
3. On the Users tab of the Login Accounts page, click the **Add User** button.
4. On the Add Login Account page, provide the mandatory information and any other information you would like available for this account. In the User Roles panel:
 - a. To create an account that can both use the features and change the configuration of Unified Management, check the **Administration (Unified Management)** box.
 - b. To create an account that can use the features but not change the configuration of Unified Management, check the **User (Unified Management)** box.
5. When you have completed the configuration for this user, click **Create & Exit**.

Authenticating a User on Client Servers

To use Unified Management to access existing client servers, a user must have:

- a user account on the management server that has permission to use Unified Management features
- an authentication account on each client server that provides access to the features they need to use on the client server

Except for the account used when the server is first added, the authentication account does not need Unified Management permissions. The permissions of the authentication account determine what a Unified Management user can do on a client server. In effect, the authentication account determines who a user is when they are on a client server.

To authenticate a Unified Management user:

1. Create an account with Unified Management permissions, or add those permissions to an existing account.
2. Log in to the management server with the account.
3. On the console menu, choose **User Settings** on the username menu. The User Settings page appears.

User Settings

Change Password

Existing Password:

New Password: i

Confirm Password:

Display Preferences

Remember Page Settings:

Resizable Table Columns:

Wrap Long Text In Tables:

Set Rows Per Page:

Default Starting Page:

Unsaved Changes Warning: Display a warning when a page has changes that have not been saved

Click To Dismiss Feedback Banners: Require a click to dismiss the feedback banner. If unchecked, banner will automatically dismiss after a short period.

Unified Server Authentication

Authentication with each unified server is required before using Unified Management features.

APC-20359	Authenticate	● Authenticated
APC-91166	Authenticate	● Authenticated

4. In the Unified Server Authentication panel, click the **Authenticate** button for the first server. The authentication dialog box appears.
5. In the authentication dialog box, provide a console user account and password on the *client server* that has the permissions you want this user to have when they access the client through the management server, then click **Submit**.

For server: APC-91166 ✕

Credentials are required to get authentication token from the remote server. Your password won't be stored.

User Name:

Password:

6. If you provide a valid account on the client server, the user is authenticated for Unified Management access to that server.
7. Click the Authenticate button for the next unauthenticated server and repeat step 5. Continue authenticating servers until they all show that they are “Connected, authenticated.”

- When all servers are authenticated with this user, click the **Save** button at the bottom of the page.

Editing Client Server Configuration

Enabling Edit mode on the Unified Management page displays additional fields on the panels for client servers. These fields vary depending upon whether the client server is connected:

- **Show/Hide Configuration** – Clicking this field toggles the Configuration fields on and off for a client server.
- **Temporary Disconnect** – For a connected server, clicking Temporary Disconnect disconnects the client server but retains its configuration information. File information about disconnected servers does not appear on the management server.
- **Reconnect** – For a temporarily disconnected server, clicking Reconnect puts the client server back under Unified Management control, and any rules previously applied via Unified Management remain centrally managed.
- **X** – Clicking this button deletes the server from Unified Management. Rules that were applied to the client server by Unified Management remain on the disconnected client server, but they become local rules and will no longer be centrally managed, even if this server is returned to Unified Management at a later time.

The actions taken when you use the fields in Edit mode generally include a confirmation dialog. If you choose to complete the action, it goes into effect when you click **Update**.

To edit the configuration for a client server:

- On the Unified Management tab of the System Configuration page, click the **Edit** button.
- On the panel for the server you want to modify, click the **Show Configuration** link.


The screenshot shows a 'Servers' section with two server panels. The left panel is for 'APC-58610 (local)' and the right panel is for 'APC-20359'. Both panels display '1/1 Connected/Total Computers', '100% Connected Computers Up To Date', and '25ms' or '15ms' Current Response Time. The right panel also shows a 'Temporary Disconnect' link and a 'Show Configuration' link highlighted with a red box and a hand cursor. A small 'X' icon is visible in the top right corner of the right panel.


- In the configuration, make whatever changes are needed. Possible changes include:
 - Changing the server URL.
 - Authenticating the connection to the client server if authentication has failed or was never completed for the currently logged in user.

- Enabling verification of the connection to the client server via a certificate.

Configure Server

Please specify a valid server URL to be used with Unified Management. Make sure to authenticate before clicking on Update button.

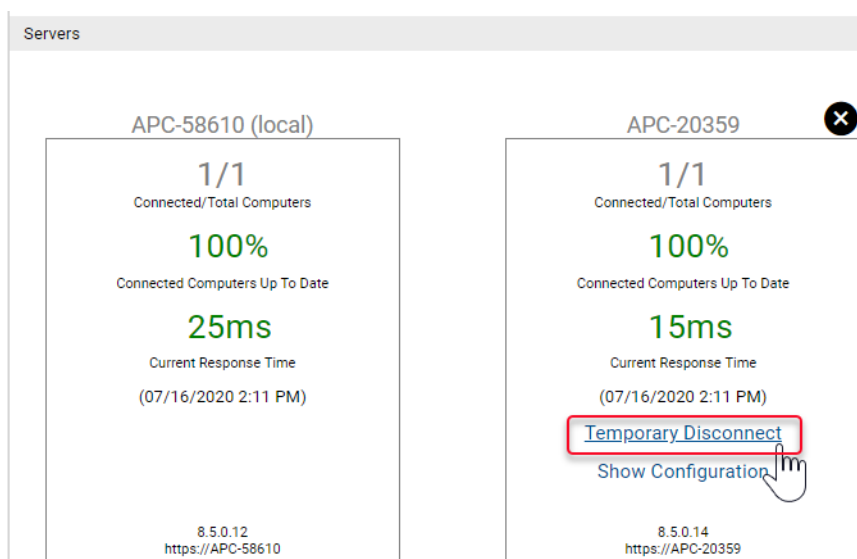
Server URL: 

Enable Certificate Verification: 

4. When you have made the necessary changes, click **Update**.

To temporarily disconnect a client server:

1. On the Unified Management tab of the System Configuration page, click the **Edit** button.
2. On the panel for the server you want to temporarily disconnect, click **Temporary Disconnect**.

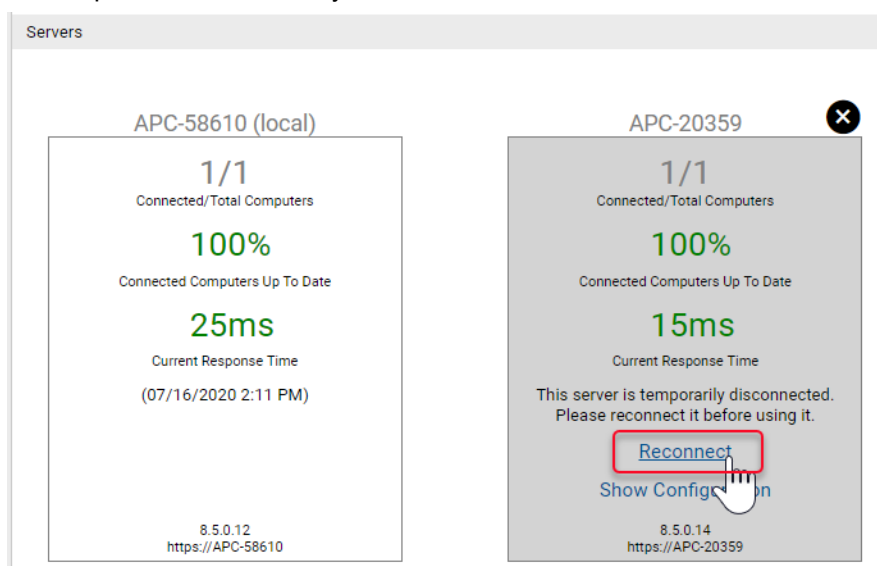


The screenshot shows a 'Servers' management interface with two server panels. The left panel, titled 'APC-58610 (local)', displays '1/1 Connected/Total Computers', '100% Connected Computers Up To Date', and '25ms Current Response Time (07/16/2020 2:11 PM)'. The right panel, titled 'APC-20359', displays '1/1 Connected/Total Computers', '100% Connected Computers Up To Date', and '15ms Current Response Time (07/16/2020 2:11 PM)'. A red box highlights the 'Temporary Disconnect' button on the right panel, which is being clicked by a mouse cursor. The background of the right panel is grayed out, and a close button (X) is visible in the top right corner of the panel area.

3. In the confirmation dialog, click **OK** to complete the disconnection. The panel background becomes gray, a status message shows that the server is disconnected, and the Temporary Disconnect button changes to Reconnect.
4. To disconnect another server, repeat steps 2 and 3 for that server's panel.
5. When you are finished disconnecting servers, click the **Update** button.

To reconnect a disconnected client server:


1. On the Unified Management tab of the System Configuration page, click the **Edit** button.
2. On the panel for the server you want reconnect, click **Reconnect**.



3. In the confirmation dialog, click **OK** to complete the reconnection. The panel background becomes white and all of the status information returns to the panel.
4. When you are finished disconnecting servers, click the **Update** button.

To remove a server from unified management:

1. On the Unified Management tab of the System Configuration page, click the **Edit** button.
2. For the server you want to delete from Unified Management, click the delete button in the upper right above the panel.



3. Review the confirmation dialog, which describes effects of deleting the server. If you still want to delete the server from Unified Management control, click **Yes**.
4. Click **Update** to complete the removal of this server.

Disabling Unified Management

If you choose to disable Unified Management on the management server, the management features are disabled, but the most recent configuration settings (that is, the information about managed servers) remain visible on the Unified Management configuration page. The following changes occur when you disable Unified Management:

- **Management Features** – Other than the configuration page, Unified Management options are removed from other pages in the console. The former management server cannot create rules that affect client servers, and the Server choice options on rules pages disappear.
- **Rule Pages** – Rules created or copied to client servers by the management server are still in effect and show that Unified Management was their source. However, they cannot be edited as Unified Management rules. They can be edited locally on any of the systems that show them.
- **File Pages** – File table pages stop showing the "Unified" banner, stop having a pulldown menu for servers, and no longer show their files grouped by server. All files not on the current server become inaccessible.

To disable Unified Management:

1. Login to the management server with an account that has the Administrator (Unified Management) user role enabled.
2. In the console menu bar, choose **Configure Unified Management** on the servername dropdown menu.
This is a shortcut to the Unified Management tab on the System Configuration page.
3. Click the **Edit** button at the bottom of the page
4. In the Unified Management field, choose **Disabled**.
5. Click **Update**.

The Unified Management feature is disabled on this server and it can no longer manage other servers. To re-enable it, use the configuration (gear) menu to open the System Configuration page and click on the Unified Management tab.

Unified Management of Rules

If you have multiple App Control Servers and have enabled Unified Management, you can centrally manage the following rule types:

- File Rules (bans and approvals)
- Custom Rules
- Registry Rules
- Memory Rules

Unified Management provides additional options for creating, distributing, and managing these rules. You can:

- distribute a new rule to multiple servers at the time you create it
- copy one or more existing rules from the management server to one or more client servers
- customize the rank of unified Custom, Registry, and Memory Rules on each server
- specify whether administrators on client servers can modify unified rules
- enable and disable unified rules on one or more servers
- view all rules (unified or local) on unified servers through the management server

With these options, you can choose how and when you distribute rules to your servers, and determine how much variation you want to allow once the rules are distributed. For example, you might want to send a rule to all servers immediately when you create it, or you might choose to test a new rule on one server first and then copy it to the other servers once you are satisfied with it. You might want some rules to be uniform on all servers while allowing others to be modified as the administrator on each server chooses.

Note

Exporting and importing rules is not the same as Unified Management, and can be done between servers that are not part of a Unified Management group. See [“Exporting and Importing Rules”](#) on page 431.

Managing Unified Rules from the Software Rules Page

The following section describes how unified rules are created and managed from the Files, Custom, Registry, and Memory tabs on the Software Rules page. It uses Custom Rules for most of the examples, but the descriptions and procedures are the same for the other unified rule types, except for ranking which does not apply to File Rules.

Although File Rules can be managed from the Software Rules page, it is more likely that you will create approvals and bans from one of the pages that list files, or from a File Details page. The procedure for creating unified rules from these pages are described in [“Changing a Unified Rule to a Local Rule”](#) on page 803.

To add (create) a unified rule from the Software Rules page:

1. Make sure you are logged in as a user with either the Administrator (Unified Management) role or the User (Unified Management) role. See [Chapter 3, “Managing Console Login Accounts,”](#) if you need to create a new user for this.
2. Go to the page for the type of rule you want to create and click the Add button for that rule type. For example, on the Custom Rules page, click **Add Custom Rule**.
3. On the Add Rule page, configure the General section, and if present, the Definition section.

Note

When you apply a rule to more than one server, client servers use default notifiers, even if a custom notifier is specified on the management server.

4. If this is the management server in a Unified Management environment, a Servers field appears in the *Rule Applies To* panel.

Add Custom Rule

General

Rule Name:

Description:

Status: Enabled Disabled

Definition

Platform:

Rule Type:

Path or File: i

Process: i

Rule Applies To

Servers: APC-58610
 All Servers
 Selected Servers

Policies: All Current and Future policies
 Selected policies

Save & Exit
 Save
 Cancel

You have three choices:

- To apply the rule only to the current server, choose the top radio button, which shows the server name.
 - To apply the rule to the current server and all servers it manages, choose **All Servers**.
 - To have the rule apply to some but not all servers, choose **Selected Servers** and check the box next to each server you want the rule applied to.
5. When you have finished specifying the servers this rule will apply to, continue by choosing the policies the rule will apply to. You have two choices:
- To apply the rule to computers in all policies, choose **All Current and Future Policies**.
 - To have the rule apply to computers in some but not all policies, choose **Selected Policies** and check the box next to each policy you want the rule applied to. Note that the server for each policy is listed.

Note: If you recently created, edited, or deleted policies on a remote system, the policy list from a remote server might not be updated on the management server immediately. With normal connectivity, policy lists from remote systems should be up to date on the management server within one minute, but slower networks might increase this gap.

6. If you chose to apply the rule to servers other than the current server, an Override Permissions section appears for Custom, Registry, and Memory Rules. The Override Permissions control what users that do not have Unified Management roles can do.

Rule Applies To

Servers: APC-58610
 All Servers
 Selected Servers

Policies: All Current and Future policies
 Selected policies

Override Permissions: No Override - Local admin cannot edit and cannot set rank
 Partial Override - Local admin cannot edit, but can set rank
 Full Override - Local admin can edit and can set rank

- **No Override** – If you choose this option, users without either User (Unified Management) or Administrator (Unified Management) *cannot* edit this rule or change its rank relative to other rules.
- **Partial Override** – If you choose this option, users without either User (Unified Management) or Administrator (Unified Management) cannot edit this rule but can change its rank relative to other rules.
- **Full Override** – If you choose this option, any user with permission to edit rules can edit and change the rank of this rule. If you override a rule (other than rank), it becomes a local rule on the server on which the override occurred, and the local rule is no longer connected to the previous (unified) rule on other servers.

Important

- If you configure a rule to allow Partial Override, keep in mind that a user on another server could make the rule ineffective by moving it to a lower rank than a different rule affecting the same action.
- Users with the Administrator (Unified Management) or User (Unified Management) role can edit unified rules regardless of the override setting.

7. When you have finished configuring the rule, click the **Save** button (to stay on the page) or **Save and Exit** button (to return to the table page for this rule type). In the Unified Management environment, this opens the Save Rule wizard.

Unified
Save Rule

1
 Rank rules

2

2
 Send rules

Rules being copied:

- Allow QA Testing path

APC-58610 Custom Rules

Drag and drop the rules below to set their rank. Use the status switch to enable the rules on the server.

Rank	Name	Status
1	Allow QA Testing path	<input checked="" type="checkbox"/>
2	Do Not Track MyApp Temp Files	<input type="checkbox"/> Enable
3	Report execution of potential risk files identified by yara	<input type="checkbox"/> Disable
4	Block execution of malicious files identified by yara	<input type="checkbox"/> Disable
5	Examine powershell script contents	<input type="checkbox"/> Enable
6	Block powershell scripts that execute memory	<input type="checkbox"/> Enable
7	[Sample] Report whenever powershell is launched with encoded arguments	<input type="checkbox"/> Disable
8	Do not treat these processes as NFT applications	<input type="checkbox"/> Enable

← Back
Next > Cancel

8. For Custom, Registry, and Memory Rules, the first page of the wizard shows rule ranking. Initially, any new rule is ranked first in the list of rules, meaning that it will take precedence over lower ranked rules for actions in which rank matters. The Save Rule wizard allows you to customize the ranking of this rule on each server. The wizard provides a separate page to rank the new rule on each server. You can drag and drop the new rule to a different position relative to the other rules:
 - a. On the first page of the wizard, which shows the ranking of rules for the Unified Management server (that is, the server you are logged into), change the rank of the new rule if you choose, and then click **Next** in the bottom right corner of the wizard. The rule ranking for the next server is displayed.
 - b. Continue examining, and if necessary, changing the ranking of the rule on each server, clicking **Next** when done with each one. When you get to the last server, the Next button changes to a **Send rules** button. Clicking this button sends the rules to each server, ranked as you have chosen for each.

Note: You can click the **Back** button in the bottom left of the wizard if you change your mind about the rank of the rule on a previous server, and you can click **Cancel** in the bottom right to go back to the Add Custom Rule page.

9. The Save Rule wizard shows the progress of rule distribution to your servers. When all servers that you specified have received the rule, the wizard shows Complete at the end of the progress bar.

Unified Save Rule

Rank rules Send rules

Progress

Begin APC-58610 APC-20359 APC-91166 Complete

APC-58610
 ✓ Successfully sent rule "Allow QA Testing path"

APC-20359
 ✓ Successfully sent rule "Allow QA Testing path"

APC-91166
 ✓ Successfully sent rule "Allow QA Testing path"

Finish

Click **Finish** to exit the wizard and return to either the Add/Edit Rule page (if you chose Save) or the table page for this rule type (if you chose Save and Exit).

Note

If errors occur during the rule creation and distribution process, they will be shown with red exclamation marks instead of blue checkmarks in the final page of the wizard. Depending upon the issue found, correcting it might be possible by simply repeating the procedure above. Some issues, however, such as connectivity failures, might require remediation on one or more of the servers under unified management.

In the rules tables, the rows for rules under Unified Management appear highlighted in green. This is true on both the management server and client servers that have received rules. In addition to highlighting, tables that show unified rules can include a Unified Server Source field, which shows the name of the management server the rule came from. This field is not displayed by default.

Software Rules

[Updateers](#)
[Rapid Configs](#)
[Publishers](#)
[Users](#)
[Directories](#)
[Files](#)
[Custom](#)
[Memory](#)
[Registry](#)
[Scripts](#)

Servers:
Saved Views:
Group By:

[Show Filters](#) | [Show Columns](#) | [Export to CSV](#) | [Refresh Table](#)

Search: Automatically apply

▼ APC-58610 69 item(s)

Showing 69 out of 69 item(s)

<input type="checkbox"/> Select 54	Rank ▲	Status	Platform	Rule Type	Name
<input type="checkbox"/>	1	<input checked="" type="checkbox"/>	Windows	Trusted Path	Allow QA Testing path
<input type="checkbox"/>	2	<input checked="" type="checkbox"/>	Windows	Performance Optimization	Do Not Track MyApp Temp Files
<input type="checkbox"/>	3	<input type="checkbox"/>	Windows	Expert	Report execution of potential risk files id
<input type="checkbox"/>	4	<input type="checkbox"/>	Windows	Expert	Block execution of malicious files identifi
<input type="checkbox"/>	5	<input checked="" type="checkbox"/>	Windows	Expert	Examine powershell script contents

In addition, rules are grouped by server on the management server rules pages. The management server rules list will be expanded by default, but you can collapse and expand the rules for any server using the arrow button to the left of its name.

Software Rules

[Updateers](#)
[Rapid Configs](#)
[Publishers](#)
[Users](#)
[Directories](#)
[Files](#)
[Custom](#)
[Memory](#)
[Registry](#)
[Scripts](#)

Servers:
Saved Views:
Group By:

[Show Filters](#) | [Show Columns](#) | [Export to CSV](#) | [Refresh Table](#)

Search: Automatically apply

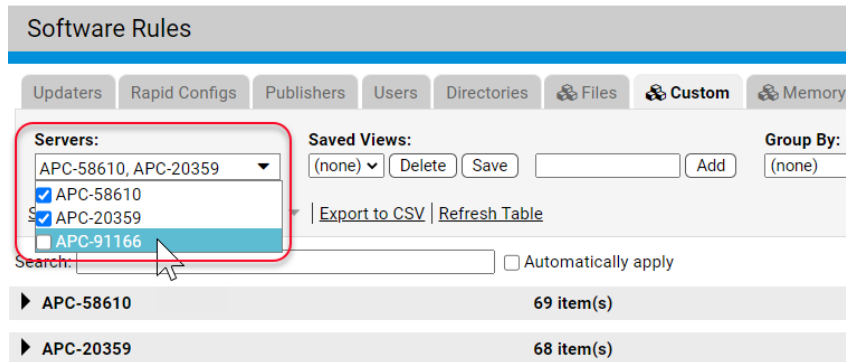
▶ APC-58610 69 item(s)
▶ APC-20359 68 item(s)
▼ APC-91166 1 item(s)

Showing 1 out of 1 item(s)

<input type="checkbox"/> Select 1	Rank ▲	Status	Platform	Rule Type	Name
<input type="checkbox"/>	1	<input checked="" type="checkbox"/>	Windows	Trusted Path	Allow QA Testing path

Showing 1 out of 1 item(s)

By default, the rules from all servers are accessible in a rules table on the management server. However, you can use the Servers menu to add or remove servers from the view. Uncheck the box next to any server you do not want included on the page.



Copying Existing Rules to Other Servers


In addition to sending a rule to multiple servers when you create it, you can copy one or more existing rules from any of the unified servers to any other unified servers. You must be logged in to the management server to do this, but the actual source of a copied rule can be the management server or one of the client servers. This process uses a wizard similar to the one used to save a new rule to multiple servers.

Notes


- If you copy a rule to a server that already has a local version of that same rule, the local version is replaced by the unified version.
- There is a *Copy this rule* command on the right menu on the Edit Rules page for Custom, Registry, and Memory Rules. This is for making copies of the rule on the same server – it does not copy rules to other servers for Unified Management.

To copy rules to another server under Unified Management:

1. In the rules table page, check the box(es) next to the rule(s) you want to copy.

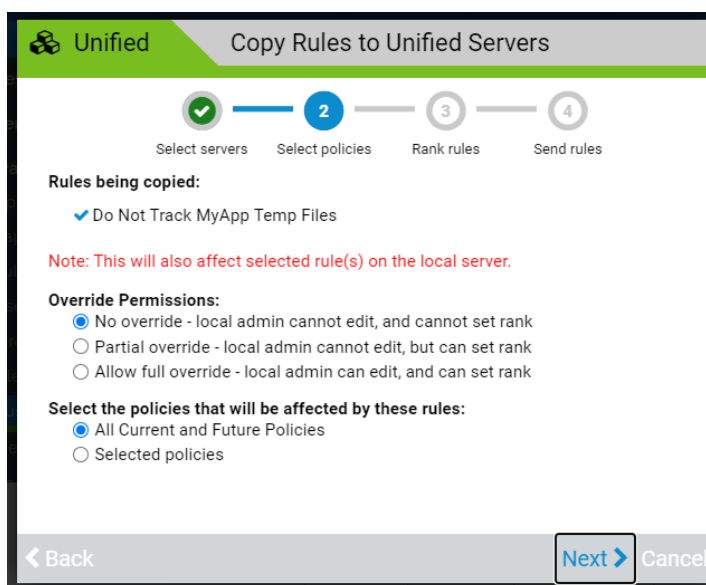


2. Click the **Copy to Unified Servers** button. This opens the Copy Rules to Unified Servers wizard.



3. On the first screen, you specify which servers will receive the rules. **All servers** is the default, but you can use the checkboxes to specify that some servers receive the rules and some don't. When you have specified which servers the rule applies to, click **Next** in the bottom right corner of the screen.

Note: Although the server that is the source of the rule is listed on this screen, its check box is grayed out on this screen and cannot be deselected since it already has the rule.



4. The next screen allows you to specify two fields:
 - a. **Override Permissions** – You can allow or block administrators who don't have permission to use or configure Unified Management from modifying the copied

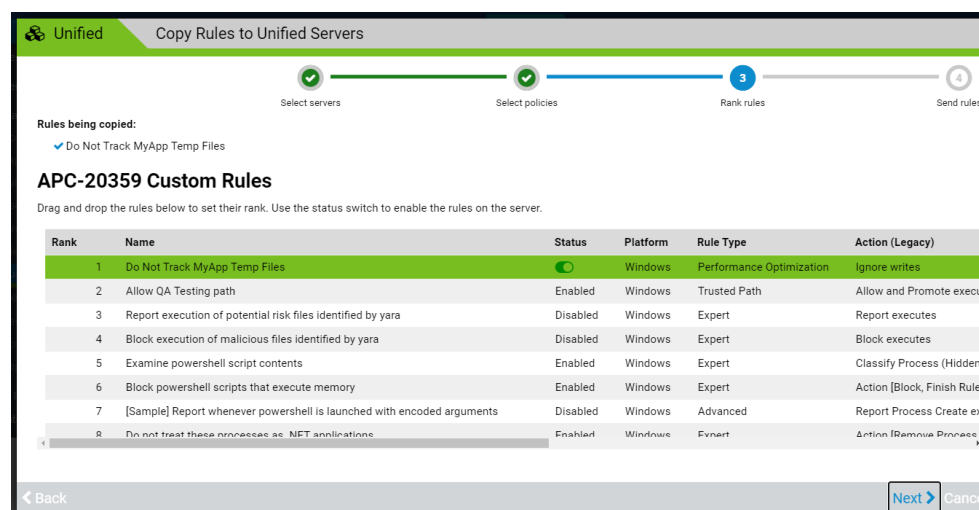
rules on their server. The choices include full override (editing rules and changing their rank), changing rank only, or no override.

Note: This permission has no effect on users that have Use Unified Management or Configure Unified Management permissions.


- b. **Policies** – You can specify that the rules apply to all policies on all servers, or you can choose individual policies to be affected by the rules.

Note: If you recently created, edited, or deleted policies on a remote system, the policy list from a remote server might not be updated on the management server immediately. With normal connectivity, policy lists from remote systems should be up to date on the management server within one minute, but slower networks might increase this gap.

- c. When you have made your override and policy choices, click **Next**.



5. The next screen allows you to customize the rank of the rules on each client server. You also can change the status of each rule from enabled to disabled, or vice versa. By default, the rules you are copying are ranked at the top of the table on the client server. You can drag the new rules to a different position in the list.
 - a. Review, and if you choose, change the rank of the new rules on each server as it appears in the wizards. Also, if you want to change the status of a rule between enabled and disabled, use the Status field toggle button. When you have made any necessary changes for this server, click **Next** to move to the next server.
 - b. When the rule rank screen for the last server appears, the Next button changes to **Send Rules**. If you are satisfied with your configuration, click this button.
6. After you click Send Rules, a progress screen shows the status of the rule copy operations. Unless errors are encountered, it will eventually show that all rules were copied to all chosen servers. When the progress bar shows Complete, click Finish to exit the wizard and return to either the Add/Edit Rule page (if you chose Save) or the table page for this rule type (if you chose Save and Exit).



If errors occur during the rule saving and distribution process, they will be shown with red exclamation marks instead of blue checkmarks in the final page of the wizard. Depending upon the issue found, correcting it might be possible by simply repeating the procedure above. Some issues, however, such as connectivity failures, might require remediation on one or more of the servers.

After a rule is copied to other servers, its row appears highlighted in green on the rules pages for all servers under unified management, including the management server itself.

▼ APC-58610		69 item(s)									
		Add Custom Rule		Copy to Unified Servers...		Export Rules		Import Rules		Showing 69 out of 69 item(s)	
<input type="checkbox"/> Select 54	Rank ▲	Status	Platform	Rule Type	Name						
<input type="checkbox"/>	1	<input checked="" type="checkbox"/>	Windows	Trusted Path	Allow QA Testing path						
<input type="checkbox"/>	2	<input checked="" type="checkbox"/>	Windows	Performance Optimization	Do Not Track MyApp Temp Files						
<input type="checkbox"/>	3	<input type="checkbox"/>	Windows	Expert	Report execution of potential risk files ide						
<input type="checkbox"/>	4	<input type="checkbox"/>	Windows	Expert	Block execution of malicious files identifi						

Editing a Unified Rule


Editing a unified rule is very similar to creating one. Users with sufficient permission can edit any field, including the rule name. As with Add Rule wizard, clicking Save or Save and Exit brings up a Save Rule wizard that allows you to modify some of the rule fields on each server.

Among the options when editing unified rules is to change the servers to which the rule applies. Changing the Servers field may create a situation in which you are *updating* a rule on some servers and *deleting* the rule from others. The action being performed on each server is reported in the Save Rule wizard.

If you edit the Servers field and uncheck the box for the management server, the rule is sent to any other servers that are still checked in the wizard, but the rule becomes local on each machine.

To edit a unified rule:

1. In the console menu, choose **Rules > Software Rules**. The Software Rules page appears.
2. On the Software Rules page, click the tab for the type of rule you want to edit, and click the View Details button next to the rule you want to edit.
3. On the Edit Rule page for that rule, make your changes and then click either **Save** (to remain on the Edit Rule page when finished) or **Save and Exit** (to return to the rule table page when finished). This opens the Save Rule wizard, which has a single screen showing the progress of rule updates on each server.



Notice in this example that while the rule was updated and sent to the first two servers, the last checkmark indicates that the rule was deleted from the final server. This would happen if that server was unchecked during the rule editing process.

4. Check the **Finish** button in the right bottom corner of the wizard to close it.

Permissions Needed for Unified Rule Editing

A user can edit a unified rule under the following conditions:

- A user must have Manage Custom/Registry/Memory Rules permissions to edit a unified rule.
- A user with User Unified Management or Configure Unified Management permission can edit a unified rule regardless of its override settings.
- A user without Use Unified Management or Configure Unified Management permission, but with Manage Custom/Registry/Memory Rules permission, can edit a unified rule if the rule's override setting is Full Override.
- A user without Use Unified Management or Configure Unified Management permission cannot edit a unified rule if the rule's override setting is No Override.

Changing a Unified Rule to a Local Rule

You might distribute some rules via Unified Management that are mandatory for all users and all servers. Other rules sent to managed servers might be suggestions or recommendations that you are willing to let local administrators to override. The ability of a local administrator to override a rule depends upon the override settings on the rule itself and the roles and permissions of the user attempting the change. Unified rules become local – that is, no longer subject to Unified Management – under the following circumstances:

- If a rule is edited on a client server, a confirmation box appears asking whether you want to override the (unified) rule on this server. If you answer Yes, the rule becomes local on that server.
- If a rule is edited on the management server, and the Servers setting is changed such that one or more client servers are no longer included, the rule is deleted from the client servers. However, if the rule is edited to not include the management server, it becomes local on any client servers that still have it.

Disabling and Enabling Unified Rules

Unified rules can be disabled and enabled on both the management server and the client servers. There are two ways to disable and re-enable unified rules:

- You can disable and enable a rule on an individual server using the toggle switch on the Rules table page.
- On the management server, you can disable and enable a rule on all unified servers using the Disable on All Servers and Enable on All Unified Servers menu choices on the Action menu of the Edit Rule page. These commands appear only on the management server.

Managing Unified File Rules from File Table and Details Pages

The Files tab of the Software Rules page shows all of the approvals and bans created at your site for specific individual files. These rules identify specific files by hash or optionally by file name (for bans only). The fundamentals of file rules are described in [“File-Specific Rules: Approvals and Bans”](#) on page 301.

Unified approvals and bans can be created in the following ways:

- From the Software Rules Files tab, open the Add File Rule page and enter the hash for a single file; for bans, you also have the option of using the file name or a specific path. The procedure for this is described in [“Managing Unified Rules from the Software Rules Page”](#) on page 792.
- From a File Details or File Instance Details page, you can choose one of the approval or ban commands on the Actions menu to create a rule for a single file.
- In a table of files (e.g., the File Catalog), you can check one or more files and choose one of the approval or ban commands on the Action menu to create one or more rules.

Note

On the File Rules tab, you can import a list of file hashes to create multiple approvals or bans at once, even if the files represented by the hashes do not yet exist on a server. However, the import dialog does not allow you to specify that these hashes are banned or approved on all servers under Unified Management. If you want to use important hashes for unified rules, you can first import the hashes to one server and then use the procedure described in [“Copying Existing Rules to Other Servers”](#) on page 798.

You can check boxes for one or more files on the Files page and use the Action menu to change their state. For a single file, the right menu on the File Details and File Instance Details pages provides the same options. These menus provide the following choices for unified management of file rules:

- **Approve Globally on all Unified Servers** – Immediately creates a hash-based rule globally approving the file(s) for all computers managed by all unified servers – no configuration is necessary.
- **Ban Globally on all Unified Servers** – Immediately creates an active hash ban for the file(s) on all computers managed by all unified servers – no configuration is necessary.
- **Approve by Policy** – Opens the Add Rule page with Approval as the Rule Type, allows you to choose both policies and servers to which this rule applies, and allows you to choose whether to allow local administrators to override this rule. You can edit other parameters, such as the rule name and its description.
- **Ban by Policy** – Opens the Add Rule page with Ban as the Rule Type, allows you to choose both policies and servers to which this rule applies, and allows you to choose whether to allow local administrators to override this rule. You can edit other parameters, such as choosing to make the rule a report-only ban.
- **Approve on Unified Servers** – Opens the Copy Rules to Unified Servers wizard, on which you can choose the servers and policies to which you want the file(s) approved, and also choose whether to allow local administrators to override this rule. No other rule parameters can be changed when you use this command.
- **Ban on Unified Servers** – Opens the Copy Rules to Unified Servers wizard, on which you can choose the servers and policies to which you want the file(s) banned, and also choose whether to allow local administrators to override this rule. No other rule parameters can be changed when you use this command.

