Carbon Black.

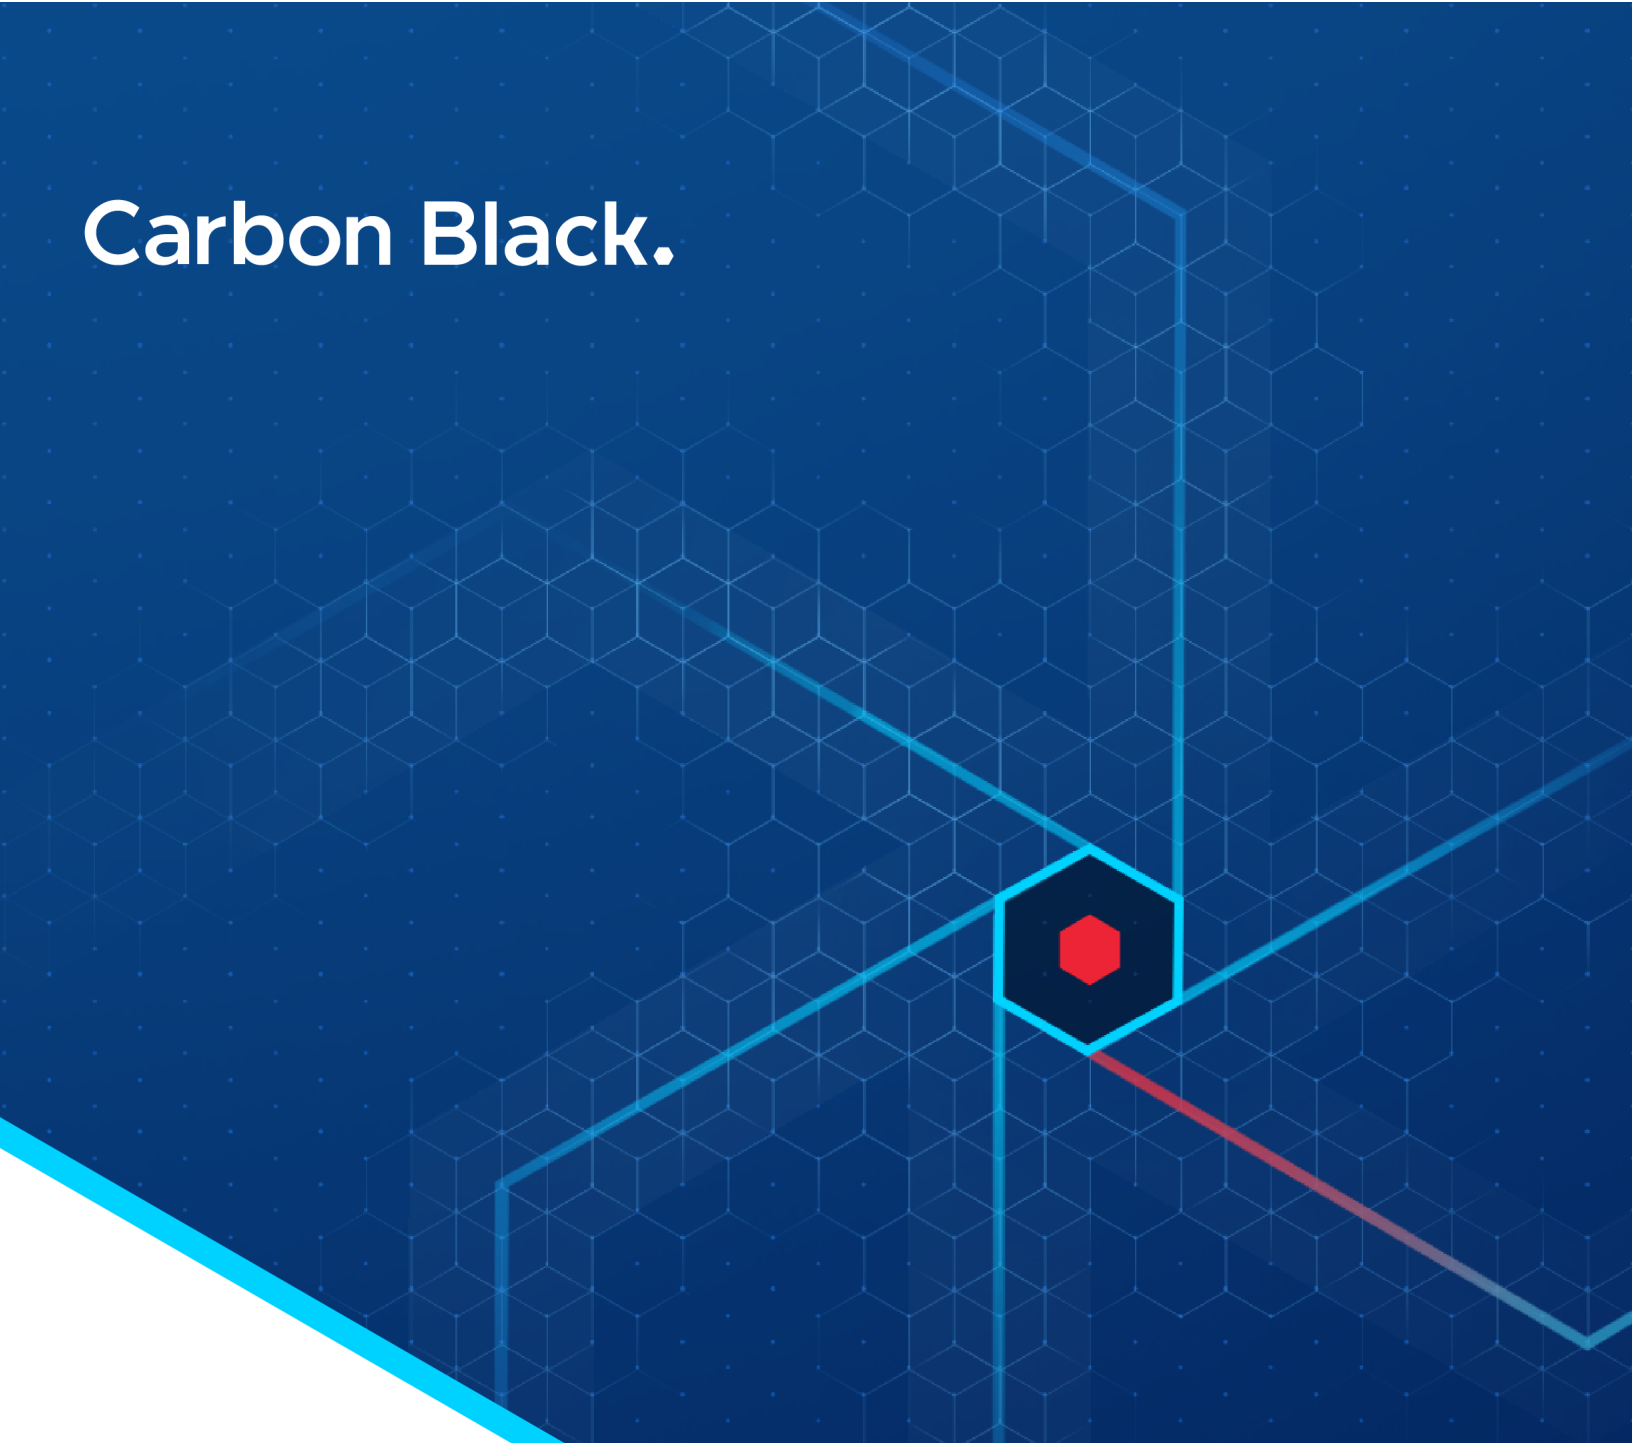

CB Response User Guide

Server/Cloud Version: 7.1

Document Date: April 2020

Copyrights and Notices

Copyright ©2011-2020 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.

This document is for use by authorized licensees of this product. It contains the confidential and proprietary information of VMware, Inc. and may be used by authorized licensees solely in accordance with the license agreement governing its use. This document may not be reproduced, retransmitted, or redistributed, in whole or in part, without the written permission of VMware. VMware disclaims all liability for the unauthorized use of the information contained in this document and makes no representations or warranties with respect to its accuracy or completeness. Users are responsible for compliance with all laws, rules, regulations, ordinances and codes in connection with the use of VMware products

THERE IS NO WARRANTY FOR THE PROGRAM, TO THE EXTENT PERMITTED BY APPLICABLE LAW EXCEPT WHEN OTHERWISE STATED IN WRITING BY VMWARE. 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 CB Response software product:

- Antlr python runtime - Copyright (c) 2010 Terence Parr
- Backbone routefilter - Copyright (c) 2012 Boaz Sender
- Backbone Upload - Copyright (c) 2014 Joe Vu, Homeslice Solutions
- Backbone Validation - Copyright (c) 2014 Thomas Pedersen, http://thedersen.com
- Backbone.js - Copyright (c) 2010–2014 Jeremy Ashkenas, DocumentCloud
- Beautifulsoup - Copyright (c) 2004–2015 Leonard Richardson
- Canvas2Image - Copyright (c) 2011 Tommy-Carlos Williams (http://github.com/devgeeks)
- Code Mirror - Copyright (c) 2014 by Marijn Haverbeke marijnh@gmail.com and others
- D3js - Copyright 2013 Mike Bostock. All rights reserved
- FileSaver - Copyright (c) 2011 Eli Grey.
- Font-Awesome - Copyright Font Awesome by Dave Gandy - http://fontawesome.io
- Fontello - Copyright (c) 2011 by Vitaly Puzrin
- Freewall - Copyright (c) 2013 Minh Nguyen.
- FullCalendar - Copyright (c) 2013 Adam Shaw
- Gridster - Copyright (c) 2012 Ducksboard
- Heredis - Copyright (c) 2009–2011, Salvatore Sanfilippo and Copyright (c) 2010–2011, Pieter Noordhuis
- Java memcached client - Copyright (c) 2006–2009 Dustin Sallings and Copyright (c) 2009–2011 Couchbase, Inc.
- Javascript Digest Auth - Copyright (c) Maricn Michalski (http://marcin-michalski.pl)
- Javascript marked - Copyright (c) 2011–2014, Christopher Jeffrey (https://github.com/chjj/)
- Javascript md5 - Copyright (c) 1998 - 2009, Paul Johnston & Contributors All rights reserved.
- Javascript modernizr - Copyright (c) 2009 - 2013 Modernizr
- Javascript zip - Copyright (c) 2013 Gildas Lormeau. All rights reserved.
- Jedis - Copyright (c) 2010 Jonathan Leibiusky
- Jmousewheel - Copyright (c) 2013 Brandon Aaron (http://brandon.aaron.sh)
- Joyride - Copyright (c) 1998 - 2014 ZURB, Inc. All rights reserved.
- JQuery - Copyright (c) 2014 The jQuery Foundation.
- JQuery cookie - Copyright (c) 2013 Klaus Hartl
- JQuery flot - Copyright (c) 2007–2014 IOLA and Ole Laursen
- JQuery Foundation - Copyright (c) 2013–2014 ZURB, inc.
- JQuery placeholder - Copyright (c) Mathias Bynens http://mathiasbynens.be/
- JQuery sortable - Copyright (c) 2012, Ali Farhadi
- Jquery sparkline - Copyright (c) 2009–2012 Splunck, Inc.
- JQuery spin - Copyright (c) 2011–2014 Felix Gnass [fgnass at neteye dot de]
- JQuery tablesorter - Copyright (c) Christian Bach.
- JQuery timepicker - Copyright (c) Jon Thornton, thornton.jon@gmail.com, https://github.com/jonthornton
- JQuery traffic cop - Copyright (c) Jim Cowart