- **Remove Approval or Ban from all Unified Servers** – Opens the Delete Rule progress wizard, deletes the rule from all connected servers, and returns the file(s) to an unapproved state.

Notes

- All of the Unified Management choices on the Files page Action menu complete their actions by showing a progress wizard page. This page reports the success or failure of rule updates on each server, and if the rule was not copied to a server for some reason, alerts you to the error condition.
- If rule for a file that you are trying to approve or ban already exists on any server, a dialog box appears to allow you to choose whether to override the existing rule.

Once you create an approval or ban rule, it appears on the File Rules page. Once a file rule has been created, all of the other unified rule management functions, such as copying rules or making them local, use the procedures described in [“Managing Unified Rules from the Software Rules Page”](#) on page 792.

Chapter 28

Monitoring System Health

This chapter introduces the System Health page, which provides App Control administrators with the ability to monitor the health and performance of the App Control Server.

Sections

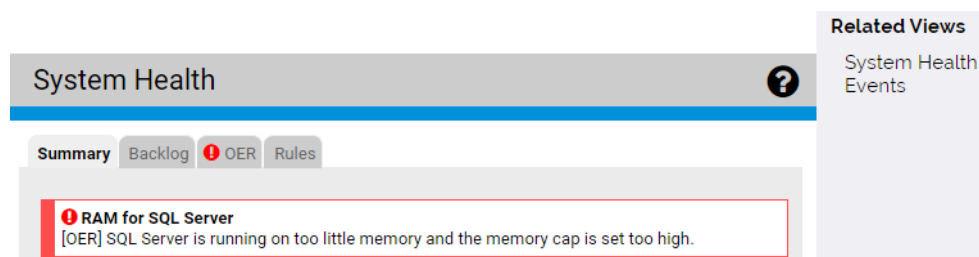
Topic	Page
Overview	807
Enabling System Health Indicators	808
Viewing the System Health Page	809
System Health Alerts	812
System Health Events	813

Overview

The System Health page provides App Control administrators with the ability to monitor factors that affect the performance of the App Control Server. It displays the output of *Health Indicators* that can warn you about problems on the App Control Server, the SQL Server, or the environment as a whole.

For example, your servers might not be in compliance with the *Operating Environment Requirements* guidelines for the number of rules or endpoints being managed. This could occur because you added more endpoints with App Control Agents to your environment. Another cause of a system health deterioration might be a change in your hardware environment, such as a change in disk capacity or RAM.

The System Health page can help you see these trends before they become a serious problem, so that you can either remedy them yourself or contact Carbon Black Support for guidance. Knowing that all of the monitored factors are healthy can also be helpful.



The System Health page contains different tab views showing the results of analysis by different health indicators. The first tab will show the overall Health Summary, which will include brief headlines for any triggered health indicators. Other tabs contain one or more related indicators. The information on the tabs may be presented as a graph, a table, simple text, or a combination of formats.

Health indicators provide feedback on critical or borderline conditions in several different ways:

- **System Health Page Triggered Indicators** – When an indicator detects that your server has an issue that affects its health, the System Health page displays a red icon and highlighting, as shown above. If a yellow icon and highlighting appears for an indicator, the factor it is reporting on is in a borderline but not critical state.
- **Alerts** – There is a built-in alert for each tab on the System Health page to warn when your system is not in compliance with App Control *Operating Environment Requirements* or other required configuration, and you can also create an alert to be triggered when a health indicator changes its severity level.
- **Events** – When the severity level of a health indicator changes, an event is recorded by the server and made available through Syslog output.

The Health Indicators displayed on the System Health page are delivered to the App Control Server through Carbon Black File Reputation. This cloud service not only delivers the initial set of indicators needed to enable the System Health page but also keeps your server up to date with any changes to existing indicators as well as new indicators that will add to your view of system health. Carbon Black File Reputation must be connected to your server for the System Health indicators to function.

Enabling System Health Indicators

System health indicators are provided by Carbon Black File Reputation. The SRS must be enabled for the initial download of Health Indicators from the cloud, and it is also used to update existing indicators when necessary and to add new indicators as they are developed.

In addition to enabling Carbon Black File Reputation on your App Control Server, you must enable Health Indicators “updates” setting on the Advanced Options tab of the System Configuration page. This switch enables both initial downloading and later updates of health indicators.

To enable System Health Indicators on a App Control Server:

1. On the console menu, click the configuration (gear) icon and choose **System Configuration** and click the **Licensing** tab.
2. On the Licensing tab, check to see whether the Carbon Black File Reputation Activation panel shows that the SRS is activated. If it is not activated, follow the activation instructions in [“Activating Carbon Black File Reputation”](#) on page 763.
3. When Carbon Black File Reputation is activated, click on the **Advanced Options** tab on the System Configuration page and click the **Edit** button at the bottom of the page.
4. In the Software Rule Options panel, check the box for **Health Indicators**. When the page is saved, this automatically downloads Health Indicators from Carbon Black File Reputation and also updates them as necessary.
5. Click the **Update** button at the bottom of the page. Download of the Health Indicators is scheduled and begins shortly. See [“Viewing the System Health Page”](#) for a description of what you will see once this feature has been activated.

Once System Health Indicators have been activated, they begin to download to your server from Carbon Black File Reputation. Depending upon your connection speed and other server activities, this might take one or two hours.

Disabling System Health Indicators

If you need to disable Health Indicator updates, go to the System Configuration page Advanced Options tab, click the **Edit** button, uncheck the **Health Indicators** box, and click the **Update** button. When you go to the System Health page after disabling indicators, the page will show a message about the feature not being enabled.

Viewing the System Health Page

Users must have a login account with “View system health indicators” permission to view this page. Accounts in the Administrators group have this permission by default. See [“Managing Console User Roles”](#) on page 101 if you need to configure another user for access to this page.


Note

The System Health page illustrations shown here were accurate at the time of publication, but because Health Indicators are delivered and updated from the Carbon Black File Reputation, the exact indicators on any tab, and their appearance and content on your version of the console, may be different.


To view the System Health page:

- On the console menu, click the configuration (gear) icon and choose **System Health**.

If the initial download of indicators is not complete, only the Summary tab appears on the System Health page, and it displays a message that the download is pending.



When the download is complete, the page has tabs for each of the available health views. The first tab shows the overall health Summary, which will include any triggered health indicators and also report any conditions that prevent proper operation of Health Indicators. The Summary tab also indicates when there are no triggered indicators.



The other tabs show the results of analysis by the different health indicators – these will vary as new indicators are made available through the SRS.

The screenshot shows the 'System Health' interface. On the left, there are three indicator cards: 'OER Summary' (Some operating environment requirements did not pass validation), 'RAM for SQL Server' (SQL Server is running on too little memory and the memory cap is set too high), and 'Resources on SQL Server Instance' (Checks whether Cb Protection is getting adequate IO and CPU resources from the SQL Server instance). A red arrow points from the 'Resources on SQL Server Instance' card to the details view on the right. The details view shows two circular gauges for 'IO' and 'CPU', both at 100% usage for 'Cb Protection'. Below the gauges, a text box explains that the health indicator checks for resource sharing with other databases. At the bottom, it shows 'OK since Jul 13 2020 02:04:41 PM' and 'Last evaluated Jul 17 2020 02:05:13 PM'. A second red arrow points from the bottom of the details view back to the indicator card.

Indicators on this tab

Details for the selected indicator

There may be multiple indicators on a System Health tab view, such as in the example above. Indicators are shown on the left side of the page, and the view on the right side of the tab shows the details for the selected indicator, which is offset to the left.

There are several conditions under which no indicators are shown on the System Health page:

- There are no health indicators available because Carbon Black File Reputation is disabled.
- The System Health feature not enabled.
- The download of health indicators from Carbon Black File Reputation is still pending.


Note

If the App Control Server determines that an indicator is not relevant in your environment, that indicator might not be displayed. In addition, if the server determines that *none* of the indicators on a tab is relevant in your environment, the tab itself will not be displayed. A change in your environment could make an indicator or even an entire tab disappear from the System Health page.

Navigating on the System Health Page

There are several ways to change views or drill down for additional information on the System Health page:

- **Change Tab Views** – You can click on any of the tabs to change the set of indicators you are viewing.
- **Change Indicator Shown on a Tab** – If there are multiple indicators on a tab, you can click on one of the other indicators on the left to change to the details shown on the right.

- **Links in the Details** – Some indicator details include links to additional information. For example, the OER Summary details view includes a link to the current *Operating Environment Requirements* document for App Control on the Carbon Black Customer Portal – note that you must have your portal login to complete navigation to this link.
- **Reload Indicator Details** – Use the reload button  in the upper right corner of the indicator details view if you want to be certain that you are viewing the most current information. Most indicators are re-evaluated every 24 hours; the OER Summary indicator is re-evaluated every 15 minutes.
- **System Health Events** – The Related Views menu on the System Health page includes a links to the Events page, filtered to show only events related to health indicators. See “[System Health Events](#)” on page 813 for more about these events.

Health Indicator State

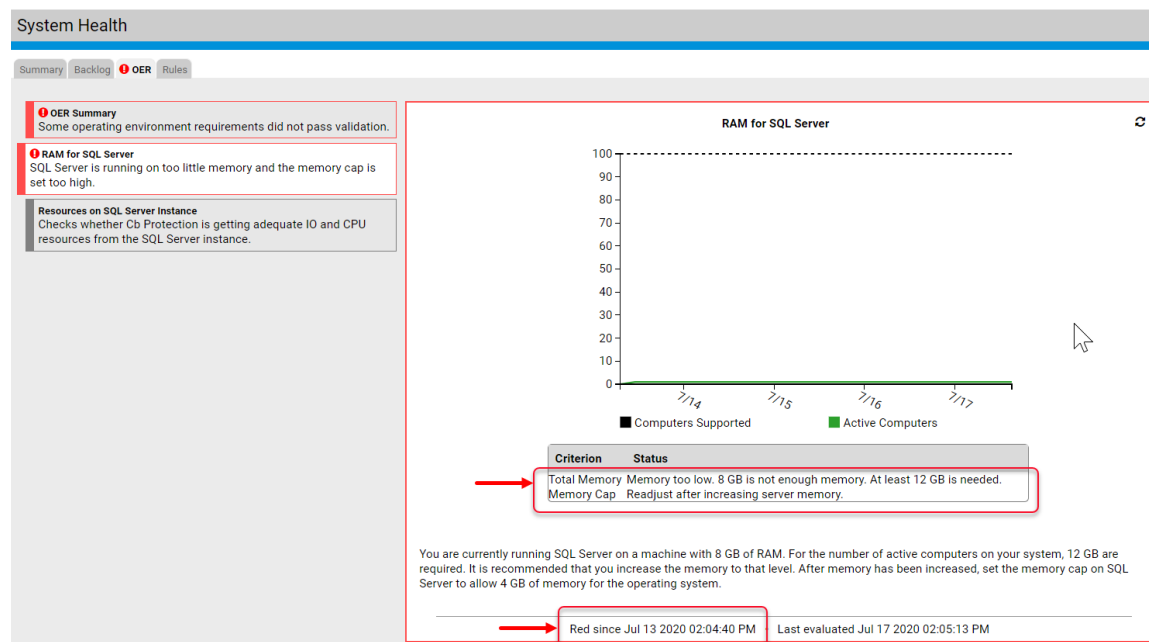
Health indicators are color coded to show the state of the parameter or resource they monitor. The colors and their states are:

- **Gray** – When an indicator is gray, the condition it is monitoring is healthy (or is strictly informational and does not have a health rating) and no action is required.
- **Yellow** – When an indicator is yellow, the condition it is monitoring is borderline, and you should follow up and take action if necessary.
- **Red** – When an indicator is red, the condition it is monitoring is in a critical state and action is required.

Note

If the App Control Server determines that an indicator is not relevant in your environment, that indicator might not be displayed.

When an indicator is showing anything less than a healthy state, it is considered “triggered”, and a triggered indicator shows its state in multiple locations, including the list of indicators on the Summary tab view, and on the tab itself and the list of indicators for the view in which the indicator appears. In addition, a description of the problem and the length of time it has existed appear in the details section of the tab view, and an alert is triggered to warn of the issue.



System Health Alerts

Alerts can notify you of system health issues. These notifications appear in the console and (if enabled) are emailed to subscribers. See [“Using App Control Alerts”](#) on page 608 for full details on alerts. These alerts may notify you of system health issues:

- **System Health Backlog Alert** – This built-in alert is triggered when the server has an excessively large backlog of unprocessed file operation messages from agents, and so cannot report the latest file states for all agents. File operations include all operations that can happen on a file of interest, such as adding the file, editing it, deleting it, or copying it.
- **System Health Environment Alert** – This built-in indicator is triggered when certain conditions in your App Control environment are reaching their limits. For example, it is triggered if one or more of the server's pathnames and filenames tables approach or exceed their size limit.
- **System Health Infrastructure Configuration Alert** – This built-in alert is triggered when any factors reported on the Infrastructure Configuration tab of the System Health page are out of compliance. It is permanently enabled, but it cannot be triggered if Health Indicators are not enabled.
- **System Health OER Alert** – This built-in alert is triggered when the environment for the server is out of compliance with certain specifications in the App Control *Operating Environment Requirements*. It is permanently enabled, although it cannot be triggered if Health Indicators are not enabled.
- **System Health Product Configuration Alert** – This built-in alert is triggered when your App Control Server is not at the latest version.
- **System Health Rules Alert** – This built-in alert is triggered when any rules on your server are ineffective. This includes rules you can correct through the console and rules that require the assistance of Carbon Black support for correction.

Note

System Health alerts only appear and can only be triggered if System Health Indicators are enabled on the Advanced tab of the System Configuration page and the related indicator has been downloaded to the server. If present, they are always enabled.

System Health Events

The App Control Server records several different events related to health indicators. You can view these events in the console, set up rules in a SIEM that respond to these events, and trigger Alerts or Event Rules based on them. There are event subtypes to inform you of changes in the indicators themselves: *Health indicator created*, *Health indicator changed*, and *Health indicator deleted*.

The event most likely to be of interest for monitoring system health is the *Health indicator severity change* subtype. An event indicating a severity change from lower to higher means that some element in your App Control environment needs your attention. On the other hand, a decrease in severity can let you know that a remediation you performed was successful. Increases in severity trigger events whose severity is Warning. Decreases in severity trigger events whose severity is Info.

The Description fields for the severity change event provide details about why the event was triggered. It also includes descriptions of the state of newly created indicators. [Table 121](#) shows the conditions that trigger *Health indicator severity change* events.

Table 121: Health Indicator Severity Change Event Conditions

Condition	Description
Indicator condition is no longer healthy	Health indicator <name> has gone to severity <severity level>. Check the health indicator for more details.
Severity increased from yellow to red	Health indicator <name> has increased in severity from <old severity> to <new severity>. Check the health indicator for more details.
Severity decreased to yellow	Health indicator <name> has decreased in severity from <old severity> to <new severity>.
Triggered indicator is now healthy	Health indicator <name> is now healthy.
New indicator condition is unhealthy	Newly created health indicator <name> has severity <severity level>. Check the health indicator for more details.
New indicator condition is healthy	Newly created health indicator <name> is healthy.

To view health indicator events in the console:

1. On the console Events page, choose **Reports > Events**.

2. In the Saved View menu, choose **System Health History**.
3. Make any other adjustments you choose to the other table view parameters, such as Max Age.

Appendix A

Live Inventory SDK: Database Views

In addition to the access provided to the Live Inventory of files and computers through the console user interface, App Control provides public views into the database. You can create your own reporting and data analysis solutions through the use of these public views. This appendix describes the available read-only database views.

Creating your own custom reports using the external database views may be useful when you want to perform complex analysis of file and computer inventory data. The SDK also facilitates:

- A special combination of filters or a file grouping not provided in the App Control Console.
- Inquiries that perform faster when done through direct database access outside of the console user interface.
- Reports that run on a specific schedule and/or need their output integrated into third-party tools.

Note

App Control also includes the API, a RESTful API that provides programmers a way to write code that interacts with App Control, either using custom scripts or from other applications. See [Appendix B, “App Control API,”](#) for more information.

Performance Considerations

The external views provide read-only access to the database and are optimized to not interfere with other App Control Server tasks. The database server is a shared resource, however, and overall performance of the App Control Server might be affected by extensive querying of external views. Consider the following general suggestions:

- Avoid running queries that take more than two minutes to complete.
- Limit total time spent querying the external database to no more than 5% of total time (e.g., a few minutes each hour).
- If possible, run queries at a time of day when App Control Agents are not very active, especially avoiding times when agents are initializing.

Contact Carbon Black Support for assistance with performance issues.

Upgrading from a Previous Version

If you used these database views in a previous release, you may need to modify some queries to match changes in this release. In the tables for each view, changes since version 6.0.2 are indicated in the following ways:

- **New** fields are indicated with a solid delta (▲) next to the name if new for **7.0.0**, a solid diamond (◆) if new for **7.0.1**, and a solid star (★) if new for **7.2.0**. Note that some fields were introduced in different builds or patches of the same version.
- **Changed** fields (field name or its values) are indicated with an open delta symbol (Δ) next to the name if changed for **7.0.0** and an open diamond (◇) if changed for **7.0.1**. A **Change Note** in the Comments column describes what has changed. Note that some fields were changed in different builds or patches of the same version.
- **Removed** fields are noted in the introduction to each view table.

App Control supports agent installation on Mac, Linux and Windows computers, so any path-related field will have operating-system-specific syntax (including delimiters).

In addition, if you are upgrading from v6.0.2 or earlier, be aware of the following global changes in terminology, which affect many of the SDK values:

Table 122: Global Terminology Changes for Post-6.0.2 Releases

Category	6.0.2 Term	7.2.2 Term
File Status	Pending	Unapproved
	Approved (Custom)	Approved by Policy
	Banned (Custom)	Banned by Policy
Computer protection level	SecCon	Enforcement Level
Enforcement Level value	20-Lockdown	High (Block Unapproved)
	30-Block-and-Ask	Medium (Prompt Unapproved)
	40-Monitor	Low (Monitor Unapproved)
	60-Visibility Only	None (Visibility)
	80-Agent Disabled	None (Disabled)

Schema Overview: bit9_public

External views represent a de-normalized view of the App Control Server live inventory. These views are suitable for reporting and analysis using data cubes. Each exposed view uses the naming convention with the prefix “Ex” for “external,” and is in the schema **bit9_public** within the database **Das**.

Specifying a Schema User

You must provide a login name for the user to whom you want to grant access to the **bit9_public** schema. Use the following script to add this login name and login manually (after the App Control Server is installed). Replace *Domain* and *cbprotectionuser* with your own values for the appropriate Windows user:

```
CREATE LOGIN [Domain\cbprotectionuser] FROM WINDOWS WITH
DEFAULT_DATABASE=[Das]
GO
CREATE USER [Domain\cbprotectionuser] FOR LOGIN
[Domain\cbprotectionuser]
GO
USE [Das]
GO
GRANT SELECT ON SCHEMA :: dbo TO
[Domain\cbprotectionuser]
GO
GRANT EXECUTE ON SCHEMA :: dbo TO
[Domain\cbprotectionuser]
GO
ALTER AUTHORIZATION ON SCHEMA::bit9_public TO
[Domain\cbprotectionuser]
GO
```


Schema Views and Diagram

[Table 123](#) shows the views available in the schema. Detail about the data in each view is shown in the subsequent tables in this topic. The full schema diagram for bit9_public appears immediately after the table.


Table 123: Schema Views for bit9_public

View Name	Description	Primary Key	Foreign Keys
ExInfo	Public properties of servers and schema in the App Control Server environment	Property_Id	(None)
ExEvents	All events shown on the Events page	Event_Id	File_Catalog_Id, Root_File_Catalog_Id, Computer_Id
ExMeters	All executions of metered files	Event_Id	Computer_Id, File_Catalog_Id
ExComputers	Metadata of all computers	Computer_Id	(None)
ExFileCatalog	Metadata for all unique hashes	File_Catalog_Id	(None)
ExFileInstances	Metadata of all file instances on all computers	File_Instance_Id	File_Instance_Group_Id, Computer_Id, File_Catalog_Id
ExDeletedFileInstances	Metadata of all deleted file instances	Deleted_File_Instance_Id	File_Instance_Group_Id, Computer_Id, File_Catalog_Id
ExFileInstanceGroups	Metadata of all file instance groups	File_Instance_Group_Id	Computer_Id, File_Catalog_Id
ExApprovalRequests	All approval requests shown on the Approval Requests page	Approval_Request_Id	File_Catalog_Id, Process_File_Catalog_Id, Root_File_Catalog_Id, Computer_Id

Schema Diagram for bit9_public



Schema Diagram for bit9_public (continued)



Details of Database Views

ExComputers

The ExComputers view provides access to the metadata of all computers running the App Control Agent at your site. To see a list of this data for all computers in the console, choose **Assets > Computers** the console menu. To see this data for a single computer, click on the name of a computer on the Computers page.

Table 124: ExComputers View Details

Field Name	Data Type	Special Values	Comments
Computer_Id	int		Primary key
Audit_Date	nvarchar		Date and time when Computer Information was collected
Computer_Information	XML		A meta-field containing data (in XML format) about the computer including number of drives and free space on each; number, model and speed of processors; and total system RAM.
▲ Platform	varchar	'Windows', 'Mac', 'Linux'	
Memory_Size	int		Size (megabytes) of installed memory on this computer
Processor_Count	int		Number of processors on this computer
Processor_Speed	float		Speed of computer processor in MHz
IP_Address	varchar		Last recorded IP address of this computer. This can be either an IPv4 or IPv6 address.
Connected	varchar	'Yes', 'No'	'Yes' if agent on this computer is connected to the App Control Server
Days_Offline	int		Number of days this computer has been offline
Computer_Uninstalled	varchar	'Yes', 'No'	'Yes' if agent has been uninstalled from this computer
Computer_Deleted	varchar	'Yes', 'No'	'Yes' if computer has been deleted from the Computers list in the App Control Server
Date_Created	datetime		Date and time this computer first connected to the App Control Server
Policy_Mode	varchar	'Control', 'Visibility', 'Agent Disabled'	Mode of the policy this computer belongs to
Agent_Version	varchar		Version of the agent installed on this computer

Field Name	Data Type	Special Values	Comments
Poll_Date	varchar		Date and time this computer last connected to the App Control Server
Policy_Assignment	varchar	'Manual', 'Automatic'	How policy is assigned to this agent (automatic means it was assigned by Active Directory mapping)
Δ Upgrade_Status	varchar	'Up to date', 'Completed', 'Not supported', 'Scheduled', 'Waiting', 'Not requested', 'Agent uninstalled', 'Reboot required', 'Blocked', 'Upgrade requested', 'Unknown'	Current upgrade status of this agent Change Note: 'Upgrade requested' was added in 7.0.0.
Δ Policy_Status	varchar	'Policy out of date', 'Approvals out of date', 'Enforcement Level out of date', 'Out of date', 'Up to date'	Current policy status of this computer Change Note: Value 'Enforcement Level out of date' was 'SecCon out of date' in 6.0.2.
Computer	nvarchar		Name of this computer
Users	nvarchar		Comma-separated list of users that have ever logged on to this computer
Δ Enforcement_Level	nvarchar	'High (Block Unapproved)', 'Medium (Prompt Unapproved)', 'Low (Monitor Unapproved)', 'None (Visibility)', 'None (Disabled)'	Enforcement Level used when this computer is online Change Note: Enforcement_Level was Online_SecCon in 6.0.2. All values changed beginning with 7.0.0.
Δ Disconnected_Level	nvarchar	'High (Block Unapproved)', 'Medium (Prompt Unapproved)', 'Low (Monitor Unapproved)', 'None (Visibility)', 'None (Disabled)'	Enforcement Level used when this computer is offline Change Note: Disconnected_Level was Offline_SecCon in 6.0.2. All values changed beginning with 7.0.0.
Computer_Tag	nvarchar		Optional custom tag assigned to this computer
OS_Short_Name	nvarchar		Short name of this computer's OS
OS_Long_Name	nvarchar		Long name of this computer's OS

Field Name	Data Type	Special Values	Comments
Policy	nvarchar		Name of the last policy for this agent
Policy_Description	nvarchar		Description of the last policy this agent has joined
Machine_Model	nvarchar		Machine model of this computer
Processor_Model	nvarchar		Processor model of this computer
Upgrade_Error	nvarchar		Agent upgrade error (if any)
Synch_Percent	int		Progress of synchronization of this computer with the App Control Server (percent)
▲ Template	varchar	'Yes','No'	'Yes' if computer is a template. 'No' if it is not (includes clones and non-cloned computers).
▲ Template_Computer_Id	int		The ID of the parent template computer. If the value is 0, the computer does not have a template parent and is not a clone. If the value is non-zero, the computer is a clone.
▲ Virtualized	varchar	'Yes','No'	'Yes' if computer is a virtual machine. 'No' if it is not.
▲ Virtual_Platform	varchar		If Virtualized is 'Yes', the platform of the virtual machine. Currently, this is either 'VMware', 'Unknown', or blank.

ExInfo

The ExInfo view provides access to data about the App Control Server and public schema (this schema) versions as well as the address of the App Control Server and other servers in its environment.

Table 125: ExInfo View Details

Field Name	Data Type	Special Values	Comments
Property_Id	int		Primary Key
Name	nvarchar	'RPCServerAddress', 'Bit9ServerVersion', 'WebServerAddress', 'DBPublicSchemaVersion',	Name of the property
Value	nvarchar		Value of the property

ExMeters

The ExMeters view provides access to data on all executions of App Control meters, which monitor each time a specified file is executed, in your environment. To see this information as it is displayed in the console, choose **Tools > Meters** in the console menu and click on the View Details button next to any meter to see information about a specific meter.

Table 126: ExMeters View Details

Field Name	Data Type	Special Values	Comments
Event_Id	bigint		Foreign key into ExEvents table for event that correspond to this meter entry. Since this value is always unique, it can also serve as a primary key.
Computer_Id	int		Foreign key into ExComputers table for computer that corresponds to this meter entry.
File_Catalog_Id	int		Foreign key into ExFileCatalog table for file that corresponds to this meter entry
Timestamp	datetime		Date and time when this meter entry was generated
Name	nvarchar		Name of the meter
Description	nvarchar		Description of the meter
Data	nvarchar		Data associated with the meter (see “type” for interpretation of this field)
ΔType	int	2 = sha1 hash, 3 = md5 hash, 4 = file name, 5 = sha256 hash 6 = sha256 fuzzy hash	Type of the Data field. This defines how the meter was created. Change Note: Some previous versions of the documentation had incorrect numerical values for this field.
User_Name	nvarchar		Name of the user that created this meter

ExEvents

The ExEvents view provides access to all events that are displayable on the Events page. This includes events related to files discovered, files blocked, files approved, unapproved files executed, system management processes, and actions by console users. To see event data as it is displayed in the console, choose **Reports > Events** in console menu; this displays the Events page.

Table 127: ExEvents View Details

Field Name	Data Type	Special Values	Comments
Event_Id	bigint		Primary Key
Computer_Id	int		Foreign key into the ExComputers for computer that sent this event
File_Catalog_Id	int		Foreign key into the ExFileCatalog table for file associated with this event
Root_File_Catalog_Id	int		Foreign key into ExFileCatalog table for a root file associated with this event
▲ File_Name	nvarchar		Name of the file related to this event
▲ Path_Name	nvarchar		File path related to this event. Paths use the OS-specific delimiter for the agent on which the file is located.
Process	nvarchar		Name of the process associated with this event
▲ Process_File_Catalog_ID	int		Foreign key into ExFileCatalog table for the process associated with this event
Timestamp	datetime		Date and time (UTC) this event was generated
IP_Address	varchar		IP address of the endpoint that originated this event
Description	nvarchar		Event description
Priority	nvarchar	'Debug', 'Info', 'Notice', 'Warning', 'Error', 'Critical'	Event priority
Type	nvarchar		Event Type
Subtype	nvarchar		Event Subtype
User_Name	nvarchar		Name of the user associated with this event

Field Name	Data Type	Special Values	Comments
▲ Rule_Name	nvarchar		Name of the App Control rule that caused the event (block/prompt/report/approval)
◆ Ban_Name	nvarchar		Name of the hash or filename ban associated with the event (empty if the ban was not named); introduced in 7.0.1 Patch 3
◆ Updater_Name	nvarchar		If an updater is associated with the event, the name of the updater; introduced in 7.0.1 Patch 3
★ Indicator_Name	nvarchar		If a threat indicator is associated with the event, the name of the indicator
★ Received_Timestamp	datetime		Date and time (UTC) this event was received by the App Control Server
Command_Line	nvarchar		Command line for the process that attempted the action recorded by this event

ExFileCatalog

The ExFileCatalog view provides access to the metadata for all unique hashes of files discovered on your computers. To see this file data as it is displayed in the console, choose **Assets > Files** in the console menu and click on the File Catalog tab.

Table 128: ExFileCatalog View Details

Field Name	Data Type	Special Values	Comments
File_Catalog_Id	int		Primary Key
Prevalence	int		Prevalence of this file – number of computers that currently have this file
First_Created	datetime		Date and time when this file was first created
File_Size	bigint		Size of this file in bytes
File_Type	varchar	'Application', 'Package', 'Script File', 'Supporting File', 'Other', 'Unknown', 'Unrecognized Executed File'	Type of this file
MD5	char		MD5 hash of this file
Sha1	char		SHA1 hash of this file

Field Name	Data Type	Special Values	Comments
Sha256	char		SHA256 hash of this file (see Sha256_Hash_Type for interpretation of this field)
Sha256_Hash_Type	int	5 = regular hash 6 = MSI fuzzy hash	Type of the Sha256_Hash. See “SHA-256” on page 239 for more details.
First_Seen_Computer_id	int		Foreign key into ExComputers table for computer on which the file was first seen
First_Seen_Name	nvarchar		File name where this file was first seen on any computer
First_Seen_Path	nvarchar		Path where this file was first seen on any computer. Uses the path delimiter for the OS of the first-seen computer.
Product_Name	nvarchar		Product name of this file
Product_Version	nvarchar		Product version of this file
Publisher	nvarchar		Publisher of this file (if file is signed with certificate)
▲◇ Publisher_State	nvarchar	‘Approved’, ‘Approved by Policy’, ‘Unapproved’, ‘Banned’, ‘Banned by Policy’	State of this publisher (if available); “none” for unsigned files Change Note: Banned and Banned by Policy were added during 7.0.1.
▲ Publisher_State_Reason	nvarchar	‘Manual’, ‘Reputation’, ‘Imported’, ‘External (API)’, ‘Unknown’	Reason the file’s publisher is approved
Publisher_or_Company	nvarchar		Publisher (if available) or Company name (if no publisher info) of this file
Company	nvarchar		Company name of this file
Installed_Program_Name	nvarchar		If this file was an installer, the name of its installed program (i.e., its name on the Add/Remove Programs page in Windows). No value for Mac or Linux files.
Trust	int	-1 = unknown, [0 – 10] valid values	Trust of this file; maximum = 10
Trust_Messages	nvarchar		More information associated with this file’s trust

Field Name	Data Type	Special Values	Comments
Threat	nvarchar	'0 - Clean', '1 - Potential risk', '2 - Malicious', 'Unknown'	Threat level of this file
Category	nvarchar		Category of this file
▲ State	nvarchar	'Unapproved', 'Approved', 'Banned', 'Approved by Policy', 'Banned by Policy', 'Mixed'	Effective global file state for this file
△ File_State	nvarchar	'Unapproved', 'Approved', 'Banned', 'Approved by Policy', 'Banned by Policy', 'Mixed'	Global file state for this file Change Note: Was Global_State in 6.0.2. Also, values changed beginning with 7.0.0.
△ File_Flags	nvarchar	Comma-separated combination of one or more of the following: 'Installer', 'Not installer (Override)', 'Installer (Override)', 'Report Only Ban'	Global file flags for this file Change Note: File_Flags was Global_Flags in 6.0.2. Also, the value 'Report Only Ban' was 'Test Banned' in 6.0.2.
▲ File_State_Reason	nvarchar	'Manual', 'Trusted Directory', 'Reputation', 'Imported', 'External (API)', 'Unknown'	Reason for the approval state of this file
▲ Approved_By_Reputation	varchar	'Yes', 'No'	Was this file approved because of its file or publisher Trust and Threat ratings in CB Reputation
Reputation_Enabled	varchar	'Yes', 'No'	Is reputation-based approval is enabled for this file
◆ Certificate_Hash	char		App Control-proprietary hash that provides unique identifier for this certificate.
◆ Certificate_State	nvarchar	'Unapproved', 'Approved', 'Banned', 'Approved by Policy', 'Banned by Policy'	Global State of the certificate for this file. Note: Invalid certificates are 'Unapproved' in this field. Unsigned certificates are null.
◆ Certificate_State_Reason	nvarchar	'Manual', 'External (API)'	State reason of the certificate (same as Publisher State Reason)
★ File_Extension	nvarchar		Extension of first seen file with this hash

ExFileInstances

The ExFileInstances view provides access to the metadata for each instance of each hash found on each computer at your site. To see this file data displayed in the console, choose **Assets > Files** in the console menu and click on the File on Computers tab. To see the complete File Instance details for any one file, from the Files on Computers tab, click on the View Details button next to the file.

Change Note: In Beginning with v7.0.1, the fields **Initialized** and **Top_Level** were removed from this view and added to **ExFileInstanceGroups**.