- JQuery UI - Copyright (c) 2014 jQuery Foundation and other contributors
- jScrollPane - Copyright (c) 2010 Kelvin Luck
- Libcurl - Copyright (c) 1996 - 2014, Daniel Stenberg, daniel@haxx.se.
- libfreeimage.a - FreeImage open source image library.
- Meld3 - Supervisor is Copyright (c) 2006–2015 Agendaless Consulting and Contributors.
- moment.js - Copyright (c) 2011–2014 Tim Wood, Iskren Chernev, Moment.js contributors
- MonthDelta - Copyright (c) 2009–2012 Jess Austin
- Mwheelintent.js - Copyright (c) 2010 Kelvin Luck
- nginx - Copyright (c) 2002–2014 Igor Sysoev and Copyright (c) 2011–2014 Nginx, Inc.
- OpenSSL - Copyright (c) 1998–2011 The OpenSSL Project. All rights reserved.
- PostgreSQL - Portions Copyright (c) 1996–2014, The PostgreSQL Global Development Group and Portions Copyright (c) 1994, The Regents of the University of California
- PostgreSQL JDBC drivers - Copyright (c) 1997–2011 PostgreSQL Global Development Group
- Protocol Buffers - Copyright (c) 2008, Google Inc.
- pyperformance - Copyright 2014 Omer Gertel
- Pyrabbit - Copyright (c) 2011 Brian K. Jones
- Python decorator - Copyright (c) 2008, Michele Simionato
- Python flask - Copyright (c) 2014 by Armin Ronacher and contributors
- Python gevent - Copyright Denis Bilenko and the contributors, http://www.gevent.org
- Python gunicorn - Copyright 2009–2013 (c) Benoit Chesneau benoitc@e-engura.org and Copyright 2009–2013 (c) Paul J. Davis paul.joseph.davis@gmail.com
- Python haigha - Copyright (c) 2011–2014, Agora Games, LLC All rights reserved.
- Python hiredis - Copyright (c) 2011, Pieter Noordhuis
- Python html5 library - Copyright (c) 2006–2013 James Graham and other contributors
- Python Jinja - Copyright (c) 2009 by the Jinja Team
- Python kombu - Copyright (c) 2015–2016 Ask Solem & contributors. All rights reserved.
- Python Markdown - Copyright 2007, 2008 The Python Markdown Project
- Python netaddr - Copyright (c) 2008 by David P. D. Moss. All rights reserved.
- Python ordereddict - Copyright (c) Raymond Hettinger on Wed, 18 Mar 2009
- Python psutil - Copyright (c) 2009, Jay Loden, Dave Daeschler, Giampaolo Rodola'
- Python psycogreen - Copyright (c) 2010–2012, Daniele Varrazzo daniele.varrazzo@gmail.com
- Python redis - Copyright (c) 2012 Andy McCurdy
- Python Seasurf - Copyright (c) 2011 by Max Countryman.
- Python simplejson - Copyright (c) 2006 Bob Ippolito
- Python sqlalchemy - Copyright (c) 2005–2014 Michael Bayer and contributors. SQLAlchemy is a trademark of Michael Bayer.
- Python sqlalchemy-migrate - Copyright (c) 2009 Evan Rosson, Jan Dittberner, Domen Kozar
- Python tempita - Copyright (c) 2008 Ian Bicking and Contributors
- Python urllib3 - Copyright (c) 2012 Andy McCurdy
- Python werkzeug - Copyright (c) 2013 by the Werkzeug Team, see AUTHORS for more details.
- QUnitJS - Copyright (c) 2013 jQuery Foundation, http://jquery.org/
- redis - Copyright (c) by Salvatore Sanfilippo and Pieter Noordhuis
- Simple Logging Facade for Java - Copyright (c) 2004–2013 QOS.ch
- Six - Copyright (c) 2010–2015 Benjamin Peterson
- Six - yum distribution - Copyright (c) 2010–2015 Benjamin Peterson
- Spymemcached / Java Memcached - Copyright (c) 2006–2009 Dustin Sallings and Copyright (c) 2009–2011 Couchbase, Inc.
- Supervisord - Supervisor is Copyright (c) 2006–2015 Agendaless Consulting and Contributors.
- Switchery - Copyright (c) 2013–2014 Alexander Petkov
- Toastr - Copyright (c) 2012 Hans Fjallemark & John Papa.
- Underscore js - Copyright (c) 2009–2014 Jeremy Ashkenas, DocumentCloud and Investigative Reporters & Editors
- Zlib - Copyright (c) 1995–2013 Jean-loup Gailly and Mark Adler

Permission is hereby granted, free of charge, to any person obtaining a copy of the above third-party software and associated documentation files (collectively, the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notices and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE LISTED ABOVE 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.

VMware Carbon Black
1100 Winter Street, Waltham, MA 02451 USA
Tel: 617.393.7400
Fax: 617.393.7499
Email: support@carbonblack.com
Web: http://www.carbonblack.com

CB Response User Guide
Product Version: 7.1
Document Revision Date: March 20, 2020

About this Document

This guide is written for both the on-premises and cloud editions of CB Response.

Sections

Topic	Page
Intended Audience	6
CB Response Terminology	6
What this Documentation Covers	7
Other Documentation	9
Community Resources	10
Contacting Support	10

Intended Audience

This documentation is for administrators and for Security Operations Center (SOC) and Incident Response (IR) personnel. It is for those setting up and maintaining security for endpoints and networks, as well as users assessing potential vulnerabilities and detecting advanced threats. Staff who manage CB Response activities should be familiar with:

- Linux, Microsoft Windows, and macOS operating systems
- Web applications
- Desktop infrastructure (especially in-house procedures for software roll-outs, patch management, and antivirus software maintenance)
- Effects of unwanted software

CB Response Terminology

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

Term	Definition
Binary	Executable file (for example, PE Windows file, ELF Linux file, or Mach-O Macintosh file) that is loaded onto a computer file in binary form for computer storage and processing purposes. CB Response only collects binaries that execute. It does not collect scripts, batch files, or computer files that are created or modified. • CB Response does collect the script or batch file names from command prompts and command lines. • CB Response also collects file names and paths as they are created or modified.
CB Response Sensor	Lightweight data gatherers installed on hosts on the deployed network. They gather event data on the hosts and securely deliver it to the CB Response server for storage and indexing.
CB Response Server	A CentOS server that exists on the deployed network. It receives data from sensors, stores and indexes that data, and provides access to the data through the CB Response console.
CB Threat Intel Feeds	Pre-configured threat intelligence feeds. These feeds contain threat intelligence data. These feeds come from various sources: • Carbon Black • Our MSSP/IR partners • Our customers • Open-source CB Threat Intel feeds provide a list of Indicators of Compromise (IOCs) and contextual information based on binary/process attributes and events (MD5, SHA-256, IP, Domain). These attributes and events are scored and rated, and then correlated with any matching files in your environment. For more information, see Chapter 15, “Threat Intelligence Feeds.”.

Term	Definition
CB Threat Intel Server	A server that is managed by Carbon Black and augments the functionality of the CB Response server.
Data File	A computer file that is a resource for storing information that requires a computer program (executable or binary file) to run. Data files are not captured by CB Response sensor.
Indicators of Compromise (IOCs)	CB Response sensors constantly monitor your computers for IOCs and send alerts to the CB Response console when detected. Queries are dynamic indicators that look at behaviors that are continuously recorded by sensors on endpoints and centrally recorded for analysis. Hashes (MD5, SHA-256), IP addresses, and domain names are static indicators that are similar to signatures. They are used to identify suspected malicious activity.
MD5	Unique cryptographic hash identifier for a binary instance in CB Response.
Process	An instance of the execution of a binary file.
Watchlist	Fully customizable searches that contain lists you can use to track specific IOCs. Watchlists are saved searches that are visible to all users. They can be used for searching either processes or binary executable files.

What this Documentation Covers

CB Response User Guide is your guide to managing CB Response, installing sensors on endpoints, and using CB Response to monitor file activity and threats on your endpoints. While this guide describes all features, access to some features requires particular user privileges. See “Managing User Accounts for On-Premise Servers” on page 51 or “Managing User Accounts for Cloud Servers” on page 65 for more information about user roles and privileges.

The following table summarizes the contents of the contents of this guide:

Chapter	Description
CB Response Overview	Introduces CB Response, explains key concepts, and suggests operating strategies for managing sensors and data to provide the visibility, detection, and response capabilities in the CB Response solution.
Getting Started	Explains how to log in and out of CB Response (on-premises) and introduces the CB Response main menu. Explains how to log in and out of CB Response Cloud, introduces the CB Response main menu, and how to set up two-factor authentication.

Chapter	Description
Managing User Accounts for On-Premise Servers	Describes how to manage access to the CB Response (on-premises) console for users and teams of users.
Managing User Accounts for Cloud Servers	Describes how to manage access to the CB Response Cloud console for users and teams of users.
Installing Sensors	Describes installing and upgrading sensors on Windows, macOS, and Linux systems.
Managing Sensors	Provides an overview of how sensors work, the information that they provide, and how to modify their configuration.
Sensor Groups	Describes creating, moving, editing, and deleting sensor groups, which determine what kind of information is provided by sensors and who can access the information.
Managing Certificates for Server-Sensor Communication	Describes how CB Response uses HTTPS and TLS to secure and authorize server-sensor communications; describes certificate management features such as certificate addition and strict validation.
Responding to Endpoint Incidents	Describes CB Response features for incident response—endpoint isolation, Live Response, and process hash banning.
Process Search and Analysis	Describes how to perform detailed process searches and in-depth analysis of the processes in search results.
Binary Search and Analysis	Explains how to search for and analyze binary metadata.
Advanced Search Queries	Describes CB Response query syntax and how to construct advanced queries to search for processes and binaries.
Ingress Filtering	Describes how to perform common tasks related to ingress filters.
Threat Intelligence Feeds	Describes CB Threat Intel feeds that, when enabled on a CB Response server, improve threat verification, detection, visibility, and analysis on your endpoints.
Configuring the Event Forwarder	Explains how to configure the Event Forwarder through the CB Response console.
Creating and Using Investigations	Describes how to work with investigations, which provide a way to group data for reporting, compliance, or retention purposes.
Watchlists	Describes creating and using watchlists, which are saved searches that are visible to all users.

Chapter	Description
Console and Email Alerts	Describes creating and managing CB Response alerts, which can be displayed in the console and also sent through email.
Using the Head-Up Display Page	Explains how to use the HUD (Head-Up Display) page, which is a customizable dashboard for CB Response users.
Netconn Metadata	Explains how to use Transport Layer Security (TLS) fingerprinting in CB Response 7.1 and later.
Sensor Parity	Contains two tables that outline availability of features on sensor systems in different operating systems.
Sensor Health Score Messages	Describes sensor health score messages that display on the Sensor Details page.

Other Documentation

Visit the Carbon Black User Exchange website at https://community.carbonblack.com to locate documentation for tasks not covered in this guide as well as other documents maintained as a knowledge base for technical support solutions. Some of these documents are updated with every newly released build, while others are updated only for minor or major version changes. Documents include:

- *CB Response Release Notes* – Provides information about new and modified features, issues resolved and general improvements in this release, and known issues and limitations. It also includes required or suggested preparatory steps before installing the server.
- *CB Response Operating Environment Requirements (OER)* – Describes performance and scalability considerations in deploying a CB Response server.
- *CB Response Server Configuration Guide (cb.conf)* – Describes the CB Response server configuration file (`cb.conf`), including options, descriptions, and parameters.
- *CB Response Server/Cluster Management Guide* – Describes how to install, manage, backup/restore, etc. a CB Response server/cluster. This guide is for on-premises CB Response installations only.
- *CB Response User Guide* – (this document) Describes the CB Response product and explains how to use all of its features and perform administration tasks.
- *CB Response Unified View User Guide* – Describes how to install and manage CB Response Unified View.
- *CB Response Integration Guide* – Provides information for administrators who are responsible for integrating CB Response with various tools, such as CB Protection, EMET, VDI, SSO, and more.
- *CB Response API* – Documentation for the CB Response REST API is located at https://developer.carbonblack.com/reference/enterprise-response. Documentation for the Python module that can be used for easy access to the REST API is hosted at https://cbapi.readthedocs.io.