Table 129: ExFileInstances View Details

Field Name	Data Type	Special Values	Comments
File_Instance_Id	bigint		Primary Key
File_Instance_Group_Id	int		Foreign key into ExFileInstanceGroups table for group that contains this file
File_Catalog_Id	int		Foreign key into ExFileCatalog table for details about this file
Computer_Id	int		Foreign key into ExComputers table for computer that has this file
Date_Created	datetime		Date and time (UTC) when file was created
File_Name	nvarchar		Name of this file
Path_Name	nvarchar		Path of this file. Uses OS-specific delimiter for the agent where the file is located.
Executed	varchar	'Yes', 'No'	'Yes' if this file was ever executed
Δ Local_State	nvarchar	'Unapproved', 'Approved', 'Banned'	Local state of this file Change Note: 'Unapproved' was 'Pending' in 6.0.2.

Field Name	Data Type	Special Values	Comments
△ Detailed_Local_State	nvarchar	'Approved (Not Persisted)', 'Unapproved (Persisted)', 'Banned by Hash', 'Locally Approved', 'Banned by Name', 'Banned by Name (Report Only)', 'Locally Approved (Auto)', 'Approved as Installer', 'Approved', 'Approved as Installer (Top Level)', 'Banned by Hash (Report Only)', 'Unapproved'	Detailed local state of this file Change Note: 'Unapproved' was 'Pending' in 6.0.2. 'Unapproved (Persisted)' was 'Pending (Persisted)' in 6.0.2.
◆ Detached_Publisher	nvarchar		Name of the detached publisher. Note that embedded publishers can be retrieved through a join with ExFileCatalog.
◆ Detached_Publisher_State	nvarchar	'Approved', 'Approved by Policy', 'Unapproved', 'Banned', 'Banned by Policy'	State of the detached publisher (if available); "none" for unsigned files
◆ Detached_Publisher_State_Reason	nvarchar	'Manual', 'Imported', 'External (API)', 'Unknown'	Reason for the state of this file's publisher
Detached_Certificate_Hash	char		App Control-proprietary hash of the detached certificate. Note that embedded certificates can be retrieved through a join with ExFileCatalog.
◆ Detached_Certificate_State	nvarchar	'Unapproved', 'Approved', 'Banned', 'Approved by Policy', 'Banned by Policy'	Global state of the detached certificate Note: Invalid certificates will be 'Unapproved' in this field. Unsigned certificates are null.
◆ Detached_Certificate_State_Reason	nvarchar	'Manual', 'Imported', 'External (API)', 'Unknown'	Reason for the state of the file's detached certificate (same as Publisher State reason)

ExDeletedFileInstances

The ExDeletedFileInstances view provides access to the metadata for each deleted file instance on each computer at your site. The App Control Server keeps track of only last deleted instance of each unique file name on each computer. This means that, if same file was created and deleted multiple times, only last deleted instance will be listed.

Change Note: Beginning with v7.0.1, the fields **Initialized** and **Top_Level** were removed from this view and added to **ExFileInstanceGroups**.

Table 130: ExDeletedFileInstances View Details

Field Name	Data Type	Special Values	Comments
Deleted_File_Instance_Id	bigint		Primary Key
File_Instance_Group_Id	int		Foreign key into ExFileInstanceGroups table for group that contains this file
File_Catalog_Id	int		Foreign key into ExFileCatalog table for details about this file
Computer_Id	int		Foreign key into ExComputers table for computer that has this file
Date_Created	datetime		Date and time (UTC) when the file was created
Date_Deleted	datetime		Date and time (UTC) when file was deleted
File_Name	nvarchar		Name of this file
Path_Name	nvarchar		Path of the file. Uses the OS-specific delimiter for the agent that had the file
◆ Detached_Publisher	nvarchar		Name of the detached publisher. Embedded publishers can be retrieved through a join with ExFileCatalog.
◆ Detached_Publisher_State	nvarchar	'Approved', 'Approved by Policy', 'Unapproved', 'Banned', 'Banned by Policy'	State of the detached publisher (if available); "none" for unsigned files
◆ Detached_Publisher_State_Reason	nvarchar	'Manual', 'Reputation', 'Imported', 'External (API)', 'Unknown'	Reason for the state of this file's publisher

Field Name	Data Type	Special Values	Comments
◆ Detached_Certificate_Hash	char		App Control-proprietary hash of the detached certificate. Embedded certificates can be retrieved through a join with ExFileCatalog
◆ Detached_Certificate_State	nvarchar	'Unapproved', 'Approved', 'Banned', 'Approved by Policy', 'Banned by Policy'	Global state of the detached certificate. Note: Invalid certificates are 'Unapproved' in this field. Unsigned certificates will be null.
◆ Detached_Certificate_State_Reason	nvarchar	'Manual', 'Imported', 'External (API)', 'Unknown'	Reason for the state of the file's detached certificate (same as Publisher State reason)

ExFileInstanceGroups

The ExFileInstanceGroups view provides access to the metadata for file instance groups found on your computers. File instance groups are groups of files associated with one primary root file, usually their installer but sometimes a file from which they were copied.

Table 131: ExFileInstanceGroups

Field Name	Data Type	Special Values	Comments
File_Instance_Group_Id	Int		Primary Key
File_Catalog_Id	Int		Foreign key into ExFileCatalog table for details about root file of this group
Computer_Id	Int		Foreign key into ExComputers table for computer that has this file group
Date_Created	datetime		Date and time (UTC) when this file group was created
Group_Type	int	0 – initialized file 1 – top-level file 2 – file installed by process 3 – file installed by installer and can be found in add/remove programs	How the group was identified by App Control

Field Name	Data Type	Special Values	Comments
Path_Name	nvarchar		Path that corresponds to the root file of this group. Paths use the OS-specific delimiter for the agent on which the file is located.
User_Name	nvarchar		User that created this group
File_Name	nvarchar		File name that corresponds to the root file of this group
Installed_Program_Name	nvarchar		If this file was an installer, this will be the installation name
◆ Initialized	varchar	'Yes', 'No'	'Yes' if the files in this group were found during initialization
◆ Top_Level	varchar	'Yes', 'No'	'Yes' if this group represents a top-level file that was not generated through an installer. 'No' if files in this group were part of an installation.

ExApprovalRequests

The ExApprovalRequests view provides access to the work flow for approval requests created by users through the App Control notifier when attempts to execute a file are blocked. This includes approval requests for files that are completely blocked from running and justifications for cases when the user responded to a prompt by allowing a file to run.

Note: This entire view was new beginning in v7.2.0.

Table 132: ExApprovalRequests

Field Name	Data Type	Special Values	Comments
Approval_Request_Id	Int		Primary Key
File_Catalog_Id	Int		Foreign key into ExFileCatalog table for file for which approval request was created
Process_File_Catalog_Id	Int		Foreign key into ExFileCatalog table for process associated with file for which approval request was created
Root_File_Catalog_Id	Int		Foreign key into ExFileCatalog table for root file associated with file for which approval request was created
Computer_Id	Int		Foreign key into ExComputers table for computer on which approval was requested

Field Name	Data Type	Special Values	Comments
Date_Created	datetime		Date and time (UTC) when this approval request was created
Date_Modified	datetime		Date and time (UTC) when this approval request was last modified
Last_Modified_By	nvarchar		User that last modified this approval request
Enforcement_Level	nvarchar	Valid enforcement levels	Enforcement Level of agent when file was blocked
Resolution	nvarchar	Not Resolved, Resolved - Publisher, Resolved - Installer, Resolved - Approved, Resolved - Rule Change, Resolved – Other, Rejected	Resolution of request
Request_Type	nvarchar	Approval, Justification	Type of approval request: Approval if file was blocked, Justification if file triggered a user-choice prompt
Requestor	nvarchar		User that created approval request on the agent
Requestor_Comments	nvarchar		Comments provided during approval request creation
Requestor_Email	nvarchar		Email address provided during approval request creation
Priority_Text	nvarchar	High, Low, Medium	Priority assigned to this request by the user creating it
Resolution_Comments	nvarchar		Comments provided during resolution of request
Status	nvarchar	New, Closed	Current status of the request
Policy	nvarchar		Name of the Policy where agent was in when block happened
Multiple_Blocks	nvarchar	Yes, No	Whether multiple blocks happened on this endpoint and hash
File_Name	nvarchar		Name of the file that was blocked on the agent
Path_Name	nvarchar		Path of the file that was blocked on the agent

Field Name	Data Type	Special Values	Comments
Process	nvarchar		Full path to the process that wrote the file that was blocked on the agent

Sample Queries

The following examples show some of the types of queries you can make with the Live Inventory SDK. Note that each query must use the **das** database.

Listing Malicious Files

If you have Carbon Black File Reputation enabled, you can use the following query to get a listing of the file names and prevalence of all malicious files determined to be on your systems that run the App Control Agent:

```
USE das
SELECT First_Seen_Path, First_Seen_Name, Sha256, Threat,
       Trust, Prevalence
FROM bit9_public.ExFileCatalog
WHERE Threat IN ('2 - Malicious', '1 - Potential risk')
ORDER BY First_Seen_Path, First_Seen_Name
```

If you run this query and there is data available, you will see output similar to the following (formatting will vary):

First_Seen_Path	First_Seen_Name	Sha256	Threat	Trust	Prev.
c:\temp\folder1	myfileapp.exe	46b8d0bc3a4db843 3fb66543c1ec03bd1 e24e0198228ac702 4c0a15658bf04fd	1 - Potential risk	2	1
c:\documents and settings\rjones	numbergen.exe	552e68dcd6c2a4d6 bf9c9dbf278967e29 04cd624c23c0aad58 c430ed7fa75acd	1 - Potential risk	1	1
c:\documents and settings\bsmith	makemess.exe	4d9ab91f5e1efbc5 abcd6ec9a0a63452 35a54cf05d6241a30 4e3bf3b40d4668	1 - Potential risk	3	1
c:\hp\bin	endprocess.exe	1effc62134ab95d29 7c34959752311e1f7 f433d07810da65b23 3bf7241ada68ad	1 - Potential risk	3	13
c:\program files\mywebapp\	f4dothis.dll	abcdea797736654a e4f74eef7371d018c 3463f24cf78aea92d afe51c7a858f19	2 - Malicious	0	1
c:\jobfiles	myway.exe	23451271912da7b6 8b407c77381ab1ff3 b59b37c1e4d9f1e41 7a1d0fcc9270dd	2 - Malicious	0	1

Listing App Control Agent Systems by Policy and Enforcement Level

You can use the following query to determine how many systems are running the agent and group the results by Policy and Enforcement Level:

```
USE das
SELECT Policy, Enforcement_Level, Disconnected_Level,
COUNT(*)
  AS Computer_Count
FROM bit9_public.ExComputers
GROUP BY Policy, Enforcement_Level, Disconnected_Level
ORDER BY Policy
```

If you run this query and there is data available, you will see output similar to the following (formatting will vary):

Policy	Connected_Enforcement_Level	Disconnected_Enforcement_Level	Count
Agent Disabled	None (Disabled)	None (Disabled)	3
Research Team	Medium (Prompt Unapproved)	Medium (Prompt Unapproved)	6
Default Policy	None (Visibility)	None (Visibility)	1
General Office	High (Block Unapproved)	High (Block Unapproved)	49
Guest Policy	High (Block Unapproved)	High (Block Unapproved)	1
IT Group	Low (Monitor Unapproved)	Low (Monitor Unapproved)	11

Listing New Unapproved Files by Policy

You can use the following query to determine how many new unapproved files have appeared during the past 24 hours and group the results by Policy:

```
USE das
SELECT Policy, COUNT(*) FROM bit9_public.ExFileInstances fi
  JOIN bit9_public.ExComputers c
  ON c.Computer_Id = fi.Computer_Id
WHERE fi.Date_Created > DATEADD(day, -1, GetUTCDate()) AND
  Local_State = 'Unapproved'
GROUP BY Policy
ORDER BY COUNT(*) DESC
```

If you run this query and there is data available, you will see output similar to the following (formatting will vary):

Policy	New Unapproved File Count
Research Team	529
General Office	101
IT Group	257

Listing New Unapproved Files by Computer and Policy

To determine how many new unapproved files have appeared during the past 24 hours and group the results by Computer and Policy:

```
USE das
SELECT c.Computer, c.Policy, COUNT(*) as Unapproved_Count
FROM bit9_public.ExFileInstances fi
  JOIN bit9_public.ExComputers c
  ON c.Computer_Id = fi.Computer_Id
WHERE fi.Date_Created > DATEADD(day, -1, GetUTCDate()) AND
Local_State = 'Unapproved'
GROUP BY c.Computer, c.Policy
ORDER BY COUNT(*) DESC
```

If you run this query and there is data available, you will see output similar to the following (formatting will vary):

Computer Name	Policy	New Unapproved File Count
MYCORP\DESKTOP-3	Research Team	307
MYCORP\LAPTOP-1	General Office	215
MYCORP\LAPTOP-4	Research Team	32
MYCORP\DESKTOP-8	IT Group	3
MYCORP\DESKTOP-10	General Office	2
MYCORP\LAPTOP-7	General Office	1

Appendix B

App Control API

The App Control API is intended for programmers who want to write code to interact with App Control, either using custom scripts or from other applications. It is a RESTful API that can be consumed over HTTPS protocol using any language that can create get URI requests and post/put JSON requests as well as interpret JSON responses.

Actions performed through the App Control API create an audit trail just as the same action performed from the console would. The appropriate API user taking the action is referenced in event.

There are two sections in this appendix:

- **API Authentication and Access Control** – This describes how to create an API user account and get the API Token necessary for API authentication of clients. It also describes how to configure permissions for the login accounts needed for such access.
- **Available Objects** – This is a listing and brief description of the objects you can access through the App Control API.

This appendix is a summary only. The full API documentation is available in two locations:

- Documentation for the REST API in your version of App Control is available through the console at <https://<yourseveraddress>/api/bit9platform/v1>.
- The App Control REST API documentation can also be found at <https://developer.carbonblack.com/reference/enterprise-protection>.

The following additional resources are available for CB API developers:

- Carbon Black provides a Python module that developers can use for easy access to the REST APIs for App Control, Carbon Black EDR, and VMware Carbon Black Cloud. The documentation for this module is available at <https://cbapi.readthedocs.io>.
- The source code for the CB API module (cbapi) for Carbon Black products is located at <https://github.com/carbonblack/cbapi-python>.
- Tutorials, blogs, and other CB API resources for App Control, Carbon Black EDR, and VMware Carbon Black Cloud are available on the Carbon Black Developer Network site at <https://developer.carbonblack.com>

Note

App Control also includes a Live Inventory SDK, which provides read-only public views into the database. You can create your own reporting and data analysis solutions through the use of these public views. See [Appendix A, “Live Inventory SDK: Database Views,”](#) for more information.

Overview

The current version of the App Control API is v1. All API calls are based at the following address: **https://<your server name>/api/bit9platform/v1**

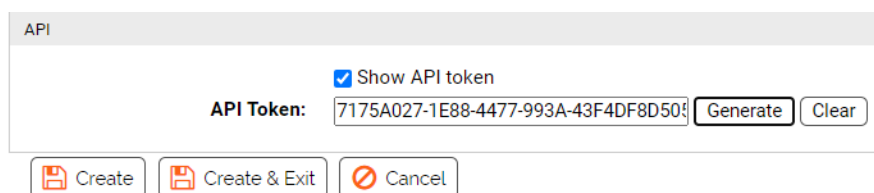
API Authentication and Access Control

App Control APIs are authenticated through an API token for the login account of the currently logged in console user. This token has to be placed inside each HTTP request's 'X-Auth-Token' header.

For access control, the best practice is to have a separate console user for each API client, with the minimum required access controls. However, the API client must have access permissions similar to what would be required to access the same objects through the console. For example, if an API client needs to access the 'event' object, the user associated with an API token used in the client must have "View events" permission. See the full API documentation on GitHub for the permissions necessary for using each object and "User Role Permissions" on page 106 for the steps needed to add or remove permissions. See also "Using the App Control API to Add a Connector" on page 842 if you intend to use the API to add a connector to App Control.

To create an API user and get its API token:

1. Review the App Control API documentation on your server or GitHub to determine the permissions needed for your API client.
2. On the console menu, click the configuration (gear) icon and choose **Login Accounts**.
3. Click the **User Roles** tab and then the **Add User Role** button to open the Add User Role page.
4. On the Add User Role page, provide a Name (for example, "API Connector Extensions"), add a Description if you choose, and check the box for each permission needed for your client. Note that some permissions depend upon others, and you must have permission to view an object if you also intend to change it.
5. When you have configured the group, click the **Enabled** button in the Status line and click the **Create & Exit** button at the bottom of the page.
6. Click the **Users** tab, and on the *Login Accounts: Users* page, click **Add User**.
7. On the Add User page, provide a user name (for example, "API HashBanScript") and password, and choose the User Role you created above.
8. Provide any other information you choose in the other fields.
9. At the bottom of the page, check the **Show API token** box and then click the **Generate** button. A string of characters appears in the API Token box.



The screenshot shows a web interface for generating an API token. At the top, there is a tab labeled "API". Below it, there is a checkbox labeled "Show API token" which is checked. Underneath, the text "API Token:" is followed by a text input field containing the alphanumeric string "7175A027-1E88-4477-993A-43F4DF8D509". To the right of the input field are two buttons: "Generate" and "Clear". At the bottom of the form, there are three buttons: "Create" (with a floppy disk icon), "Create & Exit" (with a floppy disk icon), and "Cancel" (with a red circle and slash icon).

10. Copy the API Token to a location in which you can copy it to your API code. Also make a record of the login user name the code is associated with.

11. Click the **Save** button at the bottom of the page.

Important

Do not use the API Token in any way that displays it in clear text. If the API Token is compromised, open the Edit Login Account page for the API user, check the Show API token box, click **Generate** to produce a new token, and then click **Save**. Then use the new token for authentication.

To disable API access for a user that currently has permission, follow the steps above but click **Clear** instead of Generate. If server hardening is required, all API access should be removed.

Available Objects

You can access the following App Control objects through the App Control API (see the full documentation for the actual object name in the API and to see which are read only):

- Approval Requests and Justifications – Access the work flow for approval requests and justifications created when users respond to a notifier.
- Certificates – Access publisher certificates found on endpoints and their state.
- Computers – Access computer-related properties for App Control Agents, change policies, upgrade agents, convert a computer to a VDI template, change debugging properties, take other advanced actions.
- Connectors – Access the configuration for network security connectors integrated with App Control.
- Events – Access events recorded by App Control.
- Files Analysis – Access files sent to network connectors for analysis; request or cancel analysis of a file.
- File Catalog – Access the record of all unique files found by agents, including metadata related to the files.
- File Instances – Access the live file inventory (Files on Computers) for files on all agent-managed systems; locally approve files.
- Deleted File Instances – Access the inventory of deleted files on all agent-managed systems.
- File Instance Groups – Access the record of file groups in the Files on Computers inventory.
- File Rules – Access rules related to unique files; create and edit Approvals and Bans.
- Files Uploaded from Agents – Access the record of files uploaded from agents to the App Control Server; request or cancel uploads.
- Metered Executions – Access the record of file executions tracked by a Meter.
- Notifications – Push notifications from a network connector (services and appliances) to the App Control Server.
- Notifier – Access notifiers that are used when a file action is blocked because of a rule.
- Pending Analysis – Access all pending analysis requests for a given external connector.
- Policy – Access policy information.

- Publisher – Access publisher information; change publisher state (Banning or Approving).
- Server Configuration – Access configuration properties for the server.
- Server Performance – Access server performance statistics.
- Updaters – Access updater information; enable or disable updaters.

Using the App Control API to Add a Connector

The App Control Connector allows you to integrate the App Control Server with one or more network security devices or services so that the external source can provide threat notifications to the server and the server can send files to the external source for detonation and/or analysis. Several connector integrations are built into the App Control Server and configurable through settings already in the console.

The App Control API provides a way to extend App Control Connector capabilities to devices and services not built into the current App Control Server. When correctly implemented, these connections add the notification and analysis capabilities, and the user interface elements necessary to configure and use them. The interface for configuring a new connector appears on the Connectors tab of the System Configuration page in the console. On this tab, you can make the following configuration choices:

- **Integration Enabled** - This checkbox enables and disables notification integration for this connector. If this box is unchecked, file analysis will also be disabled automatically.
- **File Analysis Enabled** - This checkbox enables and disables file analysis for this connector, the connector has this capability. This setting appears only if the connector allows file analysis.
- **Upload Location** – If File Analysis is enabled, you can customize the upload location for this connector. This option appears only if the connector allows file analysis.

When configured, the new connector appears in the console interface wherever built-in connectors would appear. For example, if a connected device or service allows analysis, the new connector appears on the Action menu of the Files pages as an analysis option. See [Appendix C, “App Control Connector,”](#) for a full description of the connector capabilities and user interface.

Notes

- To add an integration with a custom network security device or service, you must activate the “Extend connectors through API” permission for the login account that will be used for access to App Control. Completing the configuration for a connector also requires permission to view and manage system configuration.
- Once the connector is implemented through the App Control API, you do not need a special license for access to its notification features. However, to upload files from a App Control-managed computer to a third-party devices or service for analysis, you do need the separately licensed File Upload feature.

Appendix C

App Control Connector

This chapter provides instructions for configuring and using the Connector, which integrates the App Control Server with one or more network security devices or services.

Sections

Topic	Page
Overview	844
Enabling Carbon Black Inspection	845
Enabling Palo Alto Networks Integration	846
Enabling VMware Carbon Black Cloud Integration	853
Enabling Check Point Integration	857
Enabling Console Account Permissions	867
External Notifications	867
Banning Externally Reported Malware	880
Analysis of Suspicious Files on Endpoints	883
Logging of Connector-related Events	886

Overview

The App Control Connector allows you to integrate the App Control Server with one or more network security devices or services, including:

- Carbon Black Inspection
- Check Point ThreatCloud Emulation Service
- Check Point Threat Emulation Private Cloud Appliances
- Palo Alto Networks™ firewalls
- Palo Alto Networks WildFire™ public and private cloud services

Note

In addition to the supported devices and services, you can integrate other services, such as Lastline, with App Control using the App Control API. These integrations are examples of API capabilities only, and not currently supported. See [Appendix B, “App Control API,”](#) for instructions on enabling API access and authentication.

By integrating these systems with App Control, when a connected device or service detects malware on an enterprise network, App Control’s real-time endpoint sensor and recorder automatically confirms the location and scope of the threat, accelerating incident response and remediation. In addition, suspicious files found by the App Control endpoint sensor can be uploaded to one of the connected appliances or network security analysis providers for further analysis.

The App Control Connector adds the following capabilities to what the App Control Server and network security devices or services offer individually:

- **External Notifications** – Notifications provided by the connected sources appear as “External Notifications” in the App Control Console, correlated with App Control endpoint data to provide immediate visibility into the priority of the alert and the scope of any infection. See [“External Notifications”](#) on page 867 for details.
- **File Banning** – Malware reported by connected sources can be manually or automatically banned by App Control. See [“Banning Externally Reported Malware”](#) on page 880 for details.
- **Registry Control** – Suspicious file or registry activity reported by connected sources can be reported or restricted by App Control custom rules. See [“Special Rules for Reporting or Banning Malware”](#) on page 881 for details.
- **Analysis of Suspicious Files** – Suspicious files discovered on endpoints by App Control Agents can be sent to connected services for analysis. See [“Analysis of Suspicious Files on Endpoints”](#) on page 883 for details.
- **Event Logging** – Events related to external notification or analysis and reported to the App Control Server become part of the App Control event log, and are also available as Syslog output. See [“Logging of Connector-related Events”](#) on page 886 for details.
- **Event Rules** – Rules can be defined that use file-related events to take actions. For example, a rule could send any newly discovered file in the App Control Server inventory to Check Point Threat Emulation Cloud Service or appliance, or to the Palo Alto Networks WildFire cloud, for analysis. Another rule might be defined that

automatically bans any file reported as malicious in an external notification. Or if App Control detects that repeated malware infections reported by a connected third-party tool, an Event Rule could be created to ban that parent process if it is not used for any required function; see [“Event Rules”](#) on page 520.

Preparing to use the Connector

The Connector is a separately licensed option of App Control. To use the connector features you must do the following:

- Install the App Control Server with the appropriate license for the Connector, or add the license after installation.
- Configure the App Control Server and any of the connected devices as described in this appendix so that they can communicate with each other.
- If you are integrating App Control with CB Inspection, install the separately licensed CB Inspection Connector.
- Confirm that one or more console user accounts have privileges related to the Connector. Accounts in the Administrator group have these permissions by default. See [“User Role Permissions”](#) on page 106 for details.

Note

Contact your Carbon Black representative to determine which versions of Check Point and Palo Alto Networks products are compatible with the App Control Connector for Network Security Devices and App Control Server hardware/software requirements.

Enabling Carbon Black Inspection

CB Inspection provides integrated analysis services from Carbon Black and its partners. App Control Servers can send files from the endpoints they monitor to CB Inspection for analysis. Once a file is analyzed, the analysis results are sent back to the server that requested them.

The analysis includes executing Windows 32-bit and 64-bit PE executables in a sandbox environment. Analysis results also include all of the metadata available for the file from the Carbon Black platform components that have seen it.

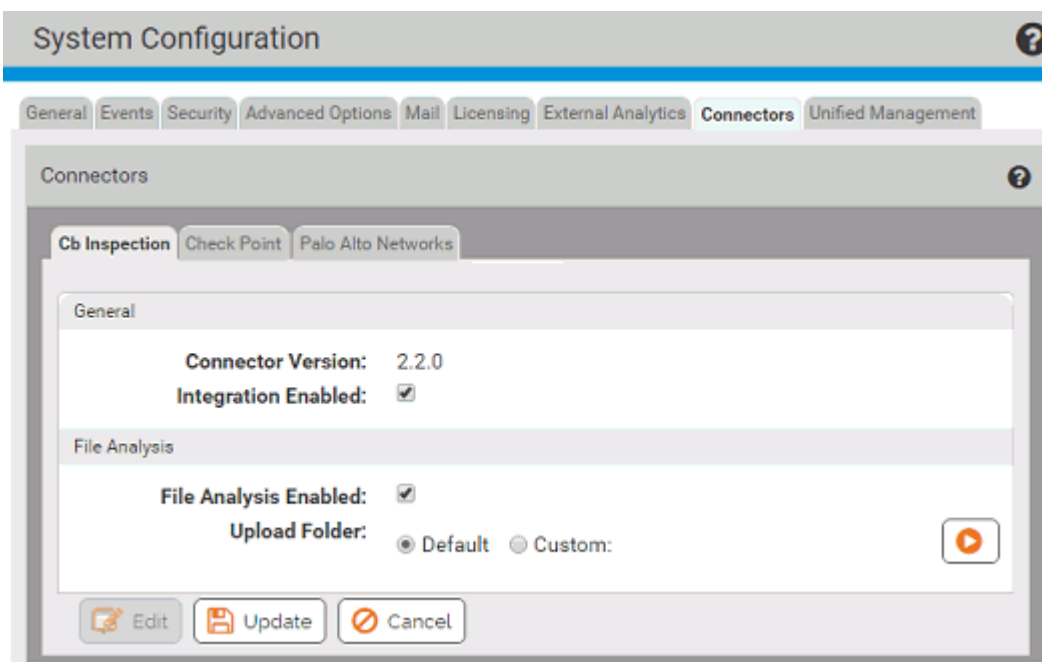
You install a separately licensed "connector" to enable use of CB Inspection with a App Control Server. The connector allows you to upload files, either by manually selecting each file or by creating Event Rules that automatically upload files matching your specifications. Once uploaded, files undergo static and dynamic analysis. Scoring information, along with details about observed behaviors, are returned via the connector, allowing you to examine any applicable telemetry from the service and drive policy based on results.

Installation and use of CB Inspection is described in the [CB Inspection User Guide](#), available on the Carbon Black [User Exchange](#). Contact your Carbon Black representative if you would like to purchase and install CB Inspection.

CB Inspection Connector Configuration

The CB Inspection Connector is enabled immediately when its installation is completed, and should be ready for use without further steps. Configuration of this Connector is shown on the Connectors tab of the App Control Console System Configuration page. There, you can do the following:

- view the version of the currently installed CB Inspection Connector
- disable and then re-enable the Connector
- change the location on the server to which agent files are uploaded



Enabling Palo Alto Networks Integration

Enabling the App Control Connector for Palo Alto Networks involves configuration steps on both the App Control Server and the Palo Alto Networks appliance. You can enable integration for notifications, for Wildfire analysis of files in your App Control inventory, or for both notification and analysis

Integrating Palo Alto Networks Appliances for Notifications

Notifications from multiple Palo Alto Networks appliances can be integrated with a App Control Server.

To enable integration of Palo Alto Networks alerts with App Control:

1. Confirm that the Palo Alto Networks firewall and App Control Server system are able to contact each other.
2. On each Palo Alto Networks appliance you plan to integrate with App Control, create a local user account with administrative read-only permissions for the App Control integration.
3. On the console menu, click the configuration (gear) icon and choose **System Configuration** and click on the **Connectors** tab and then the **Palo Alto Networks** tab.
4. Click the **Edit** button at the bottom of the page.

The screenshot shows the 'System Configuration' page in the App Control console. The 'Connectors' tab is selected, and the 'Palo Alto Networks' connector is active. The configuration is divided into three sections: General, Appliances, and File Analysis. In the General section, the 'Integration Enabled' checkbox is currently unchecked. In the Appliances section, the 'Initial Import' field is set to 7 days, and the 'Enable Additional Lookups' checkbox is unchecked. In the File Analysis section, the 'File Analysis' radio buttons are set to 'Off'. At the bottom of the page, there are three buttons: 'Edit', 'Update', and 'Cancel'.

5. Check the **Integration Enabled** checkbox. This is the primary switch for the Palo Alto Networks integration.
6. In the Appliances panel, go to the Initial Import field and enter the number of days of historical notification data you want to import to App Control. The default value is 7 days. This value affects only appliances from which no data has been received yet. If App Control already has data from an appliance, data import will resume with the time of the last data received.

7. If you want to get a full malware report for each notification that has a file reference, check the *Enable Additional Lookups* box.
Important: The Initial Import you configured will happen all at one time. If *Enable Additional Lookups* is enabled, be sure to choose an Initial Import time period that will not cause the number of WildFire cloud queries to exceed your licensed daily limit.
8. The Appliances section of the Palo Alto Networks Integration Settings page allows you to add and delete appliances to the App Control integration.

The screenshot shows the 'Appliances' configuration page. At the top, 'Integration Enabled' is checked. Under 'Appliances', 'Initial Import' is set to 7 Previous Day(s) and 'Enable Additional Lookups' is unchecked. Below this is a table of appliances. The first appliance has the following settings: Address (empty), Import Threat Log (checked), Threat Log Filter (severity neq informational and severity neq low and), Import WildFire Log (checked), WildFire Log Filter (category neq benign), User Name (empty), and Password (empty). A 'Test' button is located at the bottom of the appliance configuration area.

For each appliance, click **Add New** and provide the following information:

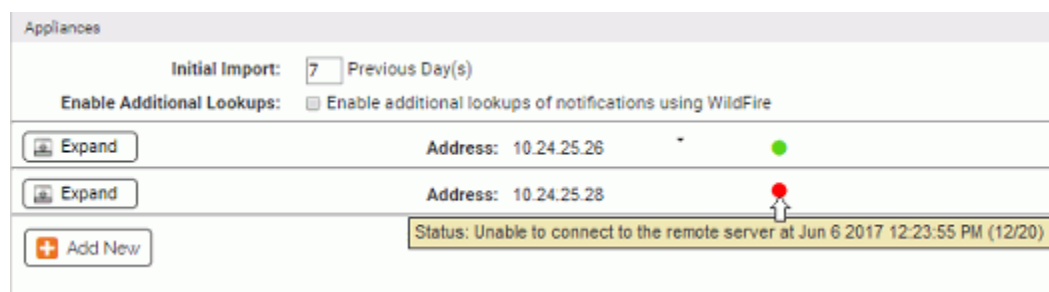
- a. **Address** – The IP address of the appliance.
 - b. **Import Threat Log** – Checking this box activates import of Threat Log data from this appliance to the External Notifications page on the App Control Server.
 - c. **Threat Log Filter** – This text field shows the filtering of Threat Log data from the appliance to limit what appears in the External Notifications table. By default, the filter eliminates notifications whose severity level is *informational*, *low*, or *medium*. You can modify the filter to get the notifications you choose; the filter syntax is the same as that used in the Palo Alto Networks Console.
 - d. **Import WildFire Log** – Checking this box activates import of WildFire Log data from this appliance to the External Notifications page on the App Control Server.
 - e. **WildFire Log Filter** – This text field shows the filtering of WildFire Log data from the appliance to limit what appears in the External Notifications table. By default, the filter eliminates notifications whose category is *benign*. You can modify the filter to get only the notifications you choose.
 - f. **User Name** and **Password** – In the User Name and Password boxes, enter the user name and password for the unique account you created in [Step 2](#).
Note: Do not use your console login credentials for either Palo Alto Networks or App Control Console in these fields.
 - g. When you have provided the address and credentials, click the **Test** button to confirm that this appliance is accessible, the credentials are appropriate, and the filter syntax is valid before saving the appliance specification.
9. If you are integrating more appliances, click the **Add New** button and provide the necessary information for another appliance.

10. The settings in the File Analysis panel determines whether files from agents managed by the App Control Server can be sent to the WildFire cloud for analysis. If you plan to enable WildFire file analysis, see [“Integrating with the WildFire Cloud for Analysis”](#) for information on configuring this section.
11. When you finish configuring the integration (and if all appliances pass the Test above), click the **Update** button at the bottom of the page.

When the notifications integration is complete, Palo Alto Networks notifications begin to appear in the App Control Console. To see the notifications, choose **Reports > External Notifications** on the console menu. You might not see notifications immediately because of pre-filtering of appliance notifications. If notifications do not appear at all, check the Events page in the console for Server errors, and also check *Bit9\Parity Server\Reporter\ParityReporter.log* for possible details of interest.

See [“External Notifications”](#) on page 867 for a full description of the notification features, including the types of notifications pre-filtered from appearing in the console.