- *CB Response connectors* – Documentation describing how to install, configure and maintain various Carbon Black connectors is located at https://developer.carbonblack.com/guide/enterprise-response/#connectors. A connector enables communication between a third-party product and CB Response server.

Community Resources

The Carbon Black User Exchange 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 Carbon Black products.

When you log into this resource, you can:

- Ask questions and provide answers to other users' questions.
- Enter a "vote" to bump up the status of product ideas.
- Download the latest 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

Carbon Black Technical Support offers several channels for resolving support questions:

Technical Support Contact Options
Carbon Black User Exchange: https://community.carbonblack.com
Carbon Black Support Site: https://www.carbonblack.com/resources/support
Email: support@carbonblack.com
Phone: 877.248.9098
Fax: 617.393.7499

When you contact technical support, provide the following information:

Required Information	Description
Contact	Your name, company, 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)

Required Information	Description
Document version	For documentation issues, the title, version and date of the manual you are using. Date and version appear on the cover page, or for longer manuals, at the end of the Copyrights and Notices section.
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

About this Document 5
Intended Audience 6
CB Response Terminology 6
What this Documentation Covers 7
Other Documentation 9
Community Resources 10
Contacting Support 10

1 CB Response Overview 26
What is CB Response? 27
System Architecture 29
CB Response Cloud 30
CB Response On-Premises 30
Data Flow Diagrams 32
CB Response Workflow Overview 35
CB Response APIs 36

2 Getting Started 37
Logging In (On-Premises) 38
Logging In and Configuring Two-Factor Authentication (Cloud Only) 38
Logging In for the First Time from an Email Invitation (Cloud) 39
Configuring Two-Factor Authentication (Cloud) 40
Logging in After Initial Login (Cloud) 43
Enabling/Disabling Two-Factor Authentication (Cloud) 44
Logging Out (Cloud and On-Premises) 44
CB Response Console Controls 45
Navigation Bar 45
Username Menu 47
EU Data Sharing Banner 49
Notifications 49
Help: User Guide and Customer Support 49

3 Managing User Accounts for On-Premise Servers 51
Overview of User Management (On-Premises) 52
Managing User Access with Teams 53
Role-based Privileges for Teams 53
Analyst & Viewer Access by Feature 54
Adding Enhanced Permissions for Analysts 56
User/Team Permissions Example 58
Creating Teams 58
Modifying Teams 59
Deleting Teams 59
Creating User Accounts for On-Premise Servers 60

Changing Passwords 61
Resetting API Tokens 62
Deleting User Accounts 62
Viewing User Activity 63
User Activity API Audit Logging 64

4 Managing User Accounts for Cloud Servers 65
Overview of Cloud User Management 66
Creating Cloud User Accounts 67
Inviting a New or Existing User to Access a Cloud Server 67
Activating an Account from an Invitation 68
Accessing Authorized Servers 70
User Account Lockout 70
Unlocking an Account 70
Viewing and Modifying Cloud User Accounts 70
Changing Security Settings, Email Address or Full Name 71
Changing Administrator / User Status 72
Resetting API Tokens 74
Viewing User Activity 74
Removing a User Account 74

5 Installing Sensors 76
Overview of Sensor Installation 77
Supported Operating Systems and Versions 77
Installing Sensors on Windows 77
HTTP Proxy Support in Windows Sensors 78
Uninstalling Windows Sensors 79
Installing Sensors on macOS Systems 80
Upgrading Sensors on macOS 81
Uninstalling Sensors on macOS 81
Installing Sensors on Linux Systems 81
Upgrading Sensors on Linux 82
Uninstalling Sensors on Linux 82
Upgrading Sensors 83
Uninstalling Sensors via the Console 83
Obtaining New Sensor Installation Packages 84

6 Managing Sensors 85
Overview of Sensor Management 86
Monitoring Sensor Status and Activity 87
The Sensors Page 87
Searching for Sensors 89
Exporting Sensor Data 90
Sensor Actions 90
Monitoring Sensor and Server Information 91
Viewing Sensor Details 93
Sensor Details Heading and Options 93
Summary 94

Vitals and Configuration 95
Sensor Vitals 95
Computer Vitals 96
Configuration 96
Team Access 97
Recent Activity 97
Activity 97
Resource Status 97
Diagnostics 98
Communication Failures 98
Driver Diagnostics 98
Reducing the Impact of Netconn Data Collection (Windows) 100
Event Diagnostics 100
Status History 100
Upgrade Status 101
Uninstall Status 101

7 Sensor Groups 102
Overview of Sensor Groups 103
Create or Edit a Sensor Group 103
General Settings 104
Sharing Settings 105
Advanced Settings 107
Permissions Settings 109
Event Collection Settings 109
Exclusion Settings (OS X/macOS only) 110
Creating Exclusions 110
Upgrade Policy Settings 112
Moving Sensors to Another Group 112
Deleting Sensor Groups 113

8 Managing Certificates for Server-Sensor Communication 114
TLS Server Certificate Management Overview 115
Certificate Management Feature Summary 116
Server-Sensor Certificate Requirements 116
How CB Response Supports Multiple Certificates 117
Using Multiple Active Certificates in a Cluster 119
Managing Certificates on the Server 120
Viewing Certificate Information in the Console 120
Substituting a Legacy Certificate during Server Installation 121
Adding Certificates through the Console 122
Choosing a Validation Option 122
Changing the Expiration Notification Period 123
Deleting Certificates 124
Upgrades from Previous Server Releases 124
Assigning Certificates to Sensor Groups 125
Assigning different certificates to different sensor groups 125
Assigning a new certificate to all sensor groups 126
Sensor Support for Certificate Management 126