Palo Alto Networks Notification Appliance Status



Once configured, the status of each Palo Alto Networks appliance sending notifications to your server is displayed on the System Configuration/Connectors/Palo Alto Networks integration page in the App Control Console. In the Appliances panel, a status indicator appears next to the address of each appliance:

- A green circle indicates that there are no issues with that appliance’s integration
- A red circle indicates a problem, and in this case, an error message will appear with the indicator.
- A light blue circle indicates that the appliance is de-activated.

For each status indicator, a tooltip provides additional information when you hover the cursor over the circle.

Modifying or Deleting an Appliance Integration

For any existing appliance integration, you can edit the configuration, for example, to enable or disable one or both of the Log imports to the App Control Server. You also can delete an appliance integration.

The screenshot shows a configuration form for an appliance integration. At the top left, there are two buttons: 'Collapse' (with a minus icon) and 'Delete' (with a trash icon). The form contains the following fields and controls:

- Address:** A text input field containing '11.22.33.44' with a green status indicator to its right.
- Import Threat Log:** A checkbox that is checked.
- Threat Log Filter:** A text input field containing 'severity neq informational and severity neq low and'.
- Import WildFire Log:** A checkbox that is checked.
- WildFire Log Filter:** A text input field containing 'category neq benign'.
- User Name:** A text input field containing 'JaneDoe'.
- Password:** A password input field with masked characters (dots).
- Test:** A button with a play icon and the text 'Test'.

To delete or edit a Palo Alto Networks appliance integration:

1. On the Connectors/Palo Alto Networks tab, click the **Edit** button at the bottom of the page.
2. Click the **Expand** button next to the appliance you want to edit. The configuration for that appliance is displayed.
3. If you want to delete an appliance from your integration, click the **Delete** button next to its address.
4. If you want to enable or disable Threat Log or WildFire log data imports to App Control, check or uncheck the appropriate checkbox.
5. If you want to change the filter for one of the log imports, edit the text in the corresponding Filter box.
6. If you have enabled an import or modified the filter, click **Test** to confirm that the appliance is accessible and the filter syntax is valid.
7. Click **Update** to save your changes.

Integrating with the WildFire Cloud for Analysis

You can enable uploading of files for analysis from App Control-managed systems to either the Palo Alto Networks WildFire public cloud or a locally installed WildFire private cloud device. In either case, the file is analyzed and the analysis results are sent back to the App Control Console.

When the WildFire integration is complete, new menu choices appear on App Control Console pages that show tables of files or file details. These *Analyze with Palo Alto Networks WildFire* commands allow uploading of files to the WildFire cloud. See [“Analysis of Suspicious Files on Endpoints”](#) on page 883 for full details on how to upload files to the WildFire cloud and how to view the results of WildFire analysis.

Note

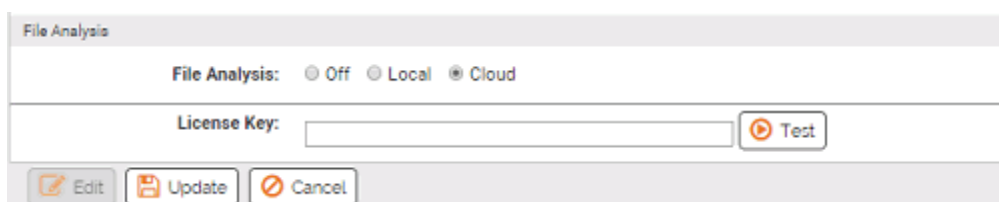
You can connect a App Control Server to one or more WildFire private cloud appliances or to the WildFire public cloud, but you cannot mix private and public cloud analysis.

Integrating with the WildFire Public Cloud

When you integrate App Control with the WildFire public cloud, the connection can either be direct or via proxy.

To enable file uploads to the WildFire public cloud for analysis:

1. If you are not already on this page, in the App Control Console, click the configuration (gear) icon and choose **System Configuration**, click the **Connectors** tab and then the **Palo Alto Networks** tab, and click the **Edit** button at the bottom of the page.
2. In the File Analysis panel, click the **Cloud** button and then enter your WildFire license key in the WildFire Key field.
Note: Files sent by App Control for analysis by the WildFire cloud service are subject to the limits in the WildFire license key.



3. Click the **Test** button next to the WildFire Key field to validate the key and the connection between the WildFire cloud and the App Control Server. If the test is not successful, use the failure message to troubleshoot the connection problem.

Note

If you need to use a proxy server for sending files from the App Control Server to the WildFire Public Cloud for analysis, you can configure this through the Licensing tab of the System Configuration page. The Carbon Black File Reputation Proxy Settings panel provides a field in which you can enter a proxy server address. This will be used for both Carbon Black File Reputation and files sent to WildFire, and the proxy will be reported when you click **Test**. See [“Activating Carbon Black File Reputation”](#) on page 763.

4. In the File Analysis panel, check the **File Analysis Enabled** checkbox.
5. If the WildFire Key test passed and you have finished entering the other required information, including checking the Integration Enabled box at the top of the page, click the **Update** button to save your changes.

WildFire Public Cloud Query Limits

Enabling WildFire public cloud analysis from App Control will increase the number of WildFire queries per day. If the number of queries sent to the WildFire cloud per day exceeds the daily limitation, consider reducing or eliminating automated file submissions or modifying the filters determining what is submitted.

The WildFire query count is incremented by the integration under the following circumstances:

- When App Control receives logs from a Palo Alto Networks appliance, the logs may reference WildFire reports. If the Enable Additional Lookups box is checked on the Palo Alto Networks Integration page, the WildFire cloud is queried for each log entry

that needs to be referenced. If your query count is exceeding the limit, you may want to disable this automatic query.

- During initial import of data from the WildFire log of the Palo Alto Networks appliance after the integration is configured, a high volume of queries may occur at one time, depending on how many days you configured for Initial Import and how many WildFire log entries exist on the firewall for that period.
- When a file is submitted from App Control to the cloud for analysis, either manually or automatically via an Event Rule, there is one WildFire query to see if the hash for that file is already known. If it is known, it will not be uploaded and so this will be the only query. If it is not known, there will be another query to submit the file and one to query for the results of the analysis.
- If an Event Rule initiates upload of a file to the WildFire cloud but the query limit for the day has already been reached, processing of that file is delayed until the next day. This allows the license count to reset. App Control initiates this delay automatically, and this state is reported as tooltip if you hover the mouse cursor over the Status field of an affected file on the Analyzed Files page.

Integrating with a WildFire Private Cloud Device

If you would rather not or cannot use a public cloud service for analysis, you can integrate App Control with a locally installed WildFire private cloud device. This also eliminates limits on the number of queries you can submit in any given time period. You can integrate multiple local WildFire appliances with a App Control Server, and analysis requests will be distributed among them.

To enable file uploads to a WildFire private cloud for analysis:

1. If you are not already on this page, in the App Control Console, click the configuration (gear) icon and choose **System Configuration**, click the **Connectors** tab and then the **Palo Alto Networks** tab, and click the **Edit** button at the bottom of the page.
2. In the File Analysis panel, click the **Local** radio button and then click the **Add New** button. The local appliance configuration fields are displayed.

The screenshot shows the 'File Analysis' configuration interface. At the top, there are three radio buttons: 'Off', 'Local' (which is selected), and 'Cloud'. Below this, there are two buttons: 'Collapse' and 'Delete'. The main configuration area contains the following fields and controls:

- Name:** A text input field.
- Appliance Enabled:** A checkbox.
- Address:** A text input field.
- API Key:** A text input field.
- Test:** A button with a play icon.

At the bottom of the panel, there is an 'Add New' button with a plus icon. Below the panel, there are three buttons: 'Edit', 'Update', and 'Cancel'.

3. In the Name field, enter the name by which you want this WildFire appliance identified in the configuration.
4. In the Address field, enter the IP address or hostname for the WildFire appliance. Enter this as an address or name, *not* as a URL with the “https” prefix.
5. In the API Key field, enter the API key for this appliance.

6. Click the **Test** button to validate the API key and the connection between the WildFire appliance and the App Control Server. If the test is not successful, use the failure message to troubleshoot the connection problem.
7. In the File Analysis panel for this device, check the **Appliance Enabled** checkbox.
8. If the license and connectivity test passed and you have finished entering the other required information, including checking the Integration Enabled box at the top of the page and configuring automatic lookups if you choose, click the **Update** button to save your changes.
9. If you want to add more private cloud appliances, click the **Add New** button and repeat the configuration for another appliance.

Enabling VMware Carbon Black Cloud Integration

If you are monitoring and managing your computers with both App Control and VMware Carbon Black Cloud, you can configure the App Control Server to connect to the Carbon Black Cloud and view information it has about activities on these jointly managed endpoints. The Carbon Black Cloud pages you can connect to from App Control depend upon whether your Carbon Black Cloud license includes Carbon Black Cloud Enterprise EDR.

Configuration of this integration requires entering search URLs for your Carbon Black Cloud server, with the search term replaced with the string “<search>”. This is done on the System Configuration Connectors tab.

The screenshot shows the 'System Configuration' interface with the 'Connectors' tab selected. Under 'Connectors', the 'Carbon Black Cloud' sub-tab is active. The configuration includes:

- Enable Carbon Black Cloud Integration:** A checkbox labeled 'Carbon Black Cloud Enabled' with an information icon.
- Do You Have Carbon Black Cloud Enterprise EDR?:** Radio buttons for 'Yes' and 'No'.
- Automatic:** A button labeled 'Populate from URL'.
- Event URL:** A text field containing 'https://server/investigate?s[c][QUERY_STRING_TYPE][0]=<search>' with a note 'Must include <search>'.
- Device URL:** A text field containing 'https://server/settings/enrollment?s[c][QUERY_STRING_TYPE][0]=<search>' with a note 'Must include <search>'.

At the bottom, there are buttons for 'Edit', 'Update', and 'Cancel'.

Table 133 describes the configuration fields and the resulting links to the Carbon Black Cloud.

Table 133: Carbon Black Cloud Integration Configuration Fields

Field	Description
Enable Carbon Black Cloud Integration	<p>Checking this box enables the integration when you save the configuration if URLs have been configured for all relevant fields.</p> <p>Uncheck the box and click Update to disable the connector. URLs are retained and are applied if you re-enable the connector.</p>
Do you have Carbon Black Cloud Enterprise EDR?	<p>The Yes radio button activates URL fields appropriate for searches of Carbon Black Cloud with Enterprise EDR. The No radio button activates URL fields appropriate for searches of Carbon Black Cloud <i>without</i> Enterprise EDR.</p>
File Event URL	<p>Link name in App Control: Carbon Black Cloud Events.</p> <p>Link location: File and File Instance Details pages</p> <p>Shows when: Carbon Black Cloud Enterprise EDR button is Yes.</p> <p>Results in Carbon Black Cloud: Shows the Investigations page with the results of a search for events involving this file, including additional data on the file provided by Enterprise EDR.</p> <p>URL: <code>https://<CarbonBlackCloudserveraddress>/EnterpriseEDR/investigate?searchWindow=ALL&query=hash:<search></code></p>
Automatic	<p>Select to populate the URL fields automatically based on your Carbon Black Cloud URL.</p>
Event URL	<p>Link name in App Control: Carbon Black Cloud Events</p> <p>Link location: File and File Instance Details pages</p> <p>Shows when: Carbon Black Cloud Enterprise EDR button is No.</p> <p>Results in Carbon Black Cloud: Shows the Investigations page with the results of a search for events involving this file.</p> <p>URL: <code>https://<CarbonBlackCloudserveraddress>/investigate?s[c][QUERY_STRING_TYPE][0]=<search></code></p>
Computer/ Device Event URL	<p>Link name in App Control: Carbon Black Cloud Events</p> <p>Link location: Computer Details page</p> <p>Shows when: Carbon Black Cloud Enterprise EDR button is No.</p> <p>Results in Carbon Black Cloud: Shows the Investigations page with the results of a search for events involving this computer, including additional file data from Enterprise EDR.</p> <p>URL: <code>https://<CarbonBlackCloudserveraddress>/EnterpriseEDR/investigate?searchWindow=ALL&query=device_name:<search></code></p>

Field	Description
Device URL	<p>Link name in App Control: Carbon Black Cloud Device</p> <p>Link location: Computer Details page</p> <p>Shows when: Always (Carbon Black Cloud Enterprise EDR can be Yes or No)</p> <p>Results in Carbon Black Cloud: Shows the All Sensors page with the results of a search for this computer.</p> <p>URL: <code>https://<CarbonBlackCloudserveraddress>/settings/enrollment?s[c][QUERY_STRING_TYPE][0]=<search></code></p>

Notes

- If you use the **Populate from URL** button on the VMware Carbon Black Cloud connector configuration screen, you only need to enter the VMware Carbon Black Cloud server address in a dialog box. URLs will be created using that address and additional strings appropriate for each link.
- If you do manually enter a URL, it must include “<search>” as shown above.

To activate VMware Carbon Black Cloud integration:

1. On the console menu, click the configuration (gear) icon and choose **System Configuration**. The System Configuration page appears.
2. Click the **Connectors** tab, and on the Connectors page click the **Carbon Black Cloud** tab.
3. Click the **Edit** button on the Carbon Black Cloud tab.
4. In the *Do you have Carbon Black Cloud Enterprise EDR?* field, click the button (**Yes** or **No**) that matches your Carbon Black Cloud environment.
Your Carbon Black Cloud Enterprise EDR choice determines the number and type of URL fields displayed. See [Table 133](#) for a description of each field and its purpose.
5. Click the **Populate from URL** button, enter the base URL address for your Carbon Black Cloud server in the dialog box, and click **OK**. This address will be used to complete URLs for each of the fields, saving you from having to type each one. You can enter this address with or without **https://** and with or without a trailing slash.
Important: The URLs constructed using your server address and shown in [Table 133](#) were accurate for the intended connections at the time of this App Control release. Because Carbon Black Cloud is constantly under development, changes to the URLs may be necessary after initial configuration. You can manually enter the URLs on this page if changes occur. Monitor the Carbon Black User Exchange for any URL changes.
6. When you have configured the URL fields, check the *Enable Carbon Black Cloud Integration* box, click the **Update** button and confirm your changes to activate the connector.

Using the Links to the Carbon Black Cloud

After configuration, there are links on the File Details, File Instance Details, and Computer Details pages, as described in [Table 133](#). Assuming the link URLs are correct, one of two things will happen when you click on a link:

- If you are already logged into the Carbon Black Cloud, the target page will display.

- If you are not already logged into the Carbon Black Cloud, the Carbon Black Cloud login page will appear and you will need to provide your credentials before going to the target page.

If the URL for the link you click on is not correct, it fails silently – there is no error message. In this case, examine the Carbon Black Cloud tab on the System Configuration / Connectors page and correct any URL errors found there. Problems could include typographic errors in the URL or choosing a Carbon Black Cloud Enterprise EDR URL when you do not have Enterprise EDR with Carbon Black Cloud.

The example below shows the links that appear on the Computer Details page when the connector is configured.

Computer Details

General

Computer Name: MYCORP\LAPTOP-7
IP Address: ff70:d1:de34:17dc:8a01
Connection Status: Connected
Health Check: Passed
Platform: Windows
Description:

Related Views

- Recent Events
- Health Check Events
- Files on this Computer
- Carbon Black Cloud Device
- Carbon Black Cloud Events

The next example shows the link that appears on the File Instance Details page when the connector is configured.

File Instance Details

Details for file on computer: MYCORP\LAPTOP-7

File Name: notepad.exe
Date Created: Feb 11 2019 04:05:03 PM
File Path: c:\\$windows.~bt\newos\windows\syswow64
Computer: MYCORP\LAPTOP-7
Platform: Windows
User Name: (none)

Related Views

- All File Instances
- File Events
- Computers with this file
- Computers without this file
- Carbon Black Cloud Events

For information about the Carbon Black Cloud and Carbon Black Cloud Enterprise EDR, see the Products section of the Carbon Black [User Exchange](#).

Enabling Check Point Integration

Enabling the App Control Connector for Check Point involves configuration steps on both the App Control Server and the Check Point cloud or appliance side. You can enable integration, configure notifications from log servers, and enable analysis of files by either the Check Point ThreatCloud Emulation Service or by a local Private Cloud appliance. The ThreatCloud Emulation Service requires a license key.

Integrating Check Point Log Servers with App Control

Notifications from multiple Check Point log servers can be integrated with a App Control Server. To perform the steps described here, you must be familiar with and have permission for advanced Check Point configuration.

To configure integration of a Check Point log server with App Control:

1. Confirm that the Check Point log server and App Control Server are able to contact each other, and that the port for the connection is open (by default, **18184**).
2. Make sure that the LEA server is running on the Check Point log server. If it is not, enable it.
3. Using the Check Point Dashboard, create an OPSEC application to be able to connect to App Control Server:
 - a. In the left lower panel click on the button for Servers and OPSEC, right click on **OPSEC Application** and choose **New OPSEC Application**.
 - b. In the Name box, enter an OPSEC Application name that clearly identifies this as an application for App Control connectivity; for example, **CbApp Control** or **Cb_App Control_Server**.
 - c. On the Host menu, choose the hostname for the Check Point log server you want to integrate with App Control.
 - d. In the Client Entities panel, check the **LEA** box.
 - e. On the LEA Permissions tab, select **Show all log fields**.
 - f. Click the **Communications** button, enter the password you will use for the SSLA certificate file, record that password for later use, and click the **Initialize** button. When the initialization is complete, click the Close button on this dialog.
 - g. Click **Close** on the OPSEC Application Properties dialog. You reopen this dialog later to copy the DN field into the App Control configuration page for Check Point.
4. In the App Control Console, click the configuration (gear) icon and choose **System Configuration** and click on the **Connectors** tab and then the **Check Point** tab.

- Click the **Edit** button at the bottom of the page.

The screenshot shows the 'System Configuration' page with the 'Connectors' tab selected. Under 'Connectors', the 'Check Point' sub-tab is active. The configuration is organized into three sections: 'General', 'Log Servers', and 'Threat Emulation'. In the 'General' section, the 'Integration Enabled' checkbox is currently unchecked. The 'Log Servers' section includes an 'Initial Import' field set to '7 Previous Day(s)' and an unchecked 'Enable Additional Lookups' checkbox. Below this is an 'Add New' button. The 'Threat Emulation' section has a 'File Analysis' radio button set to 'Off'. At the bottom of the configuration panel are three buttons: 'Edit', 'Update', and 'Cancel'.

- Check the **Integration Enabled** checkbox. This is the primary switch for the Check Point integration and must be checked for any of the integration features to be activated.
- In the Log Servers panel, go to the Initial Import field and enter the number of days of historical notification data you want to import to App Control. The default value is 7 days. This value affects only log servers from which no data has been received yet. If App Control already has data from a log server, data import will resume with the time of the last data received.
- The Enable Additional Lookups checkbox determines the level of information received from Check Point log servers. If you want to get the full malware report for each file referenced in the threat emulation notification, check this box. Note that the lookup will occur on the ThreatCloud Emulation Service or local Threat Emulation Private Cloud appliance, depending on the configuration.

Important: The Initial Import you configure will happen all at one time. If *Enable Additional Lookups* is enabled during initial import, there can be a significant performance impact. Also, if you enable additional lookups, be sure to choose an Initial Import value that will not cause the number of Check Point Threat Emulation Cloud queries to exceed your licensed limit. See [“ThreatCloud Emulation Lookup Limits”](#) on page 866 for more details.
- In the Log Servers section of the App Control Check Point page, click **Add New** to open the configuration panel for a new server. This panel allows you to add, configure, and delete integration and connectivity between Check Point log servers and the App Control Server.

General

Integration Enabled:

Log Servers

Initial Import: Previous Day(s)

Enable Additional Lookups: Enable additional lookups of notifications using Threat Emulation

Address:

Enabled:

Port:

Authentication Type: SSLCA None

SIC DN:

Password Or File: SIC One-time Password SIC Cert File

SIC One-Time Password:

Import Filters: Anti-Bot
 Anti-Virus
 Threat Emulation
 Custom Filter

10. In the Address field, provide the IP address of the log server.
11. Check the Enabled box to enable the connection between the App Control Server and this log server. You may enable or disable this integration without losing the other configuration data.
12. In the Port field, enter the port to use for connecting the App Control Server with the Check Point log server. By default, this is **18184**.
13. The Authentication field allows you to choose secure or unsecure communication between the servers. If you want the servers to communicate in clear text, click the **None** radio button and skip to step [14](#).

If you want to use secure communications, click the **SSLCA** radio button (the default) and provide the following information:

- a. **SIC DN** – This Security Internal Communication (SIC) distinguished name is required for secure communication between the App Control Server and the Check Point log server. In the Check Point Dashboard, open the Edit dialog for the OPSEC Application you created for App Control. Copy the DN field from the

- `Bit9\Integrations\CheckPoint\Bin\B9ConnectorCP.bt9`

See “[External Notifications](#)” on page 867 for a full description of the notification features, including the types of notifications pre-filtered from appearing in the console.

Custom Import Filters for Check Point

There are three standard options for determining the data that is imported from Check Point log servers to App Control: **Anti-bot**, **Anti-Virus**, and **Threat Emulation**. If you need special filters, you can choose the Custom Filter checkbox in the Import Filters field of the Connectors/Check Point tab on the System Configuration page. This disables the other filter checkboxes and opens a Custom Filter text box. All filtering is done on the log server side, reducing network traffic between the log server and the App Control Server.

The screenshot shows a configuration window for a Check Point connector. At the top left are 'Collapse' and 'Delete' buttons. The configuration fields are:

- Address:** [Empty text box]
- Enabled:**
- Port:** [18184]
- Authentication Type:** SSLCA None
- SIC DN:** [Empty text box]
- Password Or File:** SIC One-time Password SIC Cert File
- SIC One-Time Password:** [Empty text box]
- Import Filters:**
 - Anti-Bot
 - Anti-Virus
 - Threat Emulation
 - Custom Filter

Below the Import Filters section is a text box containing the following filter expression:

```
product=Anti Malware&severity>=medium|
product=New Anti Virus,Anti Virus&severity>=medium|
product=Threat Emulation&verdict=Malicious
```

At the bottom right is a 'Test' button with a play icon.

Initially, the Custom Filter box shows filters that perform the same filtering as would take place if the Anti-bot, Anti-Virus, and Threat Emulation checkboxes were checked. Examining this default custom filter can be useful in understanding the filter syntax. You add to or edit these filters or start from a blank filter window. You can resize the filter window by “dragging” the bottom right corner while holding the left mouse button down. Filters are constructed from a Check Point log attribute, an operator, and the value you want to require for the attribute. For example:

```
severity>=medium
```

requires that a notification has a severity of at least *medium* to be imported into the External Notifications table.

Filter conditions may be combined using the AND and OR operators shown in the operators table. For example:

```
product=Threat Emulation&verdict=Malicious
```

requires that the product is *Threat Emulation* and the verdict is *Malicious*.

When you finish entering the custom filter description in the box, use the **Test** button to validate the filter syntax. A successful test validates the syntax and confirms that the App Control Server has connectivity with the Check Point server, but it does not indicate that any of the Check Point data would actually match the filter.

Table 134 shows the operators for filters and examples of their use.

Note

Please note the following custom filter limitations:

- Quotes are not supported.
- Brackets are not supported to force the operator priority or to group expressions.
- Time filtering is not supported.
- Filtering by network address may not use a network mask.

Table 134: Check Point Custom Filter Operators

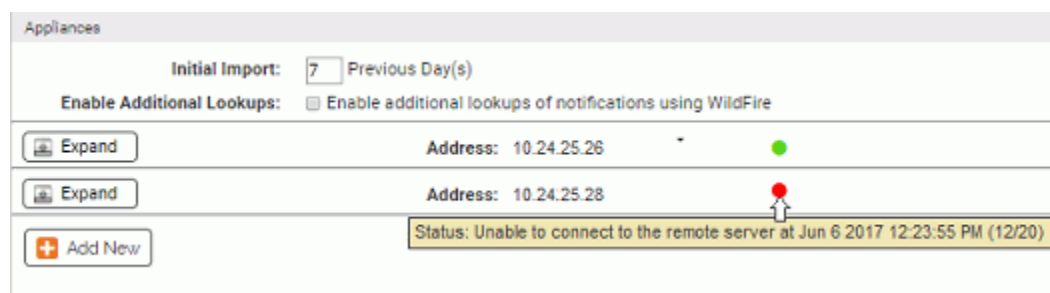
Operator	Description	Example
=	Equals	src=10.0.3.5
=	Attribute exists (if no parameter)	verdict=
!=	Does not equal	verdict!=benign
!=	Attribute does not exist (if no parameter)	verdict!=
>=	Greater than or equal	severity>=medium
<=	Less than or equal	severity<=low
%=	Contains string	emulated_on%=Windows 7
~=	Does not contain string	emulated_on~=Windows XP
= ,	Belongs to (used for a group)	product=Anti Virus,New Anti Virus
&	And (for combining filters)	severity>medium&verdict=benign Note: If AND and OR are combined, AND must always be the inner operation and OR must be the outer operation. For example: product=Anti Malware&severity>=medium product=New Anti Virus&severity>=medium
	Or (for combining filters)	severity>medium&verdict=benign verdict=malicious

Table 135 shows some of the Check Point log attributes that can be used in custom filters. Please consult Check Point LEA field guide for the complete list.

Table 135: Examples of Check Point Log Attributes Used for Filters

Check Point Log Attribute (case sensitive)	Description
src	Source hostname or IP address
dst	Destination hostname or IP address
orig	Hostname of firewall generating the log entry
severity	Severity of log entry (n/a, low, medium, high, critical)
Confidence Level	Confidence level of the event
verdict	Threat emulation verdict (benign, malicious)
src_user_name	If Identity Awareness is used on Check Point, source username
product	Software product (blade) used to generate the log entry (Anti Malware, Anti Virus, New Anti Virus, Threat Emulation)
Protection name	Malware name reported by Check Point
file_type	File type (PDF, EXE etc.)
file_size	File size (in bytes)
analyzed_on	Location where the threat emulation analysis was performed (Check Point Threat Cloud or local emulation appliance hostname)

Check Point Log Server Status



Once configured, the status of each Check Point log server integrated with App Control is displayed on the *System Configuration/Connectors/Check Point* page in the App Control Console. In the Log Servers panel, a status indicator appears next to the address of each appliance:

- A green circle indicates that there are no issues with that log server's integration
- A red circle indicates a problem, and in this case, an error message will appear with the indicator.
- A light blue circle indicates that the log server is de-activated.

For each status indicator, a tooltip provides additional information when you hover the cursor over the circle.

Log server connectivity errors are also shown as “Server Error” events on the App Control Events page and on the External Notifications page.

Modifying or Deleting a Log Server Integration

For any existing appliance integration, you can edit the configuration, for example, to enable or disable notifications from one or more Log Servers to the App Control Server. You also can delete a Log Server from the integration, or disable the entire integration with Check Point.

Address:

Enabled:

Port:

Authentication Type: SSLCA None

SIC DN:

Password Or File: SIC One-time Password SIC Cert File

SIC Cert File:

Import Filters: Anti-Bot
 Anti-Virus
 Threat Emulation
 Custom Filter

To delete or edit a Check Point log server integration:

1. On the Connectors/Check Point tab, click the **Edit** button at the bottom of the page.
2. Click the **Expand** button next to the appliance you want to edit. The configuration for that appliance is displayed.
3. If you want to delete an appliance from your integration, click the **Delete** button next to its address.
4. If you want to enable or disable log data imports, check or uncheck the appropriate checkbox.
5. If you want to change the filter for one of the log imports, check or uncheck the standard filter boxes, or edit the text in the Custom Filter box.
6. Make any other necessary changes.
7. Click **Test** to confirm that the appliance is accessible and the filter is valid.
8. Click **Update** to save your changes.

Integrating with Check Point for File Analysis

You can enable uploading of files from App Control-managed systems to the Check Point Threat Emulation Service or a local threat emulation appliance for analysis and then

receive the results back in the App Control Console. The configuration for this is in the Threat Emulation panel on the Check Point configuration page.

Note

App Control correlates only data from the portion of the Check Point report entitled "unexpected activities by time", which is at the bottom of the Check Point HTML report.

Connecting to a Threat Emulation Appliance

To enable file uploads to a Check Point threat emulation appliance:

1. On the App Control Console menu, click the configuration (gear) icon and choose **System Configuration**, click the **Connectors** tab, and then click the **Check Point** tab.
2. Click the **Edit** button.
3. In the Threat Emulation panel, click the **Local** radio button and then click **Add New**.

The screenshot shows the 'Threat Emulation' configuration window. At the top, there are radio buttons for 'File Analysis': 'Off', 'Local' (selected), and 'Cloud'. Below this are 'Collapse' and 'Delete' buttons. To the right, there are input fields for 'Name:' and 'Address:'. A list of analysis environments follows, each with a checkbox: 'Win81-64bit;Office 2013;Adobe 11:', 'Win7-64bit;Office 2013;Adobe 11:', 'Win7;Office 2003/7;Adobe 9:', 'Win7;Office 2010;Adobe 11:', 'Win7;Office 2013;Adobe 11:', and 'WinXP;Office 2003/7;Adobe 9:'. At the bottom right, there is a 'Test' button with a play icon.

4. Enter the name by which you want this Threat Emulation appliance identified in the App Control configuration.
5. Enter the IP address for the Threat Emulation appliance.
6. Check the box for each Analysis Environment in which you want files submitted from the App Control Server to be analyzed. Be sure to choose only analysis environments that are configured on the threat emulation appliance. App Control cannot determine programmatically which environments are supported on the local threat emulation appliance.
7. Click the **Test** button to validate the address of the Threat Emulation appliance and the connection between the appliance and the App Control Server. If the test is not successful, use the failure message to troubleshoot the connection problem. One possible issue is a port mismatch. While the default for connection to the ThreatCloud Emulation Service is 443, local appliances must use **18194**.

Note: The test on this page does not detect all problems with a Check Point configuration. For example, if you configure a non-existent environment, the test will not reveal that, and actions that require that environment will simply fail.

8. If the test passed and you have finished entering the other required information, click the **Update** button to save your changes.

When the analysis configuration is complete, new menu choices appear on App Control Console pages that show tables of files or file details. These **Analyze File with Check Point** commands allow uploading of files to Check Point. See [“Analysis of Suspicious Files on Endpoints”](#) on page 883 for full details on how to upload files to Check Point and how to view the results of Check Point analysis.

Connecting to the ThreatCloud Emulation Service

To enable file uploads to the Check Point cloud for analysis:

1. On the App Control Console menu, click the configuration (gear) icon and choose **System Configuration**, click the **Connectors** tab, and then click the **Check Point** tab.
2. Click the **Edit** button.
3. In the Threat Emulation panel, click the **Cloud** radio button.
4. In the Threat Emulation panel, enter your Check Point Threat Emulation Cloud Service license key.

Threat Emulation

File Analysis: Off Local Cloud

License Key:

5. Click the **Test** button next to the License Key field to validate the key and the connection between Check Point and the App Control Server. If the test is not successful, use the failure message to troubleshoot the connection problem.

Note

If you need to use a proxy server for sending files from the App Control Server to the Check Point Threat Emulation Service for analysis, you can configure this through the Licensing tab of the System Configuration page. The Carbon Black File Reputation Proxy Settings panel provides a field in which you can enter a proxy server address. This will be used for both Carbon Black File Reputation and files sent to Check Point, and the proxy will be reported when you click **Test**. See [“Activating Carbon Black File Reputation”](#) on page 763.

6. If the License Key test passed and you have finished entering the other required information, click the **Update** button to save your changes.

When the analysis configuration is complete, new menu choices appear on App Control Console pages that show tables of files or file details. These **Analyze File with Check Point** commands allow uploading of files to Check Point. See [“Analysis of Suspicious Files on Endpoints”](#) on page 883 for full details on how to upload files to Check Point and how to view the results of Check Point analysis.

ThreatCloud Emulation Lookup Limits

If you add the ThreatCloud Emulation Service analysis to your App Control-Check Point integration, be aware of the limits on the number of Check Point queries. An increase in queries will be especially noticeable if you check Enable Additional Lookups, which will request a full report from the ThreatCloud Emulation Service on a file referenced in an

external notification if the file is considered malicious and the report was not already looked up. See [“Enabling Automatic Threat Emulation Lookups”](#) for additional details.

If you find that the combined number of queries per day from your Check Point log servers appliances exceeds the lookup limits for the ThreatCloud Emulation Service, please contact Check Point for a license key extension.

Enabling Automatic Threat Emulation Lookups

When you check Enable Additional Lookups on the Check Point configuration page, you affect both the volume of lookups and the level of detail you receive in the reports that are returned.

If you are using the ThreatCloud Emulation Service, automatic lookups count against the hourly and monthly limits specified in your license key. Enabling automatic lookups when you first enable notifications could quickly exhaust the daily limit, especially if you request input of several days previous notifications.

If you are using a local Threat Emulation appliance, there is no limit to the lookups.

The content of reports received when files are submitted for analysis varies as follows:

- **Enable Additional Lookups *disabled*** (Default) – Notifications from the Threat Emulation log contain the top-level malware file, its hash, and the file size.
- **Enable Additional Lookups *enabled*** – Notifications from the Threat Emulation log cause an automatic lookup if their Verdict is malicious and if the notification was not already looked up. The results of the lookup provide file name and registry entries and also expanded file names and registry modifications, but not file hashes or sizes for the files.

Enabling Console Account Permissions

To use the Connector features, a App Control Console user must have certain permissions enabled in their user account. In addition to general administrative privileges for access to the configuration pages, the list below shows permissions specifically needed for access to Connector features.

- Tools: View file uploads (enabled by default for Administrator accounts)
- Tools: Submit files for analysis (enabled by default for Administrator accounts)

Full descriptions of these permissions and instructions on how to add them to a console user's account are described in [“User Role Permissions”](#) on page 106.