Special Requirement for Windows Sensors 127
Upgrading to Sensors that Allow Certificate Management 128

9 Troubleshooting Sensors 129
Troubleshooting Windows Sensor Installations 130
Using Control Codes to Generate Logs of Diagnostic Data 130
Debugging Sensor Communications 132
Troubleshooting Linux Sensor Installations 133
General Logging 133
Installation Verification 134
Installation Failures 135
Sensor Communication History 135
Manual Sensor Daemon Start and Stop 135
Determine Server URL 136
Initiate an Immediate Checkin to the Server 136
Driver Debug Parameters 136
Sensor Version 6.1.x Driver Debug Parameters 136
Sensor Version 6.2.x Driver Debug Parameters 136
Daemon Debug Options 137
Debugging Parameters for 6.1.x Sensors 137
Debugging Parameters for 6.2.x Sensors and Later 138
Determine Sensor Version 138
Trigger a Diagnostic Data Dump 138
Troubleshooting OSX Sensor Installations 139
Installation Verification 139
Installation Failures 139
Communications Logging 139
Manual Sensor Daemon Start and Stop 140
Determining Sensor Version 140
Determine Server URL 140
Initiate an Immediate Checkin to the Server 140
Initiate a Diagnostic Data Dump 140
Diagnostic Uploads Utility 140
Automatic Crash Data Upload 141
Manual Upload Option (Command Line Utility) 141
Enabling Sensor Diagnostics Uploads 142
File Transfer and Security 143
Data Collected by Sensor Diagnostics 143

10 Responding to Endpoint Incidents 145
Overview of Incident Response 146
Isolating an Endpoint 147
Isolation Exclusions 149
Using Live Response 150
Enabling and Configuring Live Response 150
Tuning Live Response Network Usage 152
Live Response Endpoint Sessions 152
Status, Error and Progress Messages 157
Ending Live Response Sessions 158

Registry Access in Live Response 158
Detached Session Management Mode 160
Extending Live Response 161
Live Response Activity Logging and Downloads 161
Banning Process Hashes 162
Creating Process Hash Bans 163
Banning a List of Hashes 164
Managing and Monitoring Hash Bans 166
The Manage Banned Hashes Page 166
Monitoring Banning Events 168
Searching for Blocked Processes 168
Enabling Alerts and Syslog Output for Banning Events 170
Disabling a Hash Ban 171
Disabling or Restricting the Hash Ban Feature 171
Disabling Bans in a Sensor Group 171

11 Process Search and Analysis 172
Overview of Process Search 173
Time Filter 173
Search Filters 174
Enable/Disable Filters 174
Select Multiple Filter Rows 175
Filter Row Percentages 176
Filter Search Fields 176
Search Field 176
Save Searches 177
Clear Saved Searches 177
Add Search Terms 177
Reset Search 178
Group By Process 179
Search Result Messages 179
Get Comprehensive Results 179
Example Process Search 179
Managing High-Impact Queries 180
Responding to Blocked Searches 180
Process Search Settings in the Console 180
Process Search Settings in cb.conf 181
Results Table 182
Results Table Options 182
Results Table Row Details 183
Process Analysis Page 184
Process Analysis Features 186
Process Summary 186
Interactive Process Tree 188
Process Execution Details 189
Binary Metadata 190
EMET Protections Enabled (Windows Only) 191
Process Event Filters 191
Event Timeline 193

Process Event Details 193
Process Event Types 196
Analysis Preview Page 199

12 Binary Search and Analysis 201
Overview of Binary Search 202
Entering Search Criteria 202
Additional Binary Search Page Features 204
High-level Result Summaries 204
Related Metadata 205
Binary Search Results Table 206
Binary Preview 207
Binary Analysis 207
Binary Overview 208
General Info 209
Frequency Data 210
Digital Signature Metadata 210
File Version Metadata 211
Observed Paths 212
Observed Hosts and Sensor IDs 212

13 Advanced Search Queries 213
Query Syntax Details 214
Terms, Phrases, and Operators 214
Restrictions on Terms 214
Whitespace 214
Parenthesis 215
Negative Sign 215
Double Quotes 215
Leading Wildcards 216
Fields in Process and Binary Searches 216
Fields in Alert and Threat Report Searches 222
Field Types 224
domain 224
ipaddr 224
ipv6addr 225
text 225
count 225
datetime 225
keyword 226
md5 226
sha256 226
ja3 226
ja3s 226
path 227
Wildcard Searches 227
Modload Path Searches 227
Regmod Path Searches 227
bool 228

sign 228
cmdline 228
Tokenization Rules 228
Tokenization Changes on Server Upgrade 229
Retention Maximization and cmdline Searches 231
Searching with Multiple (Bulk) Criteria 231
Searching with Binary Joins 233
Example Searches 234
Process Search Examples 234
Binary Search Examples 238
Threat Intelligence Search Examples 239

14 Ingress Filtering 240
Overview of Ingress Filtering 241
Viewing and Configuring Ingress Filters 241
Adding an Ingress Filter 242
Regex Filters 243

15 Threat Intelligence Feeds 244
Overview of Threat Intelligence Feeds 245
Threat Intelligence Feed Scores 246
Firewall Configuration for Feeds 246
Managing Threat Intelligence Feeds 246
Checking for New Threat Intelligence Feeds 248
Syncing Threat Intelligence Feeds 248
Data Sharing Settings 248
Enabling, Disabling, and Configuring a Feed 251
On-Demand Feeds from CB Threat Intel 253
Creating and Adding New Feeds 254
Searching for Threat Reports 255
Threat Report Searches and Results 256
Threat Report Details 258
Ignoring Future Reports 259

16 Configuring the Event Forwarder 260
Overview of the Event Forwarder 261
Configuring the Event Forwarder in the Console 261
Edit and Status 261
Events 262
Output 263
Splunk 263
S3 264
HTTP 265
Syslog 266
Certificates and Credentials 266