External Notifications

Enabling the App Control Connector adds an External Notifications page to the console. This page is a table of notifications from network security devices and services. Each row in the table includes key information such as file hashes and source IP addresses. If the file or computer referenced in a notification is also in App Control endpoint data, that data can be correlated with the notification.

In addition to notifications, this page will show an error message if there is a problem receiving notifications from any of the configured connected devices or services.

Notifications from Palo Alto Networks are pre-filtered to eliminate those not likely to be of interest for security analysis purposes. If a Threat Log notification has a Severity equal to “informational”, “low”, or “medium”, by default it is not included in the notifications delivered to the App Control Server. Also, WildFire Log notifications with a Category of “benign” are filtered out by default. Check Point notifications are also pre-filtered.

A daily check is done on the total number of notifications from all sources. If the daily check finds that this number is excessive, the oldest notifications in the logs are trimmed. Note, however, that the number of notifications may exceed the limit by a considerable amount before trimming is scheduled, such as when notifications are first enabled.

In addition to trimming notifications after they reach a numeric limit, the server deletes notifications past a maximum age. Initially, the numeric limit is 200,000 notifications and the age limit is six months. These may be modified in the future.

To open the External Notifications table in the console:

- Choose **Reports > External Notifications** on the console menu.

	Time	Vendor	Severity	Type	Malware Name	Total Files	Cb Computers
<input type="checkbox"/>	Jun 9 2017 01:38:01 PM	Palo Alto Networks		wildfire-result		8	1
<input type="checkbox"/>	Jun 9 2017 01:38:01 PM	Palo Alto Networks		wildfire-result		27	1
<input type="checkbox"/>	Jun 9 2017 01:38:01 PM	Cb Inspection	critical	malicious_file	np.grabber	26	1
<input type="checkbox"/>	Jun 9 2017 01:33:38 PM	Cb Inspection	low	clean_file		1	1
<input type="checkbox"/>	Jun 9 2017 09:51:57 AM	Cb Inspection	high	potential_risk_file		1	1
<input type="checkbox"/>	Jun 8 2017 01:33:18 PM	Palo Alto Networks		wildfire-result		1	1

Because of the data correlation with the App Control Server, external notifications can be prioritized immediately by their impact on systems running agents. When a malware notification is received from a connected network security source, you can determine:

- Whether the malware is present on any of your systems
- Whether it has ever executed on any of the systems
- How much it has spread (i.e., on how many computers)
- Details on the system identified as the source for this malware, including what kind of user activity there was on the system and other system activity

The External Notifications table includes several ways to drill down for additional information:

- The View Details button opens the External Notification Details page for the notification in its row. The details page includes all of the information stored in your App Control database for this notification. See “[External Notification Details](#)” on page 874 for more information. It also includes a link to open the full XML details file for the notification. See “[Showing XML Details](#)” for more information on this page.

- If there is a number greater than zero in the Total Files or New and Modified Files column, clicking on the number also opens the External Notification Details page.
- If the Malware MD5, SHA-1 or SHA-256 hash is listed in the table and identifies a file inventoried by your App Control Server, clicking on the hash opens the File Details page for that file.
- In any of the Cb Files columns, if the number of files shown is 1, clicking on the number opens the File Details page for that file. If it is 2 or greater, clicking on the number opens the External Notification Details page with the Known Files tab showing.
- In the Cb Computers column, if the number of computers shown is 1, clicking on the number opens the Computer Details page for that computer. If it is 2 or greater, clicking on the number opens the Computers table.
- If the Source or Destination Address column shows an address for a system that has the agent installed, clicking on the address opens the Computer Details page for that computer.
- The History button opens the Notification Details page with the History tab showing. The History tab includes the 20 most recent actions related to this notification.

Table 136 shows the information available in the External Notifications table. Not all of these columns appear in the table by default.

Table 136: External Notifications Table Columns

Column	Description
Vendor	Vendor whose product sent the external notification. Currently Check Point or Palo Alto Networks (other vendors might appear if you have upgraded from previous App Control versions).
Appliance	Name of the external appliance or service that provided the notification; has link to appliance or service console URL. For Check Point, if the notification came from a private threat emulator, its name is shown here.
Product	External appliance or service product name, if provided; has link to appliance console URL.
Version	External appliance, agent, or report version; has link to appliance console URL.
Time	Date and time when the malware was detected on the network.
Severity	Severity of notification. Scale varies by vendor.
Type	Type of notification (not the name). For Check Point this can be any of the configured Check Point software products (blades) that can deliver a notification. For Palo Alto Networks this can be: wildfire, spyware, virus, vulnerability, wildfire-result. Other notification types might appear if you implemented a connector for a different device or service in previous App Control versions.

Column	Description
Source IP	The IP address from which the malware originated.
Source Address	Source Address is the address from which the malware originated, from one of the following sources: <ul style="list-style-type: none"> • If the address is for a computer known to your App Control Server, the hostname listed for this source in the Bit9 database is used. In this case, the name is linked to the Computer Details page. • If the computer is unknown to your server, the server performs a reverse DNS lookup, and if the hostname can be resolved in this way, it will be used here and will persist. • If Bit9 cannot resolve the hostname, a URL is shown, as resolved by the provider • If no resolution is possible, an IP address is shown. This would be the case if malware was attempting a callback.
Source URL	URL of the computer on which the malware was originated, as resolved by the provider.
Source Username	Name of user logged into the system at the Source Address. Appears for Check Point, Microsoft and Palo Alto Networks integrations if Active Directory is integrated with the appliance or service.
Destination IP	IP address to which the malware was targeted.
Destination Address	Address to which the malware was targeted, resolved as described for Source Address.
Destination Username	Name of user logged into the system at the Destination Address. Appears for Check Point and Palo Alto Networks integrations if Active Directory is integrated with the appliance or service.
Malicious	Shows whether the notification identifies malicious files (Yes/No).
Malware Name	Malware name reported in notification (can be multiple, comma separated).
Malware MD5	Top-level MD5 hash reported in notification.
Malware SHA1	Top-level SHA1 hash reported in notification. Appears for Check Point notifications.
Malware File	Top-level filename reported in notification.
Application	Application reported in the notification.
Analysis Environment	Operating System environment used for file analysis. For Palo Alto Networks and Check Point, may also include information about key applications in the environment, such as Office.
Registry Keys	Number of registry key modifications reported in the notification.
Directories	Number of directory modifications reported in the notification.
New and Modified Files	Number of files created or modified by this malware as reported in this notification.
Total Files	Total number of unique files in this notification.

Column	Description
Received Time	Date and time this notification was received by the App Control Server.
Modified Time	Date and time when this notification was last modified (i.e., its status changed).
Cb Status	Status of the notification in App Control (Notified, Escalated, Resolved, Closed).
Cb Known Files	Number of unique files in this notification known to the App Control Server. May change based on the Correlate with App Control option on the External Notifications page.
Cb Executed Files	Number of files in this notification known to the App Control Server and executed on an endpoint. May change based on the Correlate with App Control option on the External Notifications page.
Cb Banned Files	Number of files in this notification known to the App Control Server and banned. May change based on the Correlate with App Control option on the External Notifications page.
Cb Computers	Number of Bit9-managed computers that have at least one file matching one of the reported MD5 hashes in this notification.
Cb Files On Computers	Total number of instances on agent-managed computers of files reported in this notification.
Cb Submitted	Indicates whether a file from this notification was submitted to an external device by this App Control Server for file analysis (Yes/No).

Action Menu on External Notifications Table Page

The Action menu on the External Notifications page includes the commands for changing the status of one or more notifications checked in the table and for retrieving more information about files referenced in them. Note that the notification management commands are strictly for convenience in managing them and have no impact on files in the notifications:

- **Escalate Notification** – This indicates that the notifications are of interest and you intend to investigate and/or take action related to them.
- **Resolve Notification** – This indicates that you have finished responding to these notifications.
- **Close Notification** – This indicates that you have resolved these notifications, made any necessary comments on the External Notification Details page, and no longer need to track them.
- **View Carbon Black Reputation Data** – If Carbon Black File Reputation is activated and an MD5 hash is included in the notification, opens the Carbon Black File Reputation website and displays any information available for these hashes in the checked notifications.

Saved Views on the Notifications Table Page

By default, the External Notifications page shows all notifications that have come to the App Control Server from a network security device. The pre-configured Saved Views may help focus the view on certain types of notifications:

- **Active Notifications** – Shows all notifications that do not have a status of Closed and were not a result of an analysis request from the App Control Console. This is the default view. See [“Managing Notification Status”](#) for a discussion of notification status.
- **Check Point Notifications** – Shows all notifications received from Check Point log servers.
- **File Analysis Results** – Shows all notifications from files that were submitted for analysis from the console.
- **Notifications with Files** – Shows any notifications that include at least one file hash, whether or not that file is known to the App Control Server.
- **Notifications with Known Files on Computers** – Shows any notifications that include at least one file hash for a file known to the App Control Server because it is or was on an agent-managed system.
- **Palo Alto Networks Notifications** – Shows the notifications received from Palo Alto Networks devices.

As with other console table pages, you can customize the view using the Show Filters and Show Columns buttons, and you can save any customized view you choose.

Notification Table Access from File Details Pages

On the File Details and File Instance Details pages, if there are any notifications from network security devices for the current file, an **External Notifications** choice appears on the Related Views menu. Clicking on this link opens the External Notifications table page filtered to show only notifications that include this file.

Choosing Correlation Level for External Notifications

A key feature of the App Control Connector is the correlation of security notifications received from external sources with the real-time file data available for agent-managed computers. In addition to the normal filtering and table column choices available for all App Control Console tables, the External Notifications page includes a menu that allows you to choose which files you would like correlated with notification data.

The *Correlate with App Control* panel includes the following choices:

- **New and Modified Files** – This choice correlates App Control information with all files reported in the notification, including the top-level malware and any files it writes or modifies.
- **Only Untrusted Files** – This choice correlates App Control information only for files in the notification for which the trust level reported by Carbon Black File Reputation is **5** or less.
- **Only Top Level Files** – This choice correlated App Control information only for top-level files reported in the notification, not files written or modified by these files.
- **Include Deleted Files** – This is a checkbox that is applied to any of the menu choices. If checked, files deleted from endpoints are included in those correlated with

notification data. This can be a good choice when you want to be sure to track malware that deletes itself after execution, which is very often the case.

Note

You also can change the Correlate with App Control choice on the Known Files and Files on Computer tab within an External Notification Details page. A change in any of these locations affects all notification tables.

MD5 hashes included in external notifications are used to correlate with files in the App Control Server inventory. If a notification does not include an MD5 hash but does provide a SHA-256 hash, the SHA-256 hash is used for correlation.

In a small number of cases, App Control creates a "fuzzy" hash in its file inventory for files that change their hash every time they are installed because they include date, location, or other context-specific information. These hashes are identified as "SHA-256 (Normalized)", and they may not be able to correlate with SHA-256 hashes reported in external notifications. This is relevant only if there is no MD5 hash in the notification and the file identified in the notification required a fuzzy SHA-256 hash in the App Control Server's file inventory.

For both the malware file and its parent process, file correlation begins immediately upon receipt of the notification by the App Control Server and continues as a background task for as long as is necessary to process the notification and synchronize with App Control file inventory processing. This is repeated for all unknown files until they are successfully correlated or until the notification is considered obsolete, normally 24 hours. This time period allows for correlation of a large number of new files whose notifications may arrive at the App Control Server before the server has processed the file into the Files on Computers inventory.

When files are successfully correlated, a Malicious file detected or Potential risk file detected event is generated containing the hashes of both the malware file and its parent process. If there are multiple files in the notification, the event is generated only for the top-level file. In the notification table and details, these hashes are links to the File Details pages for the respective files.

Note

When a file keeps the same name but it is modified and so its hash changes, correlation can be attempted with the new hash, but if the new hash does not appear in a notification, the correlation will fail.

Notifications from Multiple Analysis Environments

External Notifications				
Saved Views: <i>(The Current View Has Unsaved Changes - Discard)</i>		Group By:		
(none) ▼		(none) ▼ Ascending ▼		
Show Filters Show Columns Export to CSV Refresh Page				
Action ▼				
<input type="checkbox"/>	Time ▼	Vendor	Analysis Environment	Malware File
<input type="checkbox"/>	Jun 13 2017 08:25:02 AM	Cb Inspection		F_011B87
<input type="checkbox"/>	Jun 13 2017 08:21:59 AM	Palo Alto Networks	Windows XP, Adobe Reader 9.4.0, Flash 10, Office 2007	F_011B87
<input type="checkbox"/>	Jun 13 2017 08:21:59 AM	Palo Alto Networks	Windows 7 x64 SP1, Adobe Reader 11, Flash 11, Office 2010	F_011B87
<input type="checkbox"/>	Jun 13 2017 08:21:59 AM	Palo Alto Networks	PE Static Analyzer	F_011B87

Check Point and WildFire (6.0 and later) can report multiple notifications for the same file, each from a different analysis environment. The Analysis Environment field is especially useful in this case since it provides information about the test environment(s) in which the file was detonated or analyzed, allowing you to determine whether or not the file was found malicious in each environment. For notifications based on detonation of a file, the environment includes not only the base operating system but also other key software. For example, one notification might show the following Analysis Environment: *Windows 7, Adobe Reader 11, Flash 11, Office 2010*

For WildFire notifications that involved static analysis, the type of analyzer is reported in this field, for example: *DOC/CDF Analyzer*.

Note

If a file is uploaded from App Control to the WildFire cloud for analysis and WildFire reports multiple notifications for the file, the file might be considered benign in some environments and malicious in others. The External Notifications table and External Notification Details pages show the individual analysis results for each Analysis Environment. However, for a file submitted to the WildFire cloud, the Analyzed Files tab of the Requested Files page shows only the combined overall results for the file as determined by WildFire.

External Notification Details

The External Notification details page includes all of the information stored in your App Control database for one notification.

To open the External Notification Details page for one notification:

1. Choose **Reports > External Notifications** on the console menu.
2. In the row for the notification of interest, click the View Details button.

External Notification Details

FireEye Details

Type: malware-object
 Malicious: yes
 Severity: crit
 Malware Name: Notepad.Grabber
 Source Address: 192.168.5.11
 Destination Address: fe70::848a:5ef6:7b45:2f3f
 Appliance: 192.168.252.200
 Product: Web MPS
 Version: 6.2.0.74484
 Analysis Environment: Microsoft WindowsXP Professional 5.1 sp2
 Time: Jun 13 2017 05:45:15 AM

Carbon Black Details

Received Time: Jun 13 2017 06:42:20 AM
 Status: Notified
 Comment:

Total Files (2) | Known Files (1) | Files On Computers (2) | Registry Keys (8) | More Details | History

Saved Views: (none) | Add | Group By: (none) | Ascending

Show Filters | Show Columns | Export to CSV | Refresh Page

Action

Sequence	Operator	File Name	Size	MD5
1	created	notepad.exe		F2C7BB8ACC97F92E987A2D4087D021B1
2	started	123.exe	28672	F2C7BB8ACC97F92E987A2D4087D021B1

2 items | Page 1/1 | 25 rows per page

Update | Cancel

Related Views

- Carbon Black App Control Server Computer

Actions

- Escalate Notification
- Resolve Notification
- Close Notification

External Pages

- FireEye Console
- Show XML Details

The Details page includes basic information about the notification plus a series of tabs with more details at the bottom of the page. The tabs vary depending upon what type of notification it is. Most of the fields on both the main page and the tabs are described in [Table 136](#) on page 869. Information about the tabs is provided in the following sections.

Total Files Tab

This tab shows all of the files reported in this notification, including files written by other files. If the same file (i.e., a file with the same hash) is written to multiple locations, it appears multiple times in the Total Files list. The table includes the following columns:

Table 137: Total Files Tab Columns

Column	Description
Sequence	Sequence of each file's appearance when a suspected malware instance is analyzed by the network security device. The first file in the sequence is the top-level process.

Column	Description
Operation	The operation performed on a file (start, create, close, etc.)
File Name	File name reported by the network security device. For Check Point, only reported for the first file.
Size	File size reported by the network security device. For Check Point, only reported for the first file.
MD5	MD5 hash of the file. For Check Point, only reported for the first file.
File Path	File path of the file name reported in the notification.
Parent File Name	File name of the parent process of this file.
Parent File Path	File path for the parent process of this file.
SHA1	SHA1 hash of the file (if reported). For Check Point, only reported for the first file.
SHA-256	SHA-256 hash of the file (if reported). Only shown for Palo Alto Networks notifications.
Known File	Is this file known to the App Control Server (Yes/No).

The Operation column provides important information about what was done for each file included in the notification. You can sort or filter on this field to determine what was done to a file. The notification might report that one file was *created* and another *overwritten* – files having these two operations are included in the New and Modified Files list. A file also might be *opened* or *terminated*.

If a file is known to your App Control Server, its listing on the Total Files tab includes a View Details button, which opens the File Details page for the file.

The Action menu for this tab includes the following commands for selected files:

- **Ban Globally** – Bans file(s) for all policies; requires no further configuration
- **Ban By Policy** – Opens a dialog box for creation of policy-specific and report-only bans
- **Remove Approval Or Ban** – Removes any active bans/approvals immediately.
- **Find By Name** – Redirects to Find files page filtered by selected file names
- **Find By Size** – Redirects to Find files page filtered to show results of a search for files matching the sizes of the selected files as reported in the external notification
- **Find By Hash** – Redirects to Find files page filtered to show results of a search by hash for the selected files as reported in the external notification
- **View Carbon Black Reputation Data** – Redirects to Carbon Black File Reputation (if activated) for report on this file by hash

Known Files Tab

This tab shows all files from this notification that are known to the App Control Server. The table includes (either by default or customization) all fields from the File Catalog. You also can add other fields that provide information about the file from the network security device, as shown on the Total Files tab. The Action menu has the same options as the

Total Files tab menu, but uses file information from the App Control inventory rather than the notification where available.

You can modify Correlation Details options on this page to customize the App Control information correlated with the notification. Your choices here affect all pages that display correlation options.

Files On Computers Tab

This tab shows all instances of the files in this notification in the App Control Server file inventory. The can include (either by default or customization) all fields from the console Files On Computers page. You also can add External File Name and External Size columns. The Action menu has the same options as the Total Files tab menu.

You can modify Correlation Details options on this page to customize the App Control data correlated with the notification (this affects all pages that display correlation options).

Registry Keys

This tab shows all relevant registry value modifications reported in the External Notification. The table for this tab includes the following columns:

Table 138: Registry Keys Tab Columns

Column	Description
Sequence	Sequence of registry access attempts when a suspected malware instance is analyzed by the network security device.
Process	Process reported by the network security device.
Process MD5	MD5 hash of the process
Process Path	Path location of the process reported by the network security device
Key	Registry key reported by the network security device (truncated to the right when displayed)
Name	Registry field name reported by the network security device
Value	Registry field value reported by the network security device
Operation	Operation on a registry key (setval, added, etc.)

If a process that attempted access to the registry key is known to the App Control Server, its listing here includes a View Details button, which opens the File Details page for this process.

The Action menu for this tab includes the following commands for selected files:

- **Ban Process Globally** – Bans process file(s) for all policies; requires no further configuration
- **Ban Process By Policy** – Opens a dialog box for creation of policy-specific and report-only bans
- **Remove Process Approval Or Ban** – Removes any active bans/approvals immediately.

- **Create Registry Rule** – Opens an Add Registry Rule page with pre-populated values to create a rule to ban this process from accessing the registry keys reported in the notification. See [“Registry Rules”](#) for more details.

More Details Tab

This tab shows additional details from the current external notification – the information included on this tab varies according to the type of the notification. The following table shows the possible fields:

Table 139: More Details Tab Fields

Field	Description
Malware type	Type of malware as reported in external notification; may be the same as Type in the External Notifications table or a more specific type, such as Backdoor, HackTool, Trojan, etc.
Anomaly	Anomaly
Application	Application targeted
HTTP Header	HTTP header(s) reported by an external notification for a web infection
Show XML Details	Opens a new browser tab with full XML notification from the external network security device. This alert is read from a file stored on the server (inside a “store” subfolder). Note: Very large XML files may cause browser performance and navigation issues when you use this link to open them. One alternative is to right-click on the link and Save Target/Link As to a location where you can open the file with a different viewer.

History Tab

The History tab provides an audit trail for external notification workflow. This includes each change of status and any comments associated with the change. In addition to clicking this tab when you are already on the Notification Details page, you display the history by clicking the History button in the Action column of the row for a notification on External Notifications table.

Showing Related Notifications

If there are any notifications related to the one currently shown on the External Notification Details page, the Related Views menu includes a **Show Related Notifications** command. A related notification is one with the same MD5 hash as the currently shown notification.

When you click on this command, the External Notifications table opens, filtered to show the related notifications, including the one from which the link was clicked.

Showing XML Details

External notifications are reported in XML format, and contain information about analyzed malware behavior. The App Control Server parses these XML notifications for efficient storage of key information in its database. In addition, the entire content of each XML notification is stored in a separate *store* folder for each network security device vendor in the Bit9 installation directory on the App Control Server (*Bit9\Integrations\PAN\store* or *Bit9\Integrations\CheckPoint\store*). Other folders containing XML may appear under *Bit9\Integrations* for custom or previously supported connectors.

Note

- Opening very large XML details files may cause browser performance and navigation issues. One alternative is to right-click the link and Save Target As or Save Link As to a location where you can open the file with a different viewer.
- If a notification from Palo Alto Networks includes reports for multiple “Analysis Environment” types, using Show XML shows only the XML details for the Analysis Environment of the current notification.

To access the full XML details for an External Notification:

- On the External Notification Details page for the notification, click **Show XML Details** in the External Pages menu. The full details appear in a separate browser window.

External Console Access

On the Notification Details page for most connectors, you can click on a command in the External Pages menu to open the console for the appliance that provided the notification. The console opens in a new browser window. If the user on the console is not already authenticated with credentials for the external appliance, the browser is redirected to a login page.

Managing Notification Status

In the console, both the External Notifications table and the External Notification Details page show a *status* field for each notification. Notification Status is strictly a means for tracking the progress of your response to a notification and does not communicate status changes back to the notification source. There is no mandatory flow of notification status, but the following might be a useful template for status work flow.

To manage the status of a notification:

1. On the console menu, choose **Reports > External Notifications** and click the View Details button next to the notification you want to review. The External Notification Details page opens.
2. On the External Notification Details page, if you intend to examine and/or take action on this notification, choose **Escalate Notification** in the Actions menu. The status changes to Escalated.

3. Research the notification using the information on the External Notification Details page, the File Details page, the Event pages, the network security device analysis of a file, or any other means appropriate for the notification. Provide any comments related to the escalation in the Comments field.
4. Take whatever action you choose to take on the files in the notification, for example, banning files or creating custom or registry rules.
Note: Bans or other rule changes do not affect the Status field of the request itself. You must change status manually.
5. Provide any comments related to the resolution in the Comments field.
6. Once you have taken action, or if you determine that no action is necessary, choose **Resolve Notification** in the External Notification Details Action menu. The status changes to Resolved.
7. When you are finished with this notification, make any final comments in the Comments field and then choose **Close Notification** in the Actions menu. The status changes to Closed and the view returns to the External Notifications table. Closing a notification removes it from the **Active Notifications** view, but it is visible if you choose a Saved View of **(none)**.

The steps above describe Status being changed from the Actions menu on the External Notification Details page. You also can change status using the Status dropdown menu on that same page, and from the Action menu on the External Notification page table.

Banning Externally Reported Malware

The App Control Server can ban files or processes reported as part of a malware notification by external network security devices. This can be done in several ways:

- **Manual file bans** of files reported in external notifications
- **Registry Rules** that ban certain processes that attempt access to registry keys, as reported in external notifications
- **Custom Rules** that ban activity in a directory reported in external notifications
- **Event Rules** that automatically ban files (or create report-only bans) when certain file-related events occur, in this case, due to external notifications

Registry, Custom, and Event rules can also be configured to *report* the actions they describe rather than banning them.

Notes

Bans of MSI files should not rely on hashes reported by a third-party source. In addition, they should not use MD5 or SHA-1 hashes from any source. See [“Approvals and Bans of MSI Files by Hash”](#) on page 303 for details.

Manually Banning Files

You manually ban files reported in external notifications much the same way you would any inventoried file. However, you can apply bans directly from the External Notification Details page Action menu, so you can ban malware identified in an external notification, whether or not it has appeared yet on a App Control-managed endpoint.

To manually ban files reported as malware in an external notification:

1. Click the View Details button next to the notification whose files you want to ban.
2. On any of the Files tabs on the External Notification Details page, check the box to the left of each file you want to ban.
3. On the Action menu, choose the ban type you want to apply to the checked files:
 - a. Choose **Ban Globally** to ban the file for all computers. This creates the ban without requiring any further interaction.
 - b. Choose **Ban by Policy** to customize the ban. This opens the Add File Rule page with information partially filled in. On this page, you can choose a fully functional ban or a Report Only ban, and you can choose specific policies to which the ban will apply. Report Only bans are useful if you want to monitor what an active ban *would* do before fully enabling it. When you have configured the ban, click **Save**.

The Action menu on the Files tabs on the External Notification Details page include the following choices for finding a file of interest:

- **Find by Name**
- **Find by Size**
- **Find by Hash**

The Files tab of the Software Rules page (**Rules > Software Rules** on the console menu) shows bans you have created. Bans manually created from an external notification are named with a prefix of "External_" followed by the file name.

Some External Notification pages allow you to ban the *process* that attempted to perform an action on an object on your systems, such as modifying a registry key or writing to a directory. You can ban those processes using the same procedure described above, except that the commands will say *Ban Process* instead of just Ban.

Special Rules for Reporting or Banning Malware

For certain notifications, standard file bans may not provide the best remediation. The App Control Connector offers several other rules to control actions that are identified as suspicious. As with bans, these rules can be created from the External Notification Details page with some of the rule data pre-populated.

Registry Rules

If a notification includes suspicious registry entries or activity, its External Notification Details page includes a Registry Keys tab. This tab provides information about the keys that might be compromised. You can select one or more of the reported keys and:

- Ban the process that tried to access the key
- Remove previously created process bans or approvals
- Create a Registry Rule to control access to the key

Bans created in this context are similar to those created on any of the Files tabs. The Registry Rule command provides different options.

To create a Registry Rule from a Notification Details page:

1. In the Notification Details page of interest, click on the **Registry Keys** tab.
2. Check the boxes next to the registry keys for which you want to create a rule.

3. On the Action menu, choose **Create Registry Rule**. The Add Registry Rule page appears, with rule name and settings pre-populated with details from the notification.
4. By default, a rule created in this way blocks writes to the named registry keys by the processes identified in the notification, and does this for all users and all policies. You can modify these settings before you save the rule. Among the options on the Write Action menu, you can choose **Report**, which means that activity at this key is reported but not blocked. If you are unsure of how best to configure a rule, see [“Creating Registry Rules”](#) on page 452. You can **Cancel** the rule without saving it if you would like to investigate rules parameters first.
Important: Rule menus have options that *Allow* activity at the named locations and even *Promote* processes to have more privileges than they previously did. If you alter the pre-populated values, be careful of the choices you make on these menus.
5. Modify the rule as you choose, and then click the **Save** button. The new rule is created and appears on the Registry tab of the Software Rules page in the console.

Custom Rules for Directory Control

Notifications that include suspicious pathname entries have a Directories tab on their External Notification Details page, providing information about the directories that might be compromised. On this tab, you can select one or more keys and:

- Ban the process that tried to access the directory
- Remove previously created process bans or approvals
- Create a Custom Rule to control access to this location

Process bans created in this context are similar to file bans created on any of the Files tabs. The Custom Rule command provides different options.

To create a Custom Rule from a Notification Details page:

1. In the Notification Details page of interest, click on the **Directories** tab.
2. Check the boxes next to the Directories for which you want to create a rule.
3. On the Action menu, choose **Create Custom Rule**. The Add Custom Rule page appears, with rule name and settings already filled in with details from the notification.
4. By default, a rule created in this way blocks writes to the named directories by the processes identified in the notification, and does this for all users and all policies. You can modify these settings before you save the rule. Among the options on the Execute Action menu, you can choose **Report**, which means that activity at this location is reported but not blocked. If you are unsure of how best to configure a rule, see [“Creating a Custom Rule”](#) on page 396. You can **Cancel** the rule without saving it if you would like to investigate rules parameters first.
Important: Some options on the rule menus that *Allow* activity at the named locations and even *Promote* processes to have more privileges than they previously did. If you alter the pre-populated values, be careful of your choices on these menus.
5. Modify the rule as you choose, and then click the **Save** button. The new rule is created and appears on the Custom tab of the Software Rules page in the console.

Analysis of Suspicious Files on Endpoints

If you have enabled integration and file analysis with an external device or service, you can submit files from the App Control Server file inventory to the connected source for analysis. With analysis enabled, the console adds **Analyze with...** commands to menus in several locations that allow you to submit files to appliances or services from Palo Alto Networks, Check Point, or CB Inspection. For Check Point, these commands have Windows-version-specific submenus so that you can choose the environment in which you want the file analyzed. The locations for these commands are:

- File Catalog, Files on Computers and Find Files Results pages Action menus (for one or more files)
- File Details and File Instance Details Advanced menus (for one file)
- Events page Action menu (for one or more files)
- Other table pages that list files

Notes

- **Unavailable files:** A file in the inventory might be unavailable, either temporarily, because it is inaccessible on the network, or permanently, because it was deleted or was a transient file. If you attempt to send such a file for analysis, when it is not found, App Control will attempt to locate another instance of the same file and send that file for analysis. If no other instance exists, the analysis request will produce an error.
- **Non-ANSI characters:** If the name or path of a file uploaded for analysis contains non-ANSI-convertible characters, the zip file used to upload it gives it a temporary name and/or path. See [“File and Path Information for Uploaded Files”](#) on page 903 for information on how renaming is done.

Platform Note: File analysis via the Connector currently is supported for files from Windows agents.

To submit files to an external service for analysis:

1. In a table that lists files, check the boxes next to files you want to submit.
2. On the Action menu choose from the available **Analyze with** commands – the available commands depend upon the appliances you have enabled for the connector:
 - a. If you have enabled the Palo Alto Networks-App Control for file analysis, you can choose **Analyze with Palo Alto Networks WildFire**.
 - b. If you have enabled the Check Point-App Control integration for file analysis, you can choose the **Analyze with Check Point** submenu and under it, the analysis environment in which you want the file analyzed, which includes the operating system and other common tools such as Microsoft Office and Adobe Acrobat (for example **win7;Office 2010;Adobe 9**).
 - c. If you have enabled other custom connectors or are still using connectors from previous versions of App Control, other **Analyze with...** choices may appear.
 - d. If you have enabled CB Inspection integration for file analysis, you can choose **Analyze with CB Inspection**.

A message will appear indicating that the files have been scheduled for upload to the analysis source you chose.

- Alternatively, you can go to a File Details or File Instance Details page for a single file and choose an **Analyze with** command on the Advanced menu.

From these pages, if a file has already been submitted to the same analysis provider, a warning is shown, but the file will be uploaded again if you click **OK** on the warning.

- To monitor the progress of the analysis, choose **Tools > Requested Files** and click on the **Analyzed Files** tab to see the table of files submitted.

Monitoring Files Submitted for Analysis

In the console, the Analyzed Files tab of the Requested Files page shows the status and (if complete) analysis results for all files submitted to external services for analysis. The default view for this page shows all files sorted by request date, but there also are Saved Views available that can provide a more targeted list of files:

- Analysis in Progress
- Completed Analysis
- Analysis Errors
- Files Submitted to Check Point
- Files Submitted to WildFire

Requested Files: Analyzed Files

Uploaded Files Analyzed Files Diagnostic Files

Saved Views: (none) Group By: (none) Ascending

Show Filters | Show Columns | Export to CSV | Refresh Page

Action

<input type="checkbox"/>	Request Date	Status	Target	Analysis Result	Computer	File Name
<input type="checkbox"/>	Jun 11 2017 10:46:00 AM	Acquiring File	Palo Alto Networks WildFire		MYCORPDT-4	gdump.exe
<input type="checkbox"/>	Jun 11 2017 10:43:27 AM	Analyzed	Palo Alto Networks WildFire	Malicious	MYCORPDT-1	icar.exe
<input type="checkbox"/>	Jun 11 2017 10:30:59 AM	Canceled	Palo Alto Networks WildFire		MYCORPDT-7	abcd.exe
<input type="checkbox"/>	Jun 11 2017 10:11:13 AM	Canceled	FireEye:win7		MYCORPDT-6	uload.exe

The table can show the following columns (not all are shown by default):

- Request Date** – When the request for file analysis was submitted for this file.
- Requester** – The user who requested the upload.
- Upload %** – The percent complete of the upload (not the analysis).
- Status** – This indicates where in the analysis process this file is. See [“Analysis Status”](#) for a description of status values.
- Analysis Results** – When the analysis is completed, this field indicates the result of the analysis (Clean, Potential Risk or Malicious).
- Computer** – The computer from which the file was uploaded.
- File Name** – The name of the file in the location from which it was uploaded.

- **File Size** – The size of the file as it appears (or appeared) on agent-managed computers.
- **MD5** – The MD5 hash of the file.
- **Date Modified** – The last time the entry for this file was changed.
- **Error** – Any error associated with the upload or submission for analysis of the file.
- **File Path** – The directory where the file resided on the source computer at the time the file was uploaded - it is not necessarily the current location of the file.
- **Last Modified By** – Who last modified the Analyzed Files entry for this file by taking a related action.
- **Prevalence** – The prevalence of this file on agent-managed computers.
- **Provider** – Palo Alto Networks, Check Point or possibly a custom or legacy connector
- **SHA-256** – The SHA-256 hash of this file.
- **Source** – The source of this analysis request. Can be "Manual" or "Event rule".
- **Source Name** – If the source was "Event rule", the name of the rule.
- **Target** – The target for the file analysis. This will be **CB Inspection, Palo Alto Networks WildFire, Check Point:**<Target Environment> or a target for a custom connector or one from an older App Control release. For Check Point analysis done on a local appliance, this field also shows the appliance name. For example: **Check Point:win7;Office2010;Adobe9:Appliance1**

Files from the App Control Agent that are targeted for analysis are not stored on the App Control Server and cannot be downloaded to the server or deleted from this table.

Analysis Status

On the Analyzed Files tab, the Status column provides feedback on the progress of a file analysis. Hovering over the Status value in the table provides additional information. The possible values are:

- **Acquiring File** – For files that must be uploaded from an endpoint before being sent to the device for analysis, this indicates that the upload has not been completed.
- **Error** – The upload or analysis failed (e.g., because the file name or path did not exist). Moving the mouse cursor over this field shows a tooltip with details of the error.
- **Canceled** – The upload was canceled by a console user.
- **Analyzing** – The file has been moved to a device for analysis.
- **Analyzed** – The App Control Server has received an XML report from the device. Once this happens, the Status value for the file becomes a link leading to Notification Details.
- **Analyzed* (1,2...)** – When Analyzed is followed by a series of numbers in parentheses, this indicates that there were multiple file analysis results from WildFire. Each result is from a different "Analysis Environment". Hovering the mouse cursor over a number shows the Analysis Environment it represents.

Status	Target	Analysis Result
Analyzed*(2)	Palo Alto Networks Wildfire	Clean
Analyzed*(1)	Windows 7;Adobe Reader 11;Flash 11;Office 2013	
Analyzed	Palo Alto Networks Wildfire	Malicious

Clicking on a number shows the specific Notification Details for that Analysis Environment. See [“Notifications from Multiple Analysis Environments”](#) on page 874 for more on the possible values.

The Analysis Results for a file that has multiple results reports the top-level analysis value provided by WildFire.

Note

If there are analysis results for a file, they appear in an External Analysis Results panel on the File Details and File Instance Details pages for that file.

Actions on the Analyzed Files tab

The Action menu on the Analyze tab provides options for you to retry an analysis request with the same or different analysis provider. It includes the following options:

- **Cancel Analysis** – Cancels checked analysis entries. If one or more checked entries cannot be canceled, this will have no effect on those files.
- **Retry Analysis** – Retries checked analysis entries. This has no effect on entries that cannot be retried (for example, because analysis is already pending on this file).
- **View Carbon Black Reputation Data** – Get information (if available) from Carbon Black File Reputation for the checked files.
- **Analyze with ...** – Options appear for each available analysis provider (CB Inspection, Check Point, and Palo Alto Networks WildFire). For Check Point, there are options to submit the file to the appropriate operating system.

When one of these actions is chosen, the submission for analysis will use an existing uploaded file if available. If not, it will first upload file, and then submit it.

Note

In addition to the Analyzed Files tab, the Requested Files page has two other tabs not described in this appendix:

- **Uploaded Files** – Shows inventoried files uploaded from App Control-managed endpoints to the App Control Server.
- **Diagnostic Files** – Shows diagnostic files uploaded to the App Control Server.

See [Appendix E, “Uploading Files from Agents,”](#) for a full description of general and diagnostic file uploads.

Logging of Connector-related Events

The Events page provides access to all recorded events related to App Control activities in your environment, including files blocked, unapproved files executed, system management processes and actions by console users. The App Control Server updates its event data in near-real-time for connected computers, with minor variations due to event volume. See [“Event Reports”](#) on page 590 for more details.

You can optionally choose to direct the App Control Syslog event output for post-processing on another system. See [“Event Management Options”](#) on page 732 for more details.

When the App Control Connector for Network Security Devices is enabled, connector-related events appear in the App Control event log. There are several key additions or changes to App Control events due to the integration with network security devices:

- **External Notification** – This event subtype (*subtype* is the most specific identifier for an event) is under the Discovery type. It is generated for external notifications (currently from Check Point or Palo Alto Networks) received by the App Control Server. However, it is not generated for an external notification that is received as a result of a file submission if a File Analysis Complete is also generated.
- **Connector Actions in Other Events** – Other events that can report connector-related activity are shown in [Table 140](#). Most of these event subtypes are also used for other purposes – descriptions that could appear for the subtype but are not related to network security device activity are not shown here. See the separate *VMware Carbon Black App Control Events Guide* for a complete description of all event types and subtypes in App Control and how to enable Syslog event output.

Table 140: Connector-Related Events in the App Control Event Log

Event Type	Event Subtype	External Notification-Related Description and Samples
Discovery	Malicious file detected	Unknown file '\$filename\$' [\$param1\$] was identified by \$param3\$ as malicious. or File '\$filename\$' [\$param1\$] was identified by \$param3\$ as malicious.
Discovery	Potential risk file detected	Unknown file '\$filename\$' [\$param1\$] from \$param3\$ was identified by \$param3\$ as potential risk. or File '\$filename\$' [\$param1\$] from \$param3\$ was identified by \$param3\$ as potential risk.
Discovery	External Notification	\$Provider\$ reported \$malware type\$ with name \$malware name\$ for file '\$filename\$' from \$src_ip\$ to \$target_ip\$
Computer Management	File Upload Requested	User '\$username\$' requested upload of file [\$hash\$] from computer '\$computer\$'. or User '\$username\$' requested upload of file '\$param1\$' from computer '\$computer\$'. or Upload of file [\$hash\$] from computer '\$computer\$' was requested by event rule '\$ruleName\$'. Note: Reported uploads could be unrelated to External Notifications.

Event Type	Event Subtype	External Notification-Related Description and Samples
Computer Management	File Upload Completed	Upload of file [hash] from computer 'computer' completed. or Upload of file 'param1' from computer 'computer' completed.
Computer Management	File Upload Canceled	User 'username' canceled upload of file [hash] from computer 'computer'. or User 'username' canceled upload of file 'param1' from computer 'computer'.
Computer Management	File Upload Error	Upload of file [hash] from computer 'computer' failed because of error 'param2'. or Upload of file 'param1' from computer 'computer' failed because of error 'param2'.
Computer Management	File Upload Deleted	User 'username' deleted uploaded file [hash]. or User 'username' deleted uploaded file 'param1'.
General Management	Event rule created	Event rule 'param1' has been created by 'userName'.
General Management	Event rule modified	Event rule 'param1' has been modified by 'userName'.
General Management	Event rule deleted	Event rule 'param1' has been deleted by 'userName'.
Server Management	File analysis requested	User 'username' requested analysis of file [hash] with 'param1'. or Analysis of file [hash] with 'param1' was requested by event rule 'ruleName'.
Server Management	File analysis completed	File 'filename' [hash] was successfully analyzed with 'param1'. Nothing suspicious was found. or File 'filename' [hash] was successfully analyzed with 'param1'. It was reported as malicious.
Server Management	File analysis canceled	User 'username' canceled analysis of file 'filename' [hash] with 'param1'.
Server Management	File analysis error	Analysis of file 'filename' [hash] with 'param1' failed because of error 'param2'.

Event Type	Event Subtype	External Notification-Related Description and Samples
Server Management	Server error	\$param1\$ Note: This is not specific to connectors but may report connector-related errors, such as failure to connect to or authenticate with a device.
Server Management	Connector restart	Connector started, build information: \$param1\$.
Server Management	Connector shutdown	Connector shutdown cleanly.

Additional Log Information

In addition to the App Control event log, you may be interested in information available in the log files for the connector integrations. This information is located in the following locations under the App Control installation folders:

- **For Check Point** – `\Bit9\Integrations\CheckPoint\B9ConnectorCP.bt9`
- **For Palo Alto Networks** – `\Bit9\Parity Server\Reporter\ParityReporter.log`

You may see other logs for custom connectors or deprecated connectors from previous App Control releases.

Appendix D

Diagnostic Files

Sections

Topic	Page
Overview	891
Uploading Agent Diagnostic Files	891
Viewing Diagnostic Files	892

Overview

The App Control Console includes a page that displays certain diagnostic files for the App Control Server and its agents. These files can be useful when you are investigating issues in your App Control environment with the assistance of Carbon Black Support.

Diagnostic files appear on the Diagnostic Files tab of the Requested Files page, and include:

- Server Installation Logs and dump files
- Agent diagnostic files requested by console users

Server installation log and dump files appear automatically on the tab when server activity causes them to be created. Agent diagnostic files must be requested through the Computers or Computer Details page. Once uploaded to the server, these files may be downloaded to any computer running the console.

Unlike the ability to select and upload any file from an agent, access to diagnostic files is available without a special license or permissions.

Note

This appendix describes *diagnostic file* uploads only. For information about uploading other files from an agent, see [Appendix E, "Uploading Files from Agents."](#) This capability requires a separate license.

Uploading Agent Diagnostic Files

Agent diagnostic file uploads can be initiated from the Computers page or the Computer Details page. On the Computers page, you can upload files from one or more computers.

To initiate a diagnostic file upload from one agent:

1. On the console menu, choose **Assets > Computers**. The Computers page appears.
2. Find the computer whose statistics or diagnostic information you want to upload and open its details page.
3. On the Computer Details page, choose **Other Actions** in the Advanced menu, and on the Other Actions menu, choose **Upload diagnostic files**.

Unless a problem is encountered, a message on the Computer Details page indicates that the upload of the file you chose has been scheduled. You can check the Diagnostic Files tab to see whether a new zip file for this agent is available yet.

To initiate diagnostic file uploads for one or more agents:

1. On the console menu, choose **Assets > Computers**. The Computers page appears.
2. Check the box next to each computer for which you want to upload diagnostics, then choose **Upload diagnostic files** on the Action menu. A confirmation dialog appears.
3. Choose **OK** on the confirmation dialog to begin the upload. A status message indicates whether the request was successful and indicates how many computers will be sending diagnostic files to the server.

Canceling or Retrying an Upload

If an upload has not been completed, you can cancel it. This might be a choice you want to make if you inadvertently chose more computers than you actually wanted when you initiated the upload, or if the file size shown in the table is excessively large.

To cancel diagnostic file uploads:

1. On the console menu, choose **Tools > Requested Files**.
2. Click on the **Diagnostic Files** tab. The table shows diagnostic files that have been uploaded, are in the process of uploading, or were requested but not uploaded.
3. Check the box next to each file whose upload you want to cancel and choose **Cancel Uploads** from the Action menu.
4. Choose **OK** on the confirmation dialog.

If an upload has failed or been canceled, you can Retry it by checking its box on the Diagnostic Files page and choosing **Retry Uploads** on the Action menu.

Viewing Diagnostic Files

Diagnostic files are listed in the table that appears on the Diagnostic Files tab of the Requested Files page.

To view diagnostic files:

1. On the console menu, choose **Tools > Requested Files**.
2. Click on the **Diagnostic Files** tab. Any diagnostic files uploaded to the server are shown in the table.

Requested Files: Diagnostic Files

Uploaded Files | Analyzed Files | **Diagnostic Files**

Saved Views: (none) Add Group By: (none) Ascending

Show Filters | Show Columns | Export to CSV | Refresh Page

Action

<input type="checkbox"/>	File Type	Request Date	Priority	Requester	Status	Computer	File Name
<input type="checkbox"/>	Server Diagnostics	Jul 13 2020 06:40:00 AM	Medium		Uploaded	System	HostPackageInstall-2020-07-13- 63721
<input type="checkbox"/>	Server Diagnostics	Jul 13 2020 06:38:00 AM	Medium		Uploaded	System	RulesInstall-2020-07-13- 63534.log
<input type="checkbox"/>	Server Diagnostics	Jul 13 2020 06:35:00 AM	Medium		Uploaded	System	backup-ServerInstall-2020713-063209.

3 items Page 1/1

When you request a diagnostic file upload from an agent, a zip file is uploaded to the server with diagnostic and log files relevant to the agent. For example, on Windows systems, the zip file includes the agent Logs folder (ProgramData\Bit9\Parity Agent\Logs) and selected log files from the Windows folder. The exact files included in the zip file vary by operating system platform.


Uploaded diagnostic files are named in the following format:

`<computername>-diagnostics-<date>-<time>.zip`

Server diagnostics files may be **.log**, **.dmp**, or other formats.

Table 141 shows the columns available for the Diagnostic Files page, some of which appear by default and some of which you must add.

Table 141: Diagnostic Files Table Columns

Column	Description
Actions	<p>The Action column includes a checkbox for choosing files on which Action menu commands will act and buttons for taking action on individual files. The Action menu on this page includes the following commands:</p> <ul style="list-style-type: none"> • Cancel Uploads – Cancel the upload of checked files (if the upload has not been completed). • Retry Uploads – Retry the upload of checked files. • Delete Uploads – Delete the table rows for checked files, and, for successful uploads, delete the files from the server. <p> Download the file (if it was successfully uploaded) from the App Control Server to the computer on which the console is being viewed.</p>
Priority	Priority in which pending files are uploaded to the server. For diagnostic files, the priority is always Medium .
Request Date	When the file upload was requested.
Requester	The console user that requested the upload, or blank if the file is a server log.
Status	<p>The status of the file upload. The possible values are:</p> <ul style="list-style-type: none"> • Uploaded – The upload completed successfully and the file is available on the server. • Uploading – The upload is in progress but not yet complete; a partial file has been received by the server. This status is likely to appear only for very large files. • Initiated – The upload task has been received by the agent where the file is located. • Queued – The upload task has not yet been sent to the agent. • Error – The upload failed. Hovering the cursor over this status displays the error message. Errors include: No file with hash, The system cannot find the path specified, The system cannot find the file specified. • Canceled – The upload was cancelled by a console user.
Computer	The name of the computer from which the file was uploaded.
File Name	The name of the uploaded file.
File Size	The size (in bytes) of the file.
Upload %	The percent of the upload that is finished. Completed uploads show 100%. Failed uploads and uploads not yet started show 0%.
Upload Date	When the file was uploaded to the server.

Column	Description
Upload Directory	The directory on the App Control Server to which the file was uploaded. Value is "(default)" for manual uploads, which use the directory configured in the System Configuration Advanced Options tab. If the upload is due to an event rule, the actual path is shown.
Error	A description of the error that prevented the file from uploading. Not shown by default.
File Path	The location on the agent computer from which the file was uploaded. Not shown by default.
Prevalence	The number of agent-managed computers reporting to your server on which this file is present.
MD5	The MD5 hash of the file.
SHA256	The SHA-256 hash of the file.
Source	Source of the request for upload. Either "Event rule" or "Manual".
Source Name	If the request was due to an event rule, the name of the rule. If the request was manual, this field is empty.

Deleting Uploaded Diagnostic Files

When you no longer need a diagnostic file, you can delete it from the server by checking the box next to its row on the Diagnostic Files page and choosing Delete Uploads from the Action menu.

Appendix E

Uploading Files from Agents

Sections

Topic	Page
Overview	896
Enabling Access to File Upload Features	896
Scheduling Uploads	897
Viewing the Uploads Table	900
Downloading Uploaded Files	903
Deleting Uploaded Files	904

Overview

In all active modes, App Control provides the ability to monitor the propagation of software and generate audit trails of activity. In some cases, information you see during monitoring might lead to a need to access the actual file involved in certain activities. With the optional Upload Files feature, you can upload a copy of any file to the App Control Server from a computer running App Control Agent 7.0.0 or later.

Access to the Upload Files feature requires application of a special license key, either for File Uploads alone or as part of the App Control Connector license. See [“Managing App Control Licenses”](#) on page 759 for instructions on applying licenses.

Notes

The ability to send a file to third-party devices or services for analysis uses the File Upload feature. However, uploads initiated by a request for analysis are not displayed in the file upload user interface, and are not discussed here. See [Appendix C, “App Control Connector,”](#) for information on the process involved in uploading files for analysis.

Diagnostic files may be uploaded from agent computers, and in special cases, from the server. These are cataloged on a separate tab from general file uploads, but much of the user interface for acting on them is the same.

Enabling Access to File Upload Features

Important

- Permission for these features is *not* granted by default to the *admin* account or users with *Administrator* or *Administrator (Unified Management)* roles. You must explicitly add these permissions.
- While other App Control features provide data *about* files on agent-managed computers, these features allow a console user with the appropriate privileges to upload the actual file. These features should be used with extreme care, and in full compliance with your organization's policy on accessing other users' files. Be sure that only those App Control Console users that absolutely need access to the features are given permission to use them.

The following permissions control access to File Upload features:

- **Tools/View file uploads** – Ability to view uploaded files on the Requested Files page.
- **Tools/Manage uploads of inventoried files** – Ability to initiate manual file uploads from agent computers, and to create event rules that upload files. This permission applies only to files considered “interesting” (i.e., executables and scripts) by App Control.
- **Tools/Manage uploads of files by pathname** – Ability to initiate manual file uploads from agent computers. This permission enables uploading of a file by its pathname, even if not in the App Control inventory.
- **Tools/Access uploaded files** – Ability to download files uploaded to the server.

See [“User Role Permissions”](#) on page 106 for details on enabling feature access.

Scheduling Uploads

Several locations in the console provide access to commands for manually uploading files, including:

- the Events page (for events showing files that exist on computers)
- the Approval Requests page
- the File Catalog and Files on Computers tables
- the Find File Results table
- the Snapshot Contents table
- the File Details and File Instance Details pages
- the Computer Details page (for uploading a file by path only)

From most of these pages, you can upload a copy of any file that has been identified as "interesting" (i.e., executable) by App Control and has been added to the live inventory. From the Computer Details page, you can upload a copy of *any* file on the computer, whether or not it exists in your file inventory. For all uploads, the original file remains on the agent computer. Note that there are separate permissions for uploading files from the inventory and uploading any file by path.

Important

Uploading files greater than 2 gigabytes is not recommended. Files in excess of 2GB may fail to upload and show a "communication error".

In addition to performing manual uploads, you can create Event Rules that upload files when certain events take place. See "[Event Rules](#)" on page 520 for more information.

When you issue a successful upload command, a message appears on the console page indicating that the upload has been scheduled. In general, uploads begin almost immediately, but there could be delays depending upon other activities on the App Control Server and the size of the file you are uploading. Also, the App Control Server needs at least read permission to upload the file, and some files that are currently opened by other programs cannot be uploaded. If the App Control Server does not have read permission for a requested file on any agent-managed computer, the Uploaded Files table shows an error message for that file.

If an upload is scheduled for a file and no computer with that file is currently connected, the upload will be attempted later. Also, if a file upload is interrupted because of an agent-side error, it will be retried.

Starting Uploads of Inventoried Files from Tables

You can schedule the upload of one or more files at a time from the pages that include tables of files (File Catalog, Files on Computers, Events, etc.). When you request an upload, the App Control Server chooses the computer from which to upload a file matching the hash of the file you specified. It first searches for an instance of the file on a currently connected computer. If there are multiple connected computers with the file, the "best" computer is chosen based on how recently it communicated with the server and

whether any other uploads are scheduled or in progress (avoiding these is preferable). If the file does not exist on a connected computer, the server schedules the upload from a disconnected computer, and will start the upload when that computer reconnects.

To initiate a file upload from a file table:

1. Navigate to the file table page, such as Files on Computers.
2. Check the box(es) next to the file(s) you want to upload to the server.
3. On the Action menu, choose **Upload file to Server**.

The screenshot shows the 'Files on Computers' page. At the top, there are tabs for 'File Catalog' and 'Files on Computers'. Below the tabs, there are filters for 'Saved Views' (set to '(none)') and 'Group By' (set to '(none)' with 'Ascending' order). There are also buttons for 'Show Filters', 'Show Columns', 'Show Snapshot', 'Export to CSV', and 'Refresh Table'. Below this, there is an 'Action' dropdown menu with a red box around 'Upload file to server'. The main table shows a list of files with columns for 'Date Created' and 'Computer'. The selected file is '53b49363a4aeb3367196fc0d668a2aaa479a8701' with a date of 'Jul 18 2020 10:47:31 AM' and computer 'WORKGROUP\APC-58609'.

4. On the confirmation dialog, click **Yes**.

A message appears on the page indicating that the upload has been scheduled.

Starting Uploads from the File Instance Details Page

You can schedule the upload of a single file from the File Instance Details page or the File Details page. The procedure is the same.

To initiate a file upload from the File Instance Details page:

1. Navigate to the File Instance Details page for the file you want to upload.
2. On the Advanced menu to the right of the file data, choose **Upload File to Server**.

The screenshot shows the 'File Instance Details' page. On the left, there is a 'General' section with the following information: 'First Seen Name: newtonsoft.json.dll', 'First Seen Date: Jul 13 2020 07:24:52 AM', 'Last Updated: Jul 14 2020 06:16:55 PM', 'First Seen Path: c:\program files\windowsapps\microsoft.yourphone_1.20062.97.0_x64__8wekyb3d8bbwe\yourphoneappproxy\...', 'First Seen Computer: WORKGROUP\APC-58609', 'First Seen Platform: Windows', 'Extension: dll', 'Global State: Unapproved', 'Global State Details: File is unapproved, Publisher is not present, Certificate is not present', 'Flags: (none)', 'Installer / Updater: No', 'Reputation Enabled: Yes', and 'File Prevalence: File exists on 1 computer'. On the right, there is an 'Advanced' menu with options: 'View Carbon Black Reputation Data', 'Upload File to Server' (highlighted with a red box), 'Mark as Installer', 'Disable Reputation Approval for This File', and 'View Carbon Black Reputation Data' (with a button icon).

A message appears on the page indicating that the upload has been scheduled.

Once you upload a file from a Details page, the Upload File to Server command on the Advanced menu changes to **Related File Uploads**. Clicking on this link opens the

Requested Files page to the Uploaded Files tab, and filters it for the SHA-256 hash of this file.

Starting Uploads by Path from the Computer Details Page

You can schedule the upload of any file on a computer from its Computer Details page, whether or not the file exists in your file inventory of "interesting" files. Unlike uploads from other console pages, you must provide the path to the file in this case – there is no list of files to choose from, and the upload is not based on a hash.


Note

The ability to upload files via the Computer Details page requires a separate account permission – *Manage uploads of files by pathname*. See [“User Roles and Permissions”](#) on page 89 for instructions on setting this permission.


Although wildcards may not be used in the path to a file, you can specify the path location using macros and registry keys. See [“Using Macros in Rules”](#) on page 410 for the list of path macros recognized by App Control.

To initiate a file upload from the Computer Details page:

1. Navigate to the Details page for the computer that has the file you want to upload.
2. On the Advanced menu to the right of the file data, choose **Other Actions**.
3. On the Other Actions menu, choose **Upload File**



4. In the File box that appears in the menu, enter the complete path to the file you want to upload and then click the **Go** button.



A message appears on the page indicating that the upload has been scheduled. If you enter a non-existent file or path, the upload is still attempted, and you will not see an error on the page from which you initiate the upload, but a record of the failed attempt will appear in the Requested Files/Uploaded Files table.

Viewing the Uploads Table

Each requested upload appears on the Uploaded Files page, even when it fails. From this page, you can view information about the uploaded file, delete the upload from the list, retry the upload, cancel uploads in progress, and view the uploaded file.

To open the Uploaded Files page:

1. On the console menu, choose **Tools > Requested Files**.
2. If the Requested Files:Uploaded Files view is not already showing, click on the **Uploaded Files** tab.


	Request Date	Priority	Requester	Status	Computer	File Name	File Size
<input type="checkbox"/>	May 09 2017 12:53:52 PM	High	admin	Uploaded	MYCORPILT-3	two.bat	92 Bytes
<input type="checkbox"/>	May 09 2017 12:53:52 PM	High	admin	Queued	MYCORPILT-3	musictool.exe	8125 Bytes
<input type="checkbox"/>	May 09 2017 12:53:52 PM	High	admin	Uploaded	MYCORPDT-5	newbee.bat	103 Bytes
<input type="checkbox"/>	May 09 2017 12:53:52 PM	High	admin	Uploaded	MYCORPDT-5	newapp.bat	88 Bytes
<input type="checkbox"/>	May 09 2017 12:53:52 PM	High	admin	Uploaded	MYCORPDT-5	riskfile-test.bat	93 Bytes

On the Uploaded Files page, in addition to the default view, you can choose from among the following Saved Views:

- Uploads in Progress
- Completed Uploads
- Upload Errors

[Table 142](#) shows the columns available for the Uploaded Files page, some of which appear by default and some of which you must add.

Table 142: Uploaded Files Table Columns

Column	Description
Actions	<p>The Action column includes a checkbox for choosing files on which Action menu commands will act and buttons for taking action on individual files. The Action menu on this page includes the following commands:</p> <ul style="list-style-type: none"> • Cancel Uploads – Cancel the upload of checked files (if the upload has not been completed). • Retry Uploads – Retry the upload of checked files. • Delete Uploads – Delete the table rows for checked files, and, for successful uploads, delete the files from the server. • Change priority to: – Change the priority of this upload request to one of the choices on the menu. The choices are Low, Medium, High, and Highest. Changing priority affects the order in which any pending files are uploaded. • View Carbon Black Reputation Data – View any data available in the Carbon Black File Reputation database for this file (identified by hash) • Analyze with ... – If any third-party analysis devices or services are integrated through the App Control Connector, you can send selected files to them for analysis. For files that were not successfully uploaded to the Uploaded Files page, choosing an Analyze command initiates a new upload, and if that is successful, the file is submitted to the third-party device. <p>Individual uploaded file rows may be acted upon by the buttons in their row. These include the standard File Details and Find File buttons found in all file tables. There is one additional button for successfully uploaded files:</p> <p> Download the file (if it was successfully uploaded) from the App Control Server to a specified location. For this, console users must have specific permission to access uploaded files.</p>
Priority	Priority in which pending files are uploaded to the server. The priority choices are Low , Medium , High , and Highest . Can be changed on the Action menu.
Request Date	When the file upload was requested.
Requester	The console user that requested the upload, or “System” if the request was due to an event rule.
Status	<p>The status of the file upload. The possible values are:</p> <ul style="list-style-type: none"> • Uploaded – The upload completed successfully and the file is available on the server. • Uploading – The upload is in progress but not yet complete; a partial file has been received by the server. This status is likely to appear only for very large files. • Initiated – The upload task has been received by the agent where the file is located. • Queued – The upload task has not yet been sent to the agent. • Error – The upload failed. Hovering the cursor over this status displays the error message. Errors include: No file with hash, The system cannot find the path specified, The system cannot find the file specified. • Canceled – The upload was cancelled by a console user.

Column	Description
Computer	The name of the computer from which the file was uploaded.
File Name	The name of the uploaded file. For most requests, the App Control Server uploads a file matching the <i>hash</i> of the requested file, so in some cases, the name shown here will not be the same as the name of the file you chose. For uploads from the Computer Details page, the file name is always the name entered in the File box during the upload request.
File Size	The size (in bytes) of the file.
Upload %	The percent of the upload that is finished. Completed uploads show 100%. Failed uploads and uploads not yet started show 0%.
Upload Date	When the file was uploaded to the server.
Upload Directory	The directory on the App Control Server to which the file was uploaded. Value is “(default)” for manual uploads, which use the directory configured in the System Configuration Advanced Options tab. If the upload is due to an event rule, the actual path is shown.
Error	A description of the error that prevented the file from uploading. For example, the error for a file that was not present at the location given (or at all) would be file not found . Not shown by default.
File Path	The location on the agent computer from which the file was uploaded. Not shown by default.
Prevalence	The number of agent-managed computers reporting to your server on which this file is present.
MD5	The MD5 hash of the file.
SHA256	The SHA-256 hash of the file.
Source	Source of the request for upload. Either “Event rule” or “Manual”.
Source Name	If the request was due to an event rule, the name of the rule. If the request was manual, this field is empty.

Diagnostic Files

The Requested Files page also has an Diagnostic Files tab that shows diagnostic files uploaded from agent-managed endpoints to the App Control Server. There are two types of diagnostic files uploadable to the server, Server Diagnostic files and Agent Diagnostic Files:

- Server Diagnostic Files can be downloaded to a console user’s own computer by clicking the download button next to the checkbox for the file in table.
- Agent Diagnostic files remain on the server and do not have a download option.

The information and actions on the Diagnostic Files tab are generally used in conjunction with Carbon Black Support.

See [Appendix D, “Diagnostic Files,”](#) for more on uploading and downloading diagnostic files.


Downloading Uploaded Files

Once files are uploaded to the App Control Server, console users with the appropriate permissions can download selected files to their local computer for further examination.

Important

This feature in particular should be used with extreme care, in full compliance with your organization's policy on accessing other users' files. Be sure that only those App Control Console users that absolutely need access to the feature are given permission to use it. The ability to download files has its own setting ("Access uploaded files") in the console user permissions settings.

To download an uploaded file:

1. In the Uploaded Files table, click the download button  in the row for the file you want to download.
2. Follow the prompts for your browser to choose to download the file.

This copies a zip file to the download location on the computer on which the console is being viewed. The zip file includes the uploaded file and the folder path from the agent computer. You can navigate down through the folders to the file.

File and Path Information for Uploaded Files

Files you upload to the server are zipped before being uploaded, and while they appear in the App Control console under their original name, uploaded files are stored in numbered zip files on the App Control Server ('1.zip', '2.zip', and so on).

The contents of the zip file includes the uploaded file and its folder path from the agent computer. Handling of the file and path names depends on the characters used in both:

- **All ANSI characters** – If the name of the file and its path contain only ANSI-convertible characters, name and path information remain the same inside the zip file.
- **Non-ANSI characters** – If the file or its path contains non-ANSI-convertible characters, the file name and path will appear differently *inside the zip file* because the program used to zip the file does not support non-ANSI-convertible characters. In this case, the upload process creates a hard-link to the file in the Windows temp folder (C:\Windows\Temp). This link is then zipped using a new, generated file name beginning with "TMP" and a new path name. The extension will be included with the new file name unless it also contains non-ANSI-convertible characters.

Renaming has no effect on the file upload feature within App Control other than assuring that files with non-ANSI-convertible names can be uploaded successfully. However, if you are directly accessing the uploaded files through their zip file – for example, if you have created your own analysis process via an API or other means or you download the zip file itself – be aware that the zipped file name might differ from the file name on an endpoint.

Caution

The binaries of the original file and the renamed copy are the same. If a malicious file is uploaded under a temporary name it will still be malicious. Take the same precautions you would take with any suspicious uploaded file.

Upload Configuration Options

Deleting Uploaded Files

You can delete individual uploaded files from the server by checking the row for each file you want to delete on the Uploaded Files page and choosing **Delete Uploads** on the Action menu. You also can configure the App Control Server to delete files uploaded to the server on a schedule. By default, uploaded files are deleted after they have been on the server for 4 weeks.

To configure automatic deletion of uploaded files:

1. On the console menu, click the configuration (gear) icon, choose **System Configuration**, and then click on the **Advanced Options** tab on the System Configuration page.
2. Click the **Edit** button at the bottom of the page.
3. Under File Uploads, make sure the Delete Uploaded Files After box is checked, and enter the number of weeks after which you want the files to be deleted.

Note: Disabling automatic deletion of uploaded files is not recommended.

File Uploads

Delete Uploaded Files After: after 4 Week(s)

Default Upload Location:

4. Click the **Update** button at the bottom of the page.

Note

The actual uploaded files are not included in App Control Server backups, although the Uploaded Files table is backed up. If you restore a App Control database and there were files listed in the Uploaded Files table, the table is restored but the files will not be available.

Changing the Uploaded File Location

The default location of zipped, uploaded files is in the *Parity Server\Files* folder of the App Control installation directory. Uploaded files are stored in numbered zip files. For example, the first file you upload might be in the following location:

```
C:\Program Files (x86)\Bit9\Parity Server\Files\1.zip
```

You can change this location if you choose by editing the Default Upload Location setting on the System Administration/Advanced Options page (see the illustration above). You must have write permission for the upload location you specify:

- If you specify a folder without a full path, the location is assumed to be relative to the *Bit9\Parity Server* directory on the App Control Server. So, for example, the default location shown above is specified on the Advanced Options page simply as *files*.
- You can specify a full path, including a drive letter, on the App Control Server.

- You can use a full UNC path to specify a location on a system other than the App Control Server. You must have network access to the system you specify.

To change the target location for uploaded files:

1. On the console menu, click the configuration (gear) icon, choose **System Configuration**, and then click on the **Advanced Options** tab on the System Configuration page.
2. Click the **Edit** button at the bottom of the page.
3. In the File Uploads panel, enter the path for the location to which you want uploaded files sent and click the **Test** button to make sure that the location exists.
Note: If you specify a directory that does not exist, clicking Test may produce a failure message. However, if you have permission to write in the directory above the location you identified, the folder will be created and files will be uploaded to that location.
4. If the location you chose passed the test, click the **Update** button at the bottom of the page.

Note

You also can use Event Rules to automatically upload files that match the file specifications in a rule, and can define a new location for each rule. See [“Event Rules”](#) on page 520.

Appendix F

Exporting Data for External Analysis

This chapter provides instructions for configuring and using App Control External Analytics, which enables the App Control Server to export data it collects from endpoints to external analysis tools. This integration can enhance your ability to analyze data and makes it possible for the external tool to analyze data from multiple sources, including other App Control Servers.

Note

For this release, App Control has implemented the External Analytics integration with Splunk, and the examples shown here are Splunk-specific. However, the general description of configuration of data export as described in this appendix should enable integration with other external analysis tools by users with expertise in the setup of those tools.

Sections

Topic	Page
Overview	907
Preparing to Use External Analytics	907
Data Format and Management	908
Enabling External Analytics in the App Control Console	910
Enabling an External Tool for Data Analytics	914
Enabling Splunk to Collect App Control Data	915
Viewing App Control Data in External Analytics Tools	917
Using the Splunk App for App Control	918

Overview

App Control provides Syslog event output that can be analyzed and displayed by multiple different tools. Beginning with release v7.2, the App Control external analytics integration feature provides another way to utilize the extensive data collected by App Control. A App Control Server can be configured to send data to external data analytics tools, such as Splunk. Integrating App Control with an external analytics tool offers the following advantages:

- **Analyze Data from Multiple Sources** – You can view App Control information in context with streams of information from other data security platforms or multiple App Control Servers. For this release, data imported to Splunk can be normalized to the CIM standard.
- **Add App Control File Data to Analysis** – Unlike Syslog-based integrations, the external analytics integration is not limited to *event* log output. You can choose to export App Control event data, the file catalog, and/or file operations data to the external tool. The type and amount of data you send is configurable in the App Control Console.
- **Use New Reporting Capabilities** – You can use the capabilities of an external tool to generate new types of reports from your App Control data.
- **Shift the Analysis Load** – You can reduce the load on the App Control database server by moving data analysis to another tool and location.
- **Link the App Control Console to External Reporting Tools** – Enabling an analytics integration can add links from certain App Control Console pages to the external analysis tool console.