17 Creating and Using Investigations 267
Overview of Investigations 268
Viewing Investigations 268

Investigations Menu Bar 268
Event Types 269
Bar Graph 269
Events Table 270
Edit Event Description 270
Child Processes 271
Creating Investigations 271
Adding Events to Investigations 271
Removing Events from Investigations 272
Adding Custom Events to Investigations 272
Deleting Investigations 272

18 Watchlists 274
Overview 275
Viewing Watchlists and their Results 275
The Watchlists Page 276
The Watchlist Details Panel 276
Built-in and Community Watchlists 277
Creating Watchlists 278
Managing Watchlists 281
Watchlist Status 281
Watchlist Expiration 282
Slow or Error-producing Watchlists 282
Editing Watchlists 283
Deleting Watchlists 284

19 Console and Email Alerts 285
Overview of Alerts 286
Enabling Console Alerts 286
Watchlist Alerts 286
Threat Intelligence Feed Alerts 287
Viewing Alert Activity on the HUD Page 287
Managing Alerts on the Triage Alerts Page 289
Displaying the Report Name 290
Reviewing Alerts 290
Alerts Table Data 291
Managing Alert Status 292
Ignoring Future Events for False Positive Alerts 293
Enabling Email Alerts 294
Configuring an Email Server 294
Enabling Specific Email Alerts 295

20 Using the Head-Up Display Page 296
Overview of HUD 297
Viewing the HUD Page 297
Customizing the HUD Page 297
Sortable Columns 297
Endpoint Hygiene Panel 298

Event Monitor Panel 298
Query Duration Panel 299
Resolution Time Panel 300
Saved Searches Panel 300
Sensors Panel 300
Unresolved Alerts Panel 302

21 Netconn Metadata 303
Overview of TLS Fingerprinting 304
How TLS Fingerprinting Works 304
CB Response Implementation of TLS Fingerprinting 305
Process Search 305
Process Analysis 305
Watchlists 305

A Sensor Parity 307
Sensor Feature Support 308
Sensor Group Feature Support 310

B Sensor Health Score Messages 311
Windows Health Events 312
Priority List 312
Driver and Component Failures 312
Cause 312
Impact 312
Severity Scale 312
Remediation 312
Memory Usage 312
Cause 312
Impact 312
Severity Scale 313
Remediation 313
GDI Handle Count 313
Cause 313
Severity Scale 313
Remediation 313
Handle Count 313
Cause 313
Severity Scale 314
Remediation 314
Disk Space 314
Cause 314
Impact 314
Severity Scale 314
Remediation 314
Event Loss 314
Cause 314
Impact 314
Severity Scale 315
Remediation 315

Event Load 315
Cause 315
Impact 315
Severity Scale 315
Remediation 315
macOS Health Events 316
Priority List 316
Memory Usage 316
Cause 316
Impact 316
Severity Scale 316
Remediation 316
Out of License 316
Cause 316
Impact 316
Severity Scale 317
Remediation 317
Upgrade Issue 317
Cause 317
Impact 317
Severity Scale 317
Remediation 317
Proxy Driver Failure 317
Cause 317
Impact 317
Severity Scale 317
Remediation 318
Procmon Driver 318
Cause 318
Impact 318
Severity Scale 318
Remediation 318
Netmon Driver 318
Cause 318
Impact 318
Severity Scale 318
Remediation 318
Linux Health Events 319
Priority List 319
Out of License 319
Cause 319
Impact 319
Severity Scale 319
Remediation 319
Apply updated license to the CB Response server. 319
Failed to get Event log Stats 319
Cause 319
Impact 319
Severity Scale 319
Driver Failure 320

Cause 320
Impact 320
Severity Scale 320
Memory Usage 320
Cause 320
Impact 320
Severity Scale 320
Remediation 320

List of Tasks

How to . . .

To access the Ingress Filters page: 241
To access the Sensor Details page: 93
To access the Sensors page: 87
To activate a new account from an invitation: 69
To activate access to a new server from an existing account: 69
To add a new certificate to a server through the console: 122
To add a new threat intelligence feed to the CB Response server: 254
To add an ingress filter on the Ingress Filters page: 242
To add an ingress filter on the Process Search page: 242
To add events to investigations: 271
To add Exclusion settings to the sensor group panel on the Sensors page: 110
To add or remove administrator status for a user: 72
To add search terms: 177
To apply one certificate to all sensor groups: 126
To ban a list of process hashes: 164
To ban a process MD5 hash from the Process Analysis page: 163
To block or allow high-impact process searches: 181
To change cloud account details: 71
To change the notification period for an expiring certificate: 124
To change the server certificate for one sensor group: 125
To change the status of a single alert: 292
To change the status of all alerts matching a search and/or filter: 292
To change the validation method for server certificates: 123
To change your password: 61
To check for new Threat Intelligence feeds: 248
To clear saved searches: 177
To configure an email server for alerts: 294
To configure forwarded events: 262
To configure the Event Forwarder (recommended sequence): 262
To configure two-factor authentication: 40
To configure watchlist expiration: 282
To create a custom event: 272
To create an investigation: 271
To create an isolation exclusion: 149
To create an on-premises user account: 60
To create an OS X/macOS event collection exclusion for a sensor group: 110
To create and attach to a Live Response sensor session: 153
To create or edit a sensor group: 103
To create teams: 58
To create watchlists from Process Search or Binary Search pages: 278
To create watchlists from the Threat Intelligence Feeds page: 280
To create watchlists from the Watchlists page: 279
To delete a CB Response team: 59
To delete a certificate from a server: 124
To delete a user account: 62
To delete a watchlist: 284
To delete investigations: 273
To delete sensor groups: 113