Data exported for external analytics is in JSON format.

Note

Available File Catalog data is described in [Chapter 7, “File, Publisher, and Application Information.”](#) The events available from App Control are described in the separate *VMware Carbon Black App Control Events Guide*.

Preparing to Use External Analytics

To use the external data analytics features, do the following:

- Configure the App Control Server to send data to a folder for external analytics.
- Enable one or more App Control Console user accounts with the privileges related to external analytics: *View System Configuration*, *Manage System Configuration*, and (to view links to and access external tools from the App Control Console) *View External Analytics Reports*. See [“User Role Permissions”](#) on page 106 for more on user privileges.
- If you plan to link back to the external tool from the App Control Console, make sure that the console users who will be using the analytics integration also have login accounts on the external tool.
- Configure your analytics tool to consume the output.

Data Format and Management

Data for external analytical tools is exported in JSON format. The JSON output from the App Control Server includes the field name with each value, making it easier both to view the raw output and to parse it later without creating indexing dependencies.

```
NetTrace-20140325-145639.bt9 - Notepad
File Edit Format View Help
[
  { "MessageTime": "2014-03-25 21:56:39.891", "HostId": 1, "HostName": "MYCORP\\BIT9SERVER1", "HostIP": "dc60::123b:9ab8:987d:3456", "Bit9Server": "bit9server1.mycorp.local", "RequestHeader": { "Method": "REPORT_AB_LIST", "MethodVersion": 11}, "Timestamp": "2014-03-25 21:06:49", "FileHash": "", "PathName": "c:\\windows\\microsoft.net\\framework64\\v2.0.50727\\mscorsvw.exe -useclsid {701c54a0-0cdd-4eca-b729f72bf83451} -comment \\\"compile worker for microsoft.web.management.iisclient, version=7.5.0.0, culture=neutral, publickeytoken=31bf3856ad364e35, processorarchitecture=msil\\\"", "FileName": "c:\\windows\\microsoft.net\\framework64\\v2.0.50727\\mscorsvw.exe", "SourceFilePathName": "", "SourceFileName": "", "Flags": 0, "ProcessPath": "c:\\windows\\microsoft.net\\framework64\\v2.0.50727\\mscorsvw.exe", "ProcessHash": "", "InstallerHash": "654013b8fd229a50017b08dec6ca19c7dda8ce0771260e057a92625201d539b1", "ProcessHashType": "", "Ieid": 2, "MsiIeid": 0, "OpType": 9, "LocalState": 8, "FileHashType": 0, "InstallerHashType": 5, "ProcessHashType": -1, "DetachedPublisher": "", "TrustedDirectoryId": 0, "ProcessKey": "00000001-00000000-00000000-000000005331EFD9"},
  { "MessageTime": "2014-03-25 21:56:40.300", "HostId": 1, "HostName": "MYCORP\\BIT9SERVER1", "HostIP": "dc60::123b:9ab8:987d:3456", "Bit9Server": "bit9server1.mycorp.local", "RequestHeader": { "Method": "REPORT_AB_LIST", "MethodVersion": 11}, "Timestamp": "2014-03-25 21:07:40", "FileHash": "", "PathName": "c:\\windows\\microsoft.net\\framework64\\v2.0.50727\\mscorsvw.exe -useclsid {63f258b3-6dd5-478f-8d2c0a36852cfe8f} -comment \\\"compile worker for microsoft.datawarehouse.vsiintegration, version=10.0.0.0 culture=neutral, publickeytoken=89845dcd8080cc91\\\"", "FileName": "c:\\windows\\microsoft.net\\framework64\\v2.0.50727\\mscorsvw.exe", "SourceFilePathName": "", "SourceFileName": "", "Flags": 0, "ProcessPath": "c:\\windows\\microsoft.net\\framework64\\v2.0.50727\\mscorsvw.exe", "ProcessHash": "", "InstallerHash": "39d3630e623da25b8444b6d3aaab16b98e7c289c5619e19a85d47b74c71449f3", "ProcessHashType": "", "Ieid": 15, "MsiIeid": 0, "OpType": 9, "LocalState": 8, "FileHashType": 0, "InstallerHashType": 5, "ProcessHashType": -1, "DetachedPublisher": "", "TrustedDirectoryId": 0, "ProcessKey": "00000001-00000000-000001B8-000000005331F01B"}
]
```

If you are using the Splunk App for App Control, App Control data imported by the Splunk Server is mapped to the CIM so that it can be integrated with other data. See [“Field Mappings to CIM in the Splunk App for App Control”](#) on page 924 for details.

Depending upon which messages you enabled for export, one or more of the following files will appear in the Export Directory configured for External Analytics:

- **Event Data** – EventTrace-<YYYYMMDD>.bt9
- **File Catalog Data** – MetadataTrace-<YYYYMMDD>.bt9
- **File Operations Data** – NetTrace-<YYYYMMDD-HHMMSS>.bt9

Name	Date modified	Type	Size
EventTrace-20170508.bt9	5/8/2017 1:52 PM	BT9 File	91 KB
MetadataTrace-20170508.bt9	5/8/2017 1:52 PM	BT9 File	24,325 KB
NetTrace-20170508-135625.bt9	5/8/2017 1:56 PM	BT9 File	5 KB

Each message log file will grow to a maximum of 512 megabytes, at which point a new log file will be created. New logs are also started when the App Control Server processes are restarted.

New File Operations data files (NetTrace) are named with both date and time as described above.

If two Event data or File Catalog data files are created on the same day, a number is appended to the second one of each. For example, the first file catalog data file created on October 29, 2013, would be named *MetadataTrace-20131029.bt9*. If that file reached its size limit that same day, the second file would be named *MetadataTrace20131029-1.bt9*.

Note

See the separate *VMware Carbon Black App Control Events Guide* for more information about event types and subtypes that may be exported.

Data Volume for Exported Analytics

- 20KB per computer per day of file catalog
- 75KB per computer per day of events
- 135KB per computer per day of file operations (volume: High)
- 115KB per computer per day of file operations (volume: Medium)
- 100KB per computer per day of file operations (volume: Low)

Limiting Export Directory Size

There is a checkbox on the console External Analytics tab of the System Configuration page that allows you to limit the amount of data in the Export Directory. Checking this box displays a field in which you can enter the number of gigabytes of data to set as the maximum export directory size (i.e., the total size of all files in the Export Directory). When the limit is reached, files are deleted by age (oldest first) until the directory size is under the limit. The lowest allowable size limit is 3 GB. The current files in each category are never deleted. The upper limit is 10 petabytes.

Note

The Export Directory size limit controls the amount of data kept in the directory on the App Control Server but does not limit the amount of data uploaded to the external analysis tool. If you need to limit the data going to the external tool for licensing or performance reasons, use the External Analytics Settings checkboxes and radio buttons on the External Analytics configuration page, as described in [“Enabling External Analytics in the App Control Console”](#) on page 910.

Local vs. Network Log Files

When log files are local and the log content is relayed to the data analytics tool by a mechanism designed for that purpose, such as the Splunk Universal Forwarder, performance impact is expected to be minimal. However, if log files are written to a network location, there could be a delay in data availability if the network latency is high.

When analytics data is written locally, it is best to have it written to a disk other than the one on which the operating system or App Control SQL database are located.

Enabling External Analytics in the App Control Console

You configure three elements of the analytics features in the App Control Console:

- On the System Configuration page External Analytics tab, you specify the location, content, and size limitation (if any) for folder into which App Control data is exported.
- On the same tab, you can provide URLs and query specifications so that console users can link to specific reports on an external analytics server.
- On the Add Custom Rule page, you can create a rule that will ignore files written to the data export directory to reduce the impact of data exports on the App Control Server.

The following procedure describes how to accomplish the first two tasks on this list. [Table 143, “External Analytics Configuration Options”](#) on page 912 provides more detail on the parameters on the External Analytics tab.

To enable External Analytics features in the App Control Console:

1. On the console menu, click the configuration (gear) icon and choose **System Configuration** and click on the **External Analytics** tab.
2. Click the **Edit** button at the bottom of the page.

3. In the General panel, check the **Enable Export** box.
4. In the Export Directory field, enter the name of the directory into which you want App Control analytics files written. This folder must be one for which the user running the App Control Server service (*ParityServer*) has write access.
Note: If you plan to write exported data to the system hosting the App Control Server, do not use a disk volume used by the operating system or SQL Server.

5. Click the **Test** button to the right of the Export Directory field to test whether the directory is valid and the server process has write access to it.
6. In the Messages fields, specify what type of information you want to export:

External Analytics Settings

General

Enable Export:

Export Directory: Test

Messages:

File Catalog
Information about files. Traffic estimate: 20 KB/day
 Export complete catalog (est. 18442 KB) plus new files
 Export only new files

File Operations
Operations related to files on specific computers. Traffic estimate: KB/day
 Volume:

Events (No events exported)
Blocks, approvals and requests. Traffic estimate: 75 KB/day
 Include entire event backlog (est. 5095 KB) plus new events
 Include event backlog going back minute(s) (est. 0 KB) plus new events
 New events only

Limit Export Directory Size:

- a. **File Catalog** – Check this box to export File Catalog data to the export directory. Checking the box displays two radio buttons: **Export complete catalog** exports the entire current contents of the File Catalog and any new additions to the catalog. **Export only new files** exports only unique, new files discovered on agents reporting to your App Control Server once this option has been enabled.
- b. **File Operations** – Check this box to export messages from agents about operations that affect files. A dropdown menu lets you determine the volume, and by extension the type, of the data that is exported. See [Table 143](#) for details.
- c. **Events** – Check this to export App Control events. See [Table 143](#) for details about the radio button options that control the amount of Event data that is exported and display the estimated size of the export where available.

Note

When setting these Message export options, consider the traffic estimate values shown for each one and any traffic limits on the external analysis device. However, also be sure you are exporting enough data to allow for useful analysis.

7. The Analytics Server Reports section allows configuration of links from the App Control Console to reports on the external analytics server. If you want to enable these links, begin in the Root URL field, by entering the root URL of the analytics tool with which you are integrating App Control.
8. In the Analytics Server Reports panel, enter and test the Relative URL and Query string for each type of report listed. Use the marker `<va1>` in the query string to represent what is being passed (file hash, machine name, user name) to the analytics tool.
9. Click **Update**.

Table 143: External Analytics Configuration Options

Field/Button	Description
Enable Export	This checkbox activates and deactivates the External Analytics integration features, including data export and links to external analytics tools.
Export Directory	This field determines the directory to which the App Control Server exports data for external analysis. The Test button allows you to confirm that the directory is valid and that the server process has write access to it. The test results appear next to the button (either 'OK' for success or a message explaining why the test failed).
Messages: File Catalog	This checkbox enables export of File Catalog data to the export directory. Checking the box displays two radio buttons that control the amount of File Catalog data that is exported: <ul style="list-style-type: none"> • Export complete catalog – This option exports the entire current contents of the File Catalog and continues exporting any new additions to the catalog. • Export only new files – This option exports only unique, new files discovered on agents reporting to your App Control Server.
Messages: File Operations	This checkbox enables export of messages from agents about operations that affect files. A dropdown menu lets you determine the volume, and by extension the type, of data that is exported: <ul style="list-style-type: none"> • Low – Export messages about file Create, Modify, Delete, Rename, and Rename Directory operations. • Medium – Export all messages in Low plus messages about file state changes (Approved, Unapproved, Banned); this includes both individual file state changes and operations that cause state changes in groups of files. • High – Export all file operations messages.
Messages: Events	This checkbox enables export of App Control events data. Checking the box displays radio buttons that control the amount of Event data that is exported, and displays the estimated size of the export where available: <ul style="list-style-type: none"> • Include entire event backlog (est. value KB) plus new events – This exports the entire existing event database and enables ongoing export of new events. • Include event backlog going back [time value] (est. value KB) plus new events – This allows you to choose a time period of past events (starting from the present) to export and enables ongoing export of new events beginning when this is enabled. • New events only – This enables ongoing export of new events only beginning when this setting is enabled.
Limit Export Directory Size	Checking this box displays a field in which you can enter the number of gigabytes of data to set as the maximum export directory size (i.e, the total size of all files in the Export Directory). When the limit is reached, files are deleted by age (oldest first) until the directory size is under the limit. The lowest allowable size limit is 3 GB. The current files in each category are never deleted.

Field/Button	Description
Root URL	<p>The root URL (optionally including the port) entered here points to the analytics server with which you are integrating the App Control Server. This is used as the base URL for links from App Control Console pages back to reports on the analytics server.</p> <p>Note: The App Control Console user must have credentials to log into the external server, and the URL provided must allow the user to log in with those credentials, even when using the App Control Console to reach it.</p>
File Details Report	<p>This defines a link to a File Investigation report on the analytics server. There are two fields to define the line: Relative URL, which is appended to the Root URL you define, and Query String, which defines the report you want from that URL.</p> <p>When defined, this File Analytics link appears in the External Pages menu on the File Details and File Instance Details pages.</p> <p>Click the Test button to the right of this line to confirm that the URL and query definition are valid.</p>
Computer Details Report	<p>This defines a link to a Computer Investigations report on the analytics server. There are two fields to define the line: Relative URL, which is appended to the Root URL you define, and Query String, which defines the report you want from that URL.</p> <p>When defined, this Computer Analytics link appears in the External Pages menu on the Computer Details page.</p> <p>Click the Test button to the right of this line to confirm that the URL and query definition are valid.</p>
User Details Report	<p>This defines a link to a Console User Search (in this case, App Control Console Login Accounts) report on the analytics server. There are two fields to define the line: Relative URL, which is appended to the Root URL you define, and Query String, which defines the report you want from that URL.</p> <p>When defined, this User Analytics link appears in the External Pages menu on the Edit Login Account page.</p> <p>Click the Test button to the right of this line to confirm that the URL and query definition are valid.</p>
Set Analytics URLs to Splunk defaults	<p>Clicking this button inserts Splunk default Relative URL and Query String definitions into the three report fields. It also inserts "http://server:8000" in the Root URL field (port 8000 is the Splunk default).</p> <p>When you replace "server" with a valid Splunk server URL, these defaults should allow access to valid Splunk reports from the App Control Console.</p>
Clear Analytics URLs	<p>Clicking this button clears all values from the Analytics Server and Analytics Server Reports fields.</p>

Editing or Disabling the External Analytics Integration

If you need to modify or disable the external analytics integration with the App Control Server, you can use the External Analytics tab on the System Configuration page in the App Control Console. The Export Directory and any additional components installed for an integration, such as the Splunk Universal Forwarder, are not deleted or uninstalled when you disable the integration through the App Control Console.

Adding a Custom Rule to Ignore Analytics Log Files

When External Analytics is enabled, there will be repeated, ongoing file write operations in the Export Directory. Normally, this would generate significant event traffic on the App Control Server if an agent is active on the server. Since this event traffic is not usually interesting to track, consider creating a custom rule to exclude tracking of files in the Export Directory. See [Chapter 14, “Custom Software Rules,”](#) for more about how these rules may be configured.

To exclude tracking of exported analytics files:

1. On the console menu, choose **Rules > Software Rules** and click on the **Custom** tab.
2. Click the **Add Custom Rule** button.
3. On the Add Custom Rule page, provide the necessary information to create a rule that will ignore writes to the Export Directory for analytics data:
 - a. **Name** – Choose a name to clearly identify the rule; for example, *Ignore Data Analytics Log Files*.
 - b. **Description** – (Optional) Add a description to further identify the rule purpose.
 - c. **Status** – Click the **Enabled** radio button.
 - d. **Platform** – Choose the platform to which the rule is applied; this is **Windows** (the default) for Export Directories that are on the App Control Server system.
 - e. **Rule Type** – Choose **Performance Optimization**.
 - f. **Path or File** – Provide the Path and Name of the folder where analytics files are written; for example, `D:\CbProtectionAnalytics`.
 - g. **Process** – Choose **Specific Process**, then enter and **Add** the processes that App Control uses to write these files. For example, if you are running a 64-bit OS and used the default App Control installation directory, you would use:

```
<ProgramFiles>\Bit9\Parity Server\ParityServer.exe  
<ProgramFiles>\Bit9\Parity Server\Reporter\ParityReporter.exe
```
 - h. **Rule Applies To** – Choose **All policies** or if you prefer just the policy that the system being written to (usually the App Control Server) belongs to.
4. When you have finished configuring the rule, click the **Save** button. The new rule is added to the Custom Rules table.

Enabling an External Tool for Data Analytics

In addition to configuring the App Control Server to export data and (optionally) connect to an analytics server for reports, you must configure a connection for the analytics server to access the exported data. The exact steps for enabling a particular external tool for access to App Control data will vary, and can include actions taken on the App Control Server system as well as those taken on the analytics server. The next section provides the steps for enabling Splunk for App Control data access.

Enabling Splunk to Collect App Control Data

To enable a Splunk server to import App Control data for analysis, you must make modifications on both the system hosting the App Control Server and the Splunk server. The summary of these steps is as follows:

- Have a Splunk Server running and network-accessible to the App Control Server.
- Set up the Splunk Server to receive messages from the Splunk Forwarder.
- Install the Splunk App for App Control on the Splunk Server.
- Install the Splunk App for App Control on any machines running Splunk Indexer that are not on the machine running Splunkweb.
- Install the Splunk forwarder on the App Control Server.
- Install the Splunk App for App Control on the Splunk Forwarder.

Note

Instructions for setting up the App Control App for Splunk are also on the Splunk web site and might be more recent than those provided here. See:

<https://splunkbase.splunk.com/app/1790/#/details>

Configuring the Splunk Server for App Control Access

You must complete several procedures on the Splunk Server to enable use of App Control data for analytics. First, configure the Splunk Server to receive forwarder data on port 9997.

To set up the Splunk server to receive Splunk Universal Forwarder messages:

1. Log into the Splunk server as an administrator-level user.
2. In the menu bar at the top of the Splunk console, choose **Settings** (Splunk 6) or **Management** (Splunk 5), then choose > **Data > Forwarding and receiving**, and in the *Forwarding and receiving* window, choose **Configure receiving**.
3. In the *Receive data* window, check to see whether port 9997 is configured. If not, click the **New** button, enter **9997** as the port to listen on, and click the **Save** button.
4. In your firewall, create a rule to allow the Splunk Server to receive data on port 9997.

The Splunk App for App Control allows Splunk to interpret data provided by App Control so that it can be analyzed and displayed by Splunk.

To install the Splunk App for App Control on the Splunk Server:

1. Log into the Splunk server as an administrator-level user.
2. Search for “App Control” through the **Find Apps Online** feature in the Splunk console, and when you find the App Control App for Splunk, download it to a convenient location on the server.
3. In the menu bar at the top of the Splunk console, choose **Apps > Manage Apps**.
4. Install the App from its zip file:

- Click on **Install app from file** and in the *Upload an app* dialog, browse to the `cb-protection-app-for-splunk_20.tar.gz` file. and then click **Upload**. The file name, especially the numbers at the end, may vary with version changes.

Note

If you have Splunk indexers that are not on the machine running the Splunkweb, also install the Splunk App for App Control on the machines hosting these indexers. The procedure for this is the same as for installing the app on the Splunk Forwarder. See “[To install the Splunk App for App Control on the App Control Server:](#)” on page 916

Installing the Splunk Forwarder and App on the App Control Server

In addition to configuring External Analytics on the App Control Console, there are two additional steps you must take on the system hosting the App Control Server, outside the console itself, to enable Splunk connectivity:

- Install the Splunk Universal Forwarder
- Install the Splunk App for App Control in a Forwarder subdirectory

The Splunk Universal Forwarder is a package that can be installed on systems so that Splunk can collect data from them, for example from log files. In this case, when the Forwarder and Splunk App for App Control are installed, the Forwarder collects data from the App Control export folder and directs it to the correct location in the Splunk infrastructure.

Important

During the Splunk Universal Forwarder installation process, *do not* enter the location of the data files on the App Control Server when prompted. The location of these files will be provided by the Splunk App for App Control.

To install the Splunk Forwarder on the App Control Server:

1. Download the forwarder from the Splunk website:
<http://www.splunk.com/download/universalforwarder>
2. Run the appropriate installer for your operating system on the App Control Server.
3. Provide the address of your Splunk Server when prompted.
4. Once the Splunk Forwarder is installed, install the Splunk App for App Control under the Splunk Forwarder installation directory, as instructed below.

To install the Splunk App for App Control on the App Control Server:

1. Search for and download the Splunk App for App Control from the Splunk apps website:
<https://splunkbase.splunk.com>

2. Copy the downloaded file, for example, `cb-protection-app-for-splunk_20.tar.gz`, to the `\etc\apps` subdirectory under the Splunk Forwarder installation directory. For example, if you are running a 64-bit OS on the App Control Server copy the file to:

```
C:\Program Files\SplunkUniversalForwarder\etc\apps\
```

Note: There may be numbers at the end of the file name that vary with app version changes.

3. Unzip and untar the file.
4. Go into the `bit9-secapp` directory and create a new directory named `local`.
5. Copy `default\inputs.conf` into the `local` directory.
6. Edit the first line of `local\inputs.conf` to point to the location of the Export Directory configured on the System Configuration/External Analytics page of the App Control Console, and save the file. For example, if the Export Directory on the App Control Server is `D:\Bit9\LogFiles`, change the first line of `inputs.conf` to the following:

```
[monitor://D:\Bit9\LogFiles\*.bt9]
```

7. At a command prompt, restart the Splunk Forwarder:

```
cd \Program Files\SplunkUniversalForwarder\bin
.\splunk.exe restart
```

When you have completed all of the tasks described in [“Enabling External Analytics in the App Control Console”](#) and [“Enabling an External Tool for Data Analytics”](#), the Bit9-Splunk integration is complete and data from App Control should begin flowing to Splunk.

Viewing App Control Data in External Analytics Tools

How an external analytics tool uses App Control data will vary according to the reporting capabilities of the tool and its ability to integrate App Control information with streams of information from other data security platforms. Users of this integration are assumed to have knowledge of how to integrate data from various sources in their analytics tool, and how to create reports that make use of that information. The specific reports that are created are up to each user.

One way to make use of reports on an external analytics tool is to link to them from the App Control Console.

Linking to an External Tool from the App Control Console

The External Analytics tab of the console System Configuration page provides fields in which links to reports on an external tool may be defined. If a Root URL and Analytics Reports for each category are configured on this page, the following links appear in the right menu of their respective pages:

- **Computer Analytics** – This appears on the Computer Details page.
- **File Analytics** – This appears on the File Details and File Instance Details pages.
- **User Analytics** – This appears on the Edit Login Account page.

The content of the pages displayed when a console user clicks one of these links is completely determined only by the URL and query definitions provided on the

configuration page. The App Control Console user must have credentials to log into the external server, and the URL provided must allow the user to log in with those credentials even when using the App Control Console to reach it.

See [“Enabling External Analytics in the App Control Console”](#) on page 910 for details of how these links are enabled. See [“Using the Splunk App for App Control”](#) for an example of what these reports might contain.

Using the Splunk App for App Control

The *Splunk App for App Control* helps Splunk present App Control data more effectively. Installing and configuring the app adds a set of dashboards specifically for displaying App Control data. It also enhances Splunk’s ability to handle App Control data in other views, for example, by identifying App Control as the source of the data, identifying the purpose of each keyword in the key/value pairs, decoding App Control-specific values, and by mapping App Control fields to the Common Information Model (CIM) so that App Control data can be combined with data from other sources.


Dashboards in the Splunk App for App Control

The Splunk App for App Control includes the following dashboards:

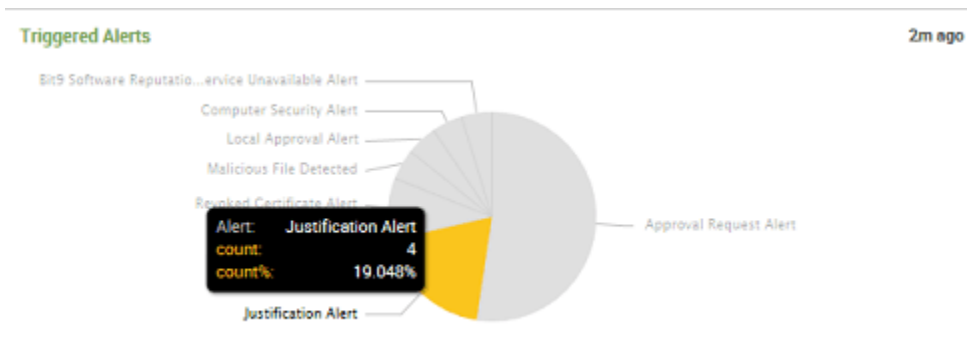
- **Deployment Activity** – Overview of information available from App Control installation.
- **Activity Details: File Activity** – Information about file creation and modification activity on App Control-managed computers.
- **Activity Details: Blocks** – Information about files blocked on App Control-managed computers.
- **Activity Details: Approvals** – Information about files approved on App Control-managed computers.
- **Activity Details: New Unapproved Files** – Information about new files that are discovered on App Control-managed computers and neither approved nor banned.
- **Activity Details: Events** – Information about events recorded on the App Control Server.
- **File Investigation** – Information suitable for a malware investigation focused on a specific file or files. If you link from the App Control Console, this provides information about the file from whose details page you linked.
- **Computer Investigation** – Information suitable for a malware investigation focused on a specific computer or computers. If you link from the App Control Console, this provides information about the computer from whose details page you linked.
- **Console Users** – Information suitable for discovering anomalous or risky actions performed by a specific App Control Console user or users. If you link from the App Control Console, this provides information about the user from whose details page you linked.
- **All Console Users** – Information about all App Control Console users.

Each of these dashboards contains panels that display information imported into Splunk from a App Control Server. Some also include a summary panel at the top. If you have used the Dashboard in the App Control Console, some of these panels will be familiar. However, here they can take advantage of the analysis and multi-source integration capabilities of Splunk. [Table 144](#) shows the panels available on the Splunk App for App Control dashboards, and identifies the dashboards on which they appear.

Panels in these dashboards may include tables of data or charts that graphically display the data, such as the display of Triggered Alerts in the following example. Some panels include both.



When you hover the mouse over a section of the chart, such as a pie chart slice or a bar in a bar chart, a legend appears describing the data represented that section.



If you click on one of these sections, the underlying data is displayed.

Events (4)		Statistics	Visualization						
Format Timeline ▾ Zoom Out Zoom to Selection Deselect									
List ▾ Format ▾ 20 Per Page ▾									
Hide Fields All Fields Selected Fields @ host 1 @ source 2 @ sourcetype 1 Interesting Fields @ ABId 1 @ ABState 1		<table border="1"> <thead> <tr> <th>i</th> <th>Time</th> <th>Event</th> </tr> </thead> <tbody> <tr> <td>▶</td> <td>3/10/14 1:54:55.000 PM</td> <td> <pre> {[-] ABId: ABState: BanName: Bit9Server: IT-Server-2 CLVersion: EventParam1: Justification Alert: 1 justification has been created. EventParam2: EventParam3: EventSubType: Alert triggered EventSubTypeId: 1104 </pre> </td> </tr> </tbody> </table>		i	Time	Event	▶	3/10/14 1:54:55.000 PM	<pre> {[-] ABId: ABState: BanName: Bit9Server: IT-Server-2 CLVersion: EventParam1: Justification Alert: 1 justification has been created. EventParam2: EventParam3: EventSubType: Alert triggered EventSubTypeId: 1104 </pre>
i	Time	Event							
▶	3/10/14 1:54:55.000 PM	<pre> {[-] ABId: ABState: BanName: Bit9Server: IT-Server-2 CLVersion: EventParam1: Justification Alert: 1 justification has been created. EventParam2: EventParam3: EventSubType: Alert triggered EventSubTypeId: 1104 </pre>							

These panels provide other standard Splunk features, such as the ability to change the time period for which data is displayed.

Table 144: Panels in Splunk App for App Control Dashboards

Dashboard	Panel	Description
Deployment Activity	Host Activity	File and event activity by agent computer.
	Triggered Alerts	Number of triggered alerts by type.
	File Blocks	Blocked files by date, computer, and product name.
	New Unapproved Files	Events reporting new unapproved files appearing on agent computers by date.
	New Files in Catalog	Unique new files added to the catalog by date.
	Approvals	File approvals by date, computer, and product name.
	File Activity	Creation and modification of files on App Control-managed systems by date, computer, and product or file name.
	Top Event Subtypes	Event subtypes listed by frequency.
Activity Details: File Activity	File Creations	File creations by date, computer, product or file name, and process.
	File Modifications	File modifications by date, computer, product or file name, and process.
Activity Details: Blocks	Blocks	File blocks by date, computer, and file name.
	Block Distributions	File blocks by file trust level, process that attempted to execute the file, and product name.
	Block Sources	File blocks by rule that blocked the action, reason (event subtype), and publisher.
Activity Details: Approvals	Approvals	File approvals by date, computer, and file name.
	Approval Distribution	File approvals by file trust level, process that generated, modified or executed the file, and product name.
	Approval Sources	File approvals by rule that approved the file, reason (event subtype), and publisher.

Dashboard	Panel	Description
Activity Details: New Unapproved Files	New Unapproved Files	New unapproved files appearing on agent computers by date.
	New Unapproved Files By Product Name	New unapproved files listed by publisher
Activity Details: New Unapproved Files (continued)	Top New Unapproved File Hashes	New unapproved files listed by hash, ranked from most to least instances
	Top New Unapproved File Names	New unapproved files listed by name, ranked from most to least instances.
	Potentially Malicious Files	New unapproved files identified as potentially malicious by Carbon Black File Reputation.
	Top Computers	Computers with new unapproved files, ranked from most to least instances.
	Known Trust Values	New unapproved files listed by Carbon Black File Reputation trust values (if known).
	Top Users	New unapproved files listed by users, ranked from most files to least files.
Activity Details: Events	Events	Events by date.
	Top Event Subtypes	Events by event subtype, ranked from most to least instances.
	Errors	Error messages.
	Top Computers	Events by agent computer referenced in event, ranked from most to least instances.
	Top Users	Events by users referenced in event, ranked from most to least instances.
	Top Event Types	Event by event type, most to least. Event types include multiple subtypes.

Dashboard	Panel	Description
File Investigation	Number of computers on which this file has been created	The prevalence of this file on App Control-managed computers reporting to this server.
	File Hashes	For file searches by name, the hashes identified for files with this name.
	File Information: First Seen on Network	The first seen name of this file on App Control-managed computers reporting to this server.
	Hash Activity	A time-based bar chart describing creations and modifications of and by files with this hash.
	Other Hashes with First Seen Name	Other hashes with the same first seen file name on App Control-managed computers reporting to this server.
	Files Modified By This File	Files for which this file is the process, presented as a simple table of events in reverse-chronological order
File Investigation (continued)	Top Hashes Modified by This File	Files for which this file is the process, sorted by the number of times a particular file (identified by hash) was modified by the specified file/hash
	Top Event Subtypes Containing This File	Event subtypes containing a reference to this file, ordered from subtypes containing the most instances of this file most to least.
	Top Rules Containing This File	Rules referencing this file, including those identifying its file/hash as the process, the installer, or the file being acted upon. The rules appear in descending order by how often they reference the specified file.

Dashboard	Panel	Description
Computer Investigation	Detection Events	Table of events related to App Control advanced threat indicators.
	Risky Behavior	Table of events related to issues with tamper protection, or the detection of potentially risky or malicious files on an agent computer.
	Risky Behavior Timeline	Amount of Risky Behavior graphed over time.
	Blocks	Chart of blocked file actions on this computer by date.
	New Files	Table of new files on the computer(s) specified in the search.
	File Activity	Chart of file creations and modifications by date.
	Approved Files	Table of files approved and the rule used for the approval.
	Events Chart	Chart of the top 10 most frequent event subtypes involving the specified computer(s) over the search time period.
	Health Checks	Table of App Control Health Check events and the results of the Health Check.
Console User Search	Events	Events that reference this user, charted by date.
	User Activity	Events that reference this user, in a table with additional detail, listed in date order.
	New or Removed Console Users	Console users that were created or deleted by this user.
	Custom Rules Actions	Creation and modification of custom rules by this user.

Dashboard	Panel	Description
Console User Search (continued)	File Approvals	File approvals by this user.
	File Bans	File bans by this user.
	Policy Management by Subtype	Policy management actions taken by this user, including policy modification and creation, and writing of agent installer files due to policy actions.
	Global Approval by Trust	Global approvals by the user, by trust.
	Globally Approved Hashes	Hashes globally approved by this user.
	Local Approval by Trust	Local approval by this user, by trust.
	Top Locally Approved Hashes	Files (by hash) locally approved by this user, most to least.
All Console Users	Events	Events by all console users, charted by date.
	Policy Management	Policy management events by date and user.
	Computer Management	Computer management events by date and user.
	Session and General Management	Session and General management events by date and user.
	Top Ten User - Global Approvals	Top ten users creating the most global file approvals.
	Top Ten User - Local Approvals	Top ten users creating the most local file approvals.

Field Mappings to CIM in the Splunk App for App Control

The Splunk Security Tool requires that data is normalized so that it can be processed and analyzed the same way, regardless of the source. The Splunk App for App Control maps the fields in App Control data analytics output to the Common Information Model (CIM). See <http://www.dmtf.org/standards/cim> for more information on the Common Information Model.