To disable a process hash ban: 171
To disable a threat intelligence feed: 253
To disable console alerts for a threat intelligence feed: 287
To disable process hash bans in a sensor group: 171
To display a table of reports from one threat intelligence feed: 256
To display online documentation from the console: 50
To display only certain search filters on the Process Search page: 174
To do a bulk IOC search on the Binary Search page: 232
To do a bulk IOC search on the Process Search page: 231
To edit watchlists: 283
To enable alerts and syslog recording of blocking events due to hash bans: 170
To enable and configure a threat intelligence feed: 251
To enable console alerts for a threat intelligence feed: 287
To enable console alerts for a watchlist: 286
To enable data sharing with Carbon Black threat intelligence feed partners: 250
To enable email alerts for a threat intelligence feed: 295
To enable email alerts for a watchlist: 295
To enable or disable Live Response via cb.conf (on-premise servers only): 151
To enable or disable Live Response via the CB Response console: 151
To enable or disable two-factor authentication: 44
To enable sharing communications: 249
To enable the display of report names: 290
To end a Live Response session: 158
To end network isolation for endpoints: 148
To execute a saved search: 177
To export sensor data from the Sensors page: 90
To ignore reports, use one of the following options: 259
To ignore the triggering event for an alert: 293
To install sensors on Linux endpoints: 81
To install sensors on macOS endpoints: 80
To install sensors on Windows endpoints: 78
To invite a user to open a CB Response Cloud account or extend it to a new server: 67
To isolate endpoints: 147
To issue a sensor control request to the sensor: 130
To log into CB Response from an email invitation: 39
To log into the CB Response console after initial login: 43
To log into the CB Response console: 38
To log out of the CB Response console: 44
To manually uninstall Linux sensors: 82
To manually uninstall macOS sensors: 81
To manually uninstall Windows sensors: 79
To modify a CB Response team: 59
To modify Live Response network usage: 152
To move sensors to a new sensor group: 112
To open a Live Response command window without a session: 160
To open the Search Threat Reports page (unfiltered): 256
To open the Triage Alerts page: 289
To perform a binary search: 203
To provide enhanced Analyst permissions to a user: 57
To remove a CB Response Cloud user account: 74
To remove an event from an investigation: 272
To reposition HUD panels: 297

To reset the API token for a user account: 62
To reset the API token for a user account: 74
To resize HUD panels: 297
To search for a sensor: 89
To search for processes that have block events: 168
To sync all threat intelligence feeds on the page: 248
To uninstall sensors using the console (all platforms): 84
To upload a certificate: 266
To upload a custom “legacy” certificate during server installation: 121
To use the time filter: 173
To view all block events for a parent process: 168
To view banned hash alerts: 170
To view server and sensor information in the Server Dashboard: 91
To view the available certificates on a server: 120
To view the CB Response Cloud servers to which you have access: 70
To view the Event Forwarder Settings page: 261
To view the HUD page: 297
To view the Threat Intelligence Feeds page: 247
To view user activity: 63

Chapter 1

CB Response Overview

This chapter introduces CB Response, explains key concepts, and suggests operating strategies for managing sensors and data to provide the visibility, detection, and response capabilities in the Carbon Black solution.

Sections

Topic	Page
What is CB Response?	27
System Architecture	29
Data Flow Diagrams	32
CB Response Workflow Overview	35
CB Response APIs	36

What is CB Response?

CB Response provides endpoint threat detection and a rapid response solution for Security Operations Center (SOC) and Incident Response (IR) teams. With CB Response, enterprises can continuously monitor and record all activity on endpoints and servers. The combination of CB Response's endpoint visibility with CB Threat Intel helps enterprises to proactively hunt for threats, customize their detection, and respond quickly. The following diagram shows how CB Response features work together to help you answer these questions:

- How did the problem start?
- What did the threat do?
- How many machines are infected?
- How can we resolve the threat?

CB Response provides these solutions:

- **Visibility** – Know what's happening on every computer at all times. With CB Response, you have immediate real-time visibility into the files, executions, network connections, and critical system resources on every machine, and the relationships between them. You can see how every file got there, what created it, when it arrived, what it did, if it made a network connection, if it deleted itself, if a registry setting was modified, and much more.
- **Detection** – See and record everything; detect attacks in real time without signatures. CB Response's threat research team analyzes threat techniques and creates Advanced Threat Indicators (ATIs) to alert you to the presence of an attack. These ATIs look for threat indicators and are not based on signatures. You can detect advanced threats, zero-day attacks, and other malware that evades signature-based detection tools—in real time. There is no wait for signature files, no testing and updating .dat files, and no sweeps, scans or polls. You get immediate, proactive, signature-less detection.

- **Response** – Use a recorded history to see an the full “kill chain” of an attack, and contain and stop attacks. When you need to respond to an alert or threat, you will instantly have the information you need to analyze, scope, contain, and remediate the problem. With the recorded details about every machine, you can “go back in time” to see what happened on any of your machines to understand the full “kill chain” of an attack. You will also have a copy of any binary that ever executed, so you can analyze it yourself, submit it to a third party, and so on. You can also contain and stop attacks by globally blocking the execution of any file automatically or with a single click.
- **Prevention via CB Protection** – Stop attacks with proactive, signature-less prevention techniques by integrating the CB Protection with CB Response. With CB Protection, you can choose from different forms of advanced endpoint protection to match your business and systems. CB Protection’s proactive “Default-Deny” approach ensures that only software you trust can run on your machines. CB Protection’s “Detect-and-Deny” technology uses ATIs to detect malware and stop its execution, and CB Protection’s unique “Detonate-and-Deny” approach automatically can send every new file that arrives on any endpoint or server to leading network security tools for “detonation.” If a tool reports finding malicious files, CB Protection can automatically stop them from running on all of your machines.

CB Response accelerates detection by going beyond signatures, and reduces the cost and complexity of incident response. Using a real-time endpoint sensor, CB Response delivers clear and accurate visibility and automates data acquisition by continuously recording and maintaining the relationships of every critical action on all machines, including events and event types such as executed binaries, registry modifications, file modifications, file executions, and network connections.

CB Response provides a cross-process event type that records an occurrence