[Table 145](#) shows the CIM mappings done in the Splunk App for App Control.

Table 145: App Control Data-to-CIM Mappings in Splunk

App Control Field	CIM Field
HostName	src_nt_host, dest_nt_host, dest, dvc_nt_host
HostIP	src_ip, dest_ip, dvc_ip
FilePath	file_path
FileHash	file_hash, hash
FileName	file_name
FileSize	file_size, size
Message	change_type
EventSubType	action
Timestamp	modtime

Index

App Control database. See database, App Control

A

acknowledging

 devices 367, 368

 files 241

 publishers 281, 284

Active Directory Integration

 AD computer metadata in App Control 164

 AD logins in console 92

 AD policy mapping 122

 AD user details in console 94

 and agent installation 134

 and Windows 2000 domain

 controllers 727

 clearing the AD server cache 130

 moving computers to another policy 170

 overview 46

 security domain for logins 727

 testing 123

AD logins in console

 disabling 93

 enabling 92

AD policy mapping rules 123

Agent

 blocked file notifiers on 538

 computer configuration 118

 connection status 161

 default data directory 139

 disabling 183, 199

 downloading installers for 134

 enabling automatic upgrade 147

 enabling management privileges 727

 file initialization for 119

 health check for 166

 installing 136

 installing on Linux computers 144

 installing on Mac computers 142

 installing on Windows computers 138

 manual upgrade on Windows

 computers 149

 non-default data directory 138

 package generation status 133

 policy status of 157

 prioritizing updates to 166

 registration with server 129

 reporting command lines on 601

 requesting update for 167

 rules out of date for 157

 securing communications with 739

 self-protection 190

 temporary policy override for 297

 uninstalling 154

 uninstalling from a Mac computer 155

 uninstalling from a Windows

 computer 154

 uninstalling on Linux computers 155

 upgrade status 153

 upgrading 147

 upgrading by policy 185

 upgrading from console 148

 uploading agent installers to the

 server 131

 using anti-virus software with (macOS/OS X) 142

 using anti-virus software with (Windows) 139

 using with anti-virus software (Linux) 145

 verifying installation 146

 viewing installer versions 133

Agent Disabled mode. See disabled mode

agent rules files

 uploading to the server 131

alerts 608

 alert history 623

 alerts page 608

 approval request 610

 baseline drift 612

 blocked file 612

 Carbon Black File Reputation

 unavailable 610

 computer in local approval 610

 computer security 611

 configuring e-mail 754

 creating 612

 deleting 619

 disabling 618

 editing 618

 event alerts 612

 file prevalence 612

 for threat detection 676

 how triggered 619

 justification 610

- new certificate 347
- on home page 58
- propagating file 612
- rapid config modified 519, 610
- resetting 621
- revoked certificate 347
- system health 812
- types 612
- updater modified 271, 610
- Alerts page 608, 613, 658
- algorithms for certificates 288
- analysis environment
 - for WildFire notifications 874
- analytics
 - exporting data for 906
- analyze file
 - on Approval Request Details page 584
- analyzing files
 - automating with event rules 520
 - using Check Point for 883
 - using WildFire for 883
- anti-virus software
 - and Agent (Linux) 145
 - and Agent (macOS/OS X) 142
 - and Agent (Windows) 139
 - enabling updaters for 266
- App Control Agent
 - diagnostic files for 890, 902
- App Control Console
 - defined 5
- App Control Server
 - installing. See App Control Installation Guide
 - overview 39
- applications
 - computers that have 257
 - viewing metadata for 256
- Applications by Publisher/Company view 226
- Applications page 256
- approval mode 156
- approval requests
 - alert for 610
 - analyzing 572, 584
 - automatic resolution email 580
 - customizing the notifier interface 587
 - enabling in Windows 566
 - how users submit 568
 - in blocked file notifiers 566
 - request details page 583
 - responding to 569
 - viewing in Console 569
- approvals
 - adding (by file) 304
 - by policy 308
 - custom 308
 - defined 261
 - local 289
 - removing 306
- approve on Enforcement Level transition (policy setting) 190
- approved (local state detail) 251
- approved as installer (local state detail) 251
- approved as top-level installer (local state details) 251
- Approved Files view 226
- approved not persisted (local state detail) 251
- approving devices 365
- approving files
 - automating using event rules 520
 - by automatic updaters 266
 - by custom rule 406
 - by file reputation 326
 - by hash 310
 - by hash for MSI files 303
 - by importing a hash list 310
 - by local approval mode 294
 - by local approval on Enforcement Level change 290
 - by publisher approval (manual) 281
 - by publisher reputation 326
 - by trusted directory 271
 - by trusted user or group 279
 - from a deployment server 271
 - overview 261
 - printer driver updates 267
 - removing approvals 306
 - removing local approval 292
- approving publishers
 - by reputation 330
 - manual 281
- archives
 - event 607
 - in trusted directories 272
- ArcSight integration

- specifying CEF as Syslog format 738
- assessment
 - threat and trust ratings 570
- B**
- backups
 - App Control database 750
 - backup missed alerts 609
 - restoring from 753
- banned by hash (local state detail) 251
- banned by hash report-only (local state detail) 251
- Banned Files view 226
- banned state 250
- banning files
 - automating using event rules 520
 - by hash 264, 310
 - by hash for MSI files 303
 - by importing a hash list 310
 - by name 264
 - by policy 304, 308
 - by publisher 282
 - from the Software Rules page 304
 - overview 48, 263
 - removing bans 306
- banning publishers 282
- bans
 - creating 263
 - custom 308
 - file name 48, 264
 - hash 48, 264
 - removing 306
 - report only 251, 264, 308
 - terminating banned processes 312
 - verifying before deployment 303
- baseline drift 634
 - adding results to a snapshot 644
 - alert for 612
 - by file category 639
 - creating and editing reports 646
 - displaying in dashboards 656
 - remediation of 643
 - snapshots for 653
 - viewing report results 638
 - viewing the list of reports 637
- block banned file hashes (policy setting) 190
- block banned file names (policy setting) 189
- block files with banned publishers or certificates (policy setting) 190
- block network executables (policy setting) 190
- block unanalyzed scripts and executables (policy setting) 189
- block unapproved executables (policy setting) 189
- block unapproved scripts (policy setting) 189
- block-and-ask. See Medium Enforcement Level
- blocked file notifiers. See notifiers
- blocking files 199
 - by custom rule 393
 - by file ban 263
 - by publisher 282
 - by script rule 379
 - on devices 361
- browsers
 - certificates warnings in 53
 - supported 53
- C**
- cache consistency check 169
- cache, AD
 - clearing 130
- cached events 602
- Carbon Black EDR
 - and Computer Details 164
 - API token 769
 - computers with sensor 156
 - integrating with App Control 768
 - sensor status 164
 - updates for macOS (OS X) sensors 267
- Carbon Black File Reputation
 - alert when unavailable 610
 - defined 46
 - enabling and disabling 763
 - file category 240
 - file category information from 226
 - file trust rating 46, 239
 - proxy settings 763
 - synchronization with 767
 - threat level 239
 - TLS 1.2 requirement 763
 - using a proxy server for 766
 - view file data from 765

- Carbon Black Technical Support 11
- Categorized Files view 226
- category. See file category
- CEF. See ArcSight integration
- certificate rules 337
- Certificates page 339
- certificates, App Control
 - for agent-server communication 739
 - using SAN in 741
- certificates, file-signing
 - alerts for 347
 - algorithm options 288
 - and publisher approvals 286
 - approval configuration options 349
 - approving and banning 348, 351
 - approving by 281
 - certificate details fields 341
 - certificate global state 352
 - certificate path 343
 - configuring approvals by 287, 748
 - cosigner 350
 - countersignature options 288
 - detached 281, 347, 350
 - discovery and control of 337
 - effect on global file state 359
 - embedded 350
 - enabling/disabling bans by policy 358
 - events for 348
 - expired 288
 - feature overview 338
 - finding child certificates 345
 - finding events for a certificate 345
 - finding files signed by 345
 - for publisher approvals 281
 - in external views 348
 - information in file details 347
 - key length options 288
 - other rules and certificate global state 358
 - path differences 350
 - path position of 341, 350
 - policy setting for 190
 - revocation checks 289
 - table of 339
 - types 350
 - viewing details 344
 - viewing for a publisher 346
- certificates, VMware Carbon Black App Control
 - Control
 - and console login 53
 - Check Point
 - analyzing files with 883
 - enabling file analysis with 864
 - enabling integration with 857
 - proxy settings or 866
 - CIM
 - mappings for Splunk 924
 - CL. See configuration list
 - CLI management privileges 727
 - and command line reporting 602
 - cloned computers
 - cleanup of 217, 219
 - deleting 217, 219
 - file inventory choices 217
 - managing 207
 - server backlog for 215
 - command lines
 - reporting in events 601
 - company
 - viewing files by 226
 - Computer Details page 159
 - computer security alert 611
 - computers
 - adding 173
 - assigning policies 129
 - changing policies 170
 - cloned 207
 - connected (viewing) 156
 - deleting 173, 747
 - details about 159
 - disconnected (viewing) 156
 - duplicate registrations 593
 - health check for 161
 - in Local Approval (viewing) 156
 - initializing 119
 - installing agent on 136
 - placing in local approval mode 296, 297
 - remote reboot of 168
 - requiring upgrade, (viewing) 156
 - restoring from local approval mode 297
 - template computers 207
 - timed Enforcement Level override for 297
 - uninstalling agent on 154
 - viewing AD details about 130
 - viewing connection status 156

- virtual machines 207
 - with/without specified files 229
 - Computers page 146, 156
 - configuration list
 - current (for server) 156
 - file state and 241
 - for an agent computer 163
 - confirm navigation dialog
 - enabling/disabling 83
 - connected computers, viewing 156
 - connected Enforcement Level 184, 202
 - connection status (agent) 161
 - Connector 843
 - console account permissions for 867
 - enabling Check Point integration 857
 - enabling Palo Alto Networks integration 846
 - connector. See Connector.
 - Console
 - 2FA for 769
 - browser certificate for 53
 - creating accounts 89
 - logging in 53
 - logging out 55
 - supported browsers 53
 - using 52
 - using an idP with 769
 - using SAML with 769
 - console menu 60
 - console menu bar 60
 - console, App Control. See App Control Console
 - console. See Console
 - control mode 183
 - enabling for a policy 181
 - licenses for 759
 - overview 50
 - cosigner certificates 350
 - countersignatures (for certificates) 288
 - countersigner certificates, see cosigner certificates
 - cryptomining protection
 - rapid config for 508
 - CSC temporary files 267
 - custom rules
 - do not track example 446
 - expert interface 488
 - exporting and importing 431
 - in visibility mode 396
 - overview 393
 - trusted paths 441
 - customer support 11
- ## D
- dashboards
 - adding portlets to 696
 - baseline drift portlets in 656
 - changing appearance of 688
 - changing color of 690
 - changing width of 690
 - copying 695
 - creating 691
 - editing 691, 695
 - home page 57
 - layout of 689
 - managing 691
 - portlets on 683
 - sharing with other users 693
 - system 686
 - viewing 681, 686
 - data analytics 906
 - preparing for 907
 - data directory
 - for macOS agents 142
 - for Windows 138
 - on Linux agents 145
 - databas, App Control
 - database limit alert 609
 - database, App Control
 - address 726
 - authorization type 726
 - configuration information 724
 - events in 732
 - external 734
 - restoring 753
 - schema version 725
 - size 726
 - verification failed alert 609
 - views via live inventory SDK 815
 - database,App Control
 - unique files 47
 - debug level
 - for an agent computer 163
 - default policy 192
 - default starting page 83
 - deleted computers 173

- deleted file state 250
 - deleted files
 - searching for 718
 - viewing 226
 - deleting files 314
 - automating with event rules 321
 - delivery optimization
 - rapid config for 508
 - detached certificates 347, 350
 - detection, threat 660
 - device paths, in software rules 409
 - devices
 - acknowledging 367, 368
 - all devices on computers 375
 - approving and banning 360
 - control in App Control 362
 - device catalog 371
 - managing 360, 365
 - managing by model 366
 - managing individual devices 371
 - per-policy control 362
 - policy settings 363
 - rules for 361
 - DFS
 - and Windows 2003/XP 139
 - diagnostic files 890, 902
 - viewing 892, 902
 - directory policies. See custom rules
 - disabled mode (agent) 183, 199
 - disconnected computers
 - and file searches 713, 718
 - and policy deletion 121
 - changing Enforcement Level 297
 - deleting 747
 - during lockdown 203
 - timed deletion 747
 - viewing 156
 - disconnected Enforcement Level 184, 202
 - display preferences 83
 - DMG files
 - for installing Mac agents 142
 - domain controllers
 - rapid config for 508
 - doppelganger protection
 - rapid config for 508
 - Download Agent Packages page 135
 - downloading
 - agent installers 134
 - downloading agent installers 134
 - downloading data to CSV files 80
 - drift reports. See baseline drift
 - duplicate computer registrations 593
 - dynamic code execution (memory rule) 480
 - dynamic tables 66
 - downloading data from 80
 - filtering results 73
 - hiding columns 76
 - Saved Views in 78
 - showing columns 76
- ## E
- email
 - address in approval request 568
 - address in SSL certificate 741
 - for alerts 608, 754
 - for approval requests 580
 - generated by block notifier link 550
 - login account user address 96
 - embedded certificates 350
 - emergency lockdown 204
 - Enforcement Level
 - and policy settings 180
 - changing 201
 - connected 184, 202
 - defined 5, 197
 - disconnected 184, 202
 - effect on policy enforcement 199
 - file blocking for active policy settings 199
 - High (Block Unapproved) 197
 - local approval 201
 - locking down all computers 203
 - Low (Monitor Unapproved) 198
 - Medium (Prompt Unapproved) 198
 - None (Disabled) 199
 - None (Visibility) 198
 - out of date on agent 157
 - overview 50
 - setting for new policies 184
 - timed overrides of 297
 - event rules 520
 - copying settings from 528
 - creating alerts for 612
 - disabling 523
 - enabling 523
 - ranking of 532

- events
 - agent health check 166
 - archives of 607
 - cached 602
 - creating alerts for 612
 - creating reports of 599
 - editing reports of 600
 - events page 595
 - external logging 734
 - home page summary 591
 - log files 732
 - logging of 732
 - overview 590
 - reporting command lines in 601
 - saved views of 591, 599
 - Syslog message severity 598
 - system health 813
 - threat detection 672
 - triggering actions with 520
 - types 591
 - events integration
 - See the separate App Control Events Guide
 - exceptions
 - for indicator sets 666
 - executables
 - advanced policy settings for 189
 - defined 44
 - Existing Files view 226
 - expert rules 488
 - expired certificates
 - and certificate approvals 349
 - and publisher approvals 288
 - export directory for external analytics 912
 - exporting data 80
 - data analytics 906
 - exporting rules to another server 431
 - external analytics
 - accessing external tools from App Control console 917
 - creating a rule to ignore logs 914
 - data format for 908
 - enabling connection for 914
 - enabling in App Control Console 910
 - export directory for 912
 - exported files 908
 - installing Splunk app for App Control 916
 - installing Splunk Universal Forwarder 916
 - viewing App Control data 917
 - external event logging 734
 - external notifications 867
 - event rules for 520
 - trimming 868
 - external views App Control database 815
- ## F
- file and path rules enforcement (policy setting) 190
 - file and path rules. See custom rules
 - file bans. See bans
 - File Catalog tab 249
 - file category
 - defined 240
 - drift by 639
 - file creation control 393
 - file details 237
 - File Details page 310, 320, 321
 - file execution control 393
 - file extensions
 - script rules and 379
 - File Group Details page 247
 - file groups
 - and initialized files 228
 - overview 234
 - viewing files in 247
 - file hash bans 264
 - File Instance Details page 242
 - initiating Find Files from 713
 - file instances
 - file name 244
 - path for 244
 - file integrity control 393
 - file inventory 41
 - excluding MS support files 230
 - of cloned computers 217
 - file name bans 48
 - file rules
 - approvals 304
 - bans 304
 - removing 306
 - file state 47, 249
 - and certificate global state 359
 - approved 249
 - banned 249

- banned (local) 250
 - defined 6
 - deleted 250
 - flags affecting 237, 249
 - global 237
 - instance states 250
 - local 250
 - local state details 251
 - locally approved 250
 - unapproved 250
 - file state reason 237
 - file tracking
 - and alerts 608
 - disabling for a path 446
 - enable/disable by policy 185
 - excluding MS support files 230
 - using baseline drift 635
 - files
 - acknowledging 241
 - analyzing with third-party devices 843
 - App Control database 47
 - approving. *See* approving files
 - banning. *See* banning files
 - baseline drift of 635
 - blocked file alerts 612
 - blocking 199
 - blocking by custom rule 393
 - blocking by device rule 361
 - blocking by script rule 379
 - categories of 226
 - deleting 314
 - diagnostic 890, 902
 - executable 44
 - existing 226
 - file groups 247
 - finding 713
 - finding computers with/without specified files 229
 - first-seen name 237
 - including deleted files in a search 718
 - initializing 119
 - installing on a locked-down computer 294
 - live inventory of 41
 - local approval 289
 - locating executables on computers 715
 - malicious 226
 - marking as installer 300
 - marking as not installer 300
 - metering executions 630
 - monitoring specific executions 629
 - on deleted computers 719
 - on disconnected computers 718
 - path for first-seen 237
 - prevalence alerts 612
 - propagation alerts 612
 - reputation 327
 - show individual files 308
 - snapshots of 653
 - threat level for 239
 - tracked in App Control 44
 - tracking drift 634
 - trust rating for 239
 - uploading from agents 895
 - viewing Carbon Black File Reputation for 765
 - viewing removed 226
 - Files on Computers tab 48, 250
 - filtering
 - data in portlets 708
 - table data in portlets 706
 - table results 73
 - Find Files page
 - overview 714
 - Saved Views in 720
 - finding computers
 - with/without specified files 229
 - finding files
 - case sensitivity 715
 - computers with/without specified files 229
 - from Computer Details page 166
 - from Find Files page 713
 - from Home Page 58
 - on computers in a policy 196
 - overview 713
 - special cases 718
 - using filters in a search 716
 - viewing all unapproved files in a policy 201
 - flags (file state) 249
 - fuzzy hashing 239
- ## G
- global state 237
 - group information (file details) 240
 - groups

- trusted for installation 279
- groups (file details) 241
- H**
- hashes
 - approving 310
 - approving a list of 310
 - banning 48, 310
 - banning a list of 310
 - fuzzy hashing 239
 - identifying unknown 763
 - MD5 239
 - SHA-1 239
 - SHA-256 239
- health check
 - for agents 161, 166
- health checks
 - system health 806
- help
 - for portlets 685
 - user guide 86
- hiding table columns 76
- High (Block Unapproved) Enforcement Level 197
- High Enforcement Level
 - installing software on computers in 294
 - switching to 201
- Home page 60
- home page 57
 - changing appearance of 688
 - changing default for new users 697
 - editing 691
 - resetting to default 697
- HP ArcSight. See ArcSight integration
- I**
- identity providers (IdPs)
 - for console login 769
- importing rules from another server 431
- indicator set details 664
- indicator sets
 - enabling and disabling 664
 - exceptions to 666
 - for threat detection 662
 - updates to 670
- information button
 - for portlets 685
 - on Active Directory Policy Mappings
 - page 115, 128
- initialization 5
 - and local approval 289
 - of cloned computers 217
 - of computers 119
 - status of 164
- initialized files
 - overview 228
 - viewing for one computer 248
- installed programs 234
- Installed Programs view 226
- installer (override) file flag 250
- installer file flag 250
- installer versions (agent) 133
- installers
 - agent 136
 - and file groups 234
 - defined 300
 - files approved as 251
 - files identified as 237
 - files marked as 242
 - in trusted directories 272
 - marking file as 300
 - recognized in trusted directories 272
 - top level 251
- installing
 - Agent 136
 - App Control Server. See App Control Installation Guide
- IPv6
 - in server address 725
- J**
- Java
 - script rules for 380
 - updater for 267
- javascript
 - script rule limitations for 379
- JSON
 - data export format 908
- justification (for user-initiated approvals)
 - alert for 610
- justifications (for user-initiated approvals)
 - alert for 610
 - customizing the notifier interface 587
 - details page 583
 - enabling in Windows 566
 - how users submit 568

- in blocked-file notifiers 566
 - responding to 580
 - viewing in Console 569
- K**
- kernel memory access (memory rule) 480
- kernels, Linux. See separate Operating Environment Requirements guide
- key length (for certificates) 288
- L**
- LEEF. See QRadar integration
- licenses, App Control 759
 - and local approval mode 294
 - Carbon Black File Reputation 763
 - for connectors 845
 - for file uploads 896
 - managing 759
 - viewing limits and usage 759
- licenses, App Control
 - adding 761
- Linux computers
 - agent data directory on 145
 - installing agent on 144
 - uninstalling agent from 155
- linux hardening
 - rapid config for 508
- live inventory
 - and baseline drift 635
 - and executable files 44
 - and finding files 713
 - database views of 815
 - defined 41
 - SDK 815
- local approval 289
 - of all unapproved files on a computer 294
 - of files 289
 - of one file 292
 - removing 292
- local approval mode 294
 - alert for 610
 - and disconnected computers 297
 - and online computers 296, 297
 - restoring computers to original policies 297
 - setting time-duration alerts 612
 - timed Enforcement Level changes 298
 - viewing computers in 156
- local file state 250
- local file state details 251
- locally approved (local state detail) 251
- locally approved auto (local state detail) 251
- locally approved state 250
- lockdown
 - Enforcement Level for 197
 - locking down all computers 203
 - restoring after 204
- lockdown. See also High Enforcement Level
- log files
 - managing 732
- logging in 53
- logging out 55
- login accounts 89
 - creating new groups 102
 - defined 6
 - deleting 100
 - disabling 101
 - groups 101
 - permissions for Connector 867
 - role-based access 102
 - setting preferences for 83
 - using AD accounts 92
- logo
 - specifying for notifier 560
- Low (Monitor Unapproved) Enforcement Level 198
 - file execution warnings in 201
 - switching to High Enforcement from 201
- M**
- Mac computers
 - agent data directory for 142
 - App Store updater 266
 - blocked file notifiers on 542
 - installing agent on 142
 - native updater support for 266
 - submitting approval requests from 569
 - Symantec Endpoint Protection updater for 268
 - tray icon on 543
 - uninstalling agent from 155
- Mac System Updates 267
- macOS. See Mac computers

macros, in software rules 410
 malicious files
 alerts for 610
 how specified
 Malicious Files view 226
 Mark as installer/not installer 300
 Medium (Prompt Unapproved) Enforcement Level 198
 memory rules 470
 copying 475
 editing notifier message for 472
 expert interface 488
 exporting and importing 431
 operating system restrictions 471
 parameters of 476
 viewing associated events 471
 memory rules enforcement (policy setting) 190
 meters (software execution) 629
 creating 630
 Microsoft .NET updates 267
 Microsoft Exchange Server
 rapid config for 508
 Microsoft Office
 rapid config for 508
 Microsoft Office Click-to-Run updates 267
 Microsoft SCCM
 rapid config for 508
 Microsoft SQL Server
 rapid config for 508
 mimikatz protection
 rapid config for 509
 modes
 overview 50
 setting for policies 183
 monitor. See Low Enforcement Level
 MSI files
 and trusted directories 272
 for installing Windows agents 138
 hash approvals and bans for 303

N

network security devices
 notifications from 867
 new certificate alerts 347
 New Unapproved Files view 226
 not installer (override) file flag 250
 notifications

 external 867
 notifiers for blocked files
 conditional messages in 554
 configuring 548
 customizing the logo for 560
 defined 194
 disabling 550, 562
 editing 548, 552
 editing by policy setting 546
 editing the source line in 559
 enabling approval requests in 566
 for terminal servers 564
 for XenApp 564
 history window for Mac 543
 information links in 550
 on Mac computers 542
 timeouts for on-screen display 550
 using tags in 552
 NT authorization
 for database server 726

O

object previews
 in table data 82
 offline computers. See disconnected computers
 online computers. See connected computers
 online help 86
 operating strategies 51
 OS X. See Mac computers

P

packages
 by publisher/company 226
 Mac .pkg files 228
 trusted 226, 276
 Palo Alto Networks
 access to console from Console 879
 enabling App Control integration with 846
 file analysis with WildFire 850
 Integration with App Control 843
 notifications from 867
 Parity Knowledge Service. See Carbon Black File Reputation
 passwords
 CLI management 727
 console 96

- console (changing) 98
- path
 - certificate 343
 - first-seen file 237
 - trusted 441
- path position, certificate 350
- path position, for certificates 341
- path rules. See custom rules
- pending files. See unapproved files
- performance optimization
 - custom rules for 393
- policies
 - AD mapping 122
 - creating 181
 - default 192
 - defined 5, 49
 - deleting 205
 - disabling enforcement 183
 - Enforcement Level for 184
 - for uninstalling an agent 154
 - mode choices 183
 - moving computers between 170
 - related views menu 196
 - setting alerts for 612
 - template 192
 - templates for 184
 - viewing unapproved files in a policy 201
 - when assigned 129
- Policies page 133, 182
- policy settings
 - and Enforcement Level 180
 - blocking for different Enforcement Levels 199
 - creating a template policy for 192
 - device control 368
 - editing 194
 - enable/disable file tracking 185
 - local approval of unapproved files on Enforcement Level change 290
 - notifiers for 546
 - options for 187
 - removable device 363
- policy specific states (file details) 240
- policy status 157
- portlets 683
 - adding to a dashboard 696
 - baseline drift 656
 - creating 701
 - deleting 700
 - editing 86, 700
 - filtering data in 708
 - filtering table data in 706
 - moving on dashboard 691
- potential risk files
 - alerts for 609
 - Carbon Black File Reputation information about 239
- powershell protection
 - rapid config for 509
- preferences, console user 83
- prevalence of files on computers 237
- printer driver updates 267
- prioritizing agent updates 166
- privileges, login account
 - and AD accounts 92
 - customizing 101
- process protection. See memory rules
- processes
 - in custom rules 423
 - in memory rules 481
 - in registry rules 460
 - in script rules 379
 - terminating if banned 312
- promote (treat as installer)
 - in custom rules 404
 - notifier option 540
- promoted process 423
- propagating files
 - setting alerts for 612
- proxy settings
 - Carbon Black File Reputation 763
 - Check Point 866
 - WildFire 851
- publishers
 - acknowledging 281, 284
 - and global file state 237
 - approving 281
 - approving by reputation 327
 - banning 282
 - certificates for 346
 - detached publisher state 244
 - in file details 238, 245, 259
 - policy setting for 190
 - publisher details 253
 - publisher state 238
 - viewing files by 226

Q

- Q1Labs. See QRadar integration
- QRadar integration
 - specifying LEEF as Syslog format 738

R

- ransomware
 - protection from 509
- rapid configs
 - alert when modified 519, 610
 - automatic updates of 518
 - enabling 511
 - overview 506
 - specifying notifiers for 516
 - table of 507
- reboot
 - of agent computers 168
- Red Hat Prelinking 268
- refresh page 67
- registration of agents 129
- registry rules 449
 - copying 454
 - editing notifier message for 455, 458, 461
 - enabling by policy 190
 - expert interface 488
 - exporting and importing 431
 - parameters of 455
 - process menu options 461
 - write actions 458
- registry rules enforcement (policy setting) 190
- removable devices. See devices
- Removed Files view 226
- Report Process Create rule
 - and command line reporting 601
- reporting problems 11
- Report-Only (for file bans) 264
- report-only ban flag 250
- reputation approvals 326
- reputation services. See Carbon Black File Reputation
- reputation-based rules 326, 327
- resizing table columns 67
- Restore page 204
- restoring
 - computers in emergency lockdown 204
 - database 753

- local-approval computers to policies 297
- revocation checks (for certificates) 289
- revoked certificate alerts 347
- role-based access. See login accounts
- rules
 - copying 454, 475
 - exporting and importing 431
- rules files
 - uploading to the server 131

S

- SAML
 - for console login 769
- SAN (subject alternative name)
 - in certificate definition 741
- Saved Views
 - creating 79
 - discarding changes to 80
 - overview 78
- script processors 379
 - restricting 509
- script rules 379
 - javascript limitations 379
 - rescanning after added 380
 - yara used in 382
- scripts
 - blocking unapproved 189
 - custom definitions of 378
 - defined 379
 - editing rules for 378
- SecCon. See Enforcement Level
- security domain
 - for AD integration 727
- self-protection. See tamper protection
- self-service approvals
 - rapid config for 509
- Server
 - restoring 753
 - status information 724
- server backlog
 - for cloned computers 215
- shared drives
 - file execution setting for 190
- shortcut links 83
- Show deleted files box
 - in Find Files results 718
- Show Individual Files box 308
- show/hide columns 67

- show/hide filters 67
 - show/hide snapshots 67
 - showing table columns 76
 - SIEM integration 734
 - See also the separate App Control Events Guide
 - silent blocks
 - in memory rules 479
 - silent blocks. See also notifiers for blocked files
 - snapshots
 - adding drift results to 644
 - creating 653
 - editing 655
 - for baseline drift reports 653
 - showing panel 67
 - software approvals. See approvals
 - software bans. See banning files and bans.
 - software metering 629
 - Software Meters page 630
 - Software Rules page 274
 - software updates
 - automatic updater support 266
 - Splunk
 - CIM mappings for App Control data 924
 - enabling data collection by 915
 - enabling data export to 910
 - installing App Control app 916
 - installing Universal Forwarder on App Control Server 916
 - viewing App Control data in 918
 - SQL Server
 - authorization for 726
 - for external event logging 734
 - SSL security
 - configuring 739
 - starting page, changing 83
 - Symantec Endpoint Protection (SEP) for Mac updater 268
 - synchronization
 - agent-server 163, 164
 - and template computers 209, 216
 - with Carbon Black File Reputation 767
 - Syslog
 - enabling for events 738
 - integrating with ArcSight 738
 - integrating with QRadar 738
 - message severity 598
 - System 812
 - system backups 750
 - System dashboard 686
 - System Health page 806
- ## T
- tags
 - for alert messages 617
 - for approval requests 566
 - for computer identification 162
 - for customizing notifiers 552
 - for expert rules 499
 - tamper protection
 - for agents (policy setting) 190, 194
 - technical support 11
 - template computers
 - converting to regular computer 221
 - creating 209
 - deleting 217
 - editing 211, 216
 - viewing table of 210
 - template policy 192
 - templates
 - for virtual machines 207
 - terminate processes for banned images (policy setting) 190
 - terminating processes with banned images (policy setting) 312
 - TGZ files
 - for installing Linux agents 144
 - threat detection 660
 - alerts for 676
 - and App Control upgrades 662
 - events for 672
 - indicator sets for 662
 - monitoring reports 671
 - responding to 677
 - suspicious files 675
 - updates to 670
 - threat level, from Carbon Black File Reputation 239
 - timeouts for notifier display 550
 - trust rating
 - for files 46, 327
 - for publishers 328
 - from Carbon Black File Reputation 239
 - trusted directories 271
 - archive files in 272

- installer files in 272
- packages recognized by App Control 272
- trusted groups 279
- trusted package
 - noted in file details 242
- Trusted Packages view 226, 276
- trusted paths 441
- trusted users 279
- two-factor authentication
 - for console login 769

U

- unapproved (local state detail) 252
- unapproved files
 - approving on Enforcement Level change 290
 - executables (blocking by policy) 189
 - finding all on computers in a policy 201
 - local state 250
 - local state detail 252
 - locally approving on a computer 294
 - scripts (blocking by policy) 189
 - unapproved (persisted) 252
 - viewing new unapproved 226
- unapproved persisted (local state detail) 252
- unapproved scripts (policy setting) 189
- unapproved state 250
- Unified Management configuration 769
- uninstalling
 - agent software 154
- updaters
 - alert when modified 271, 610
 - enabling 266
 - for Carbon Black EDR sensors on macOS (OS X) 267
- upgrade status
 - agents 153
- upgrading Agent 147
 - manual upgrades 149
- uploading agent installers 131
- uploading agent rules files to the server 131
- uploading files from agents 895
 - automating using event rules 520
 - changing upload location 904
 - non-ANSI file names 903

- uploads
 - of diagnostic files 890, 902
- URLs
 - for downloading agent installers 134
 - in notifier link 550
- user assistance 86
- user passwords
 - Console (changing) 83
- user preferences 83
- users, App Control Console. See login accounts
- users, trusted. See trusted users

V

- version number
 - agent config list 163
 - Server 56
 - server config list 156
- virtual machines
 - identifying in computer details 162
 - managing 207
- virtual platform
 - in computer details 162, 164
- virtualization
 - session 564
- Visibility and Control mode. See control mode
- visibility mode 183, 198
 - and custom rules 396
 - licenses for 759
- Visibility Only mode. See visibility mode
- Visual Studio
 - rapid config for 509
- VMware
 - identifying in computer details 162, 164
 - managing clones 207
- VMware Carbon Black App Control Server
 - version number 56

W

- warnings
 - about non-upgraded agents 147
 - file execution 201
 - license limit 761
- Watchlist
 - Carbon Black EDR 768
- wildcards, in software rules 409
- WildFire

- analyzing files with 883
- integrating with App Control 850
- multiple notifications from 874
- proxy settings 851
- WIM files
 - enabling trusted directory analysis for 272
- Windows 2000 domain controllers and AD integration 727
- Windows App Store
 - rapid config for 509
- Windows computers
 - agent data directory for 138
 - enabling file approval requests for 566
 - installing agent on 138
 - manual agent upgrades on 149
 - submitting approval requests from 568
 - uninstalling agent from 154
- Windows Defender updates 268
- Windows Hardening
 - rapid config for 509
- Windows Installer Transform files (not supported) 138
- Windows updates 175, 268
- WMI Protection
 - rapid config for 509

Y

- yara
 - in script rules 382
 - rules out of date 158
 - yara tags in rules 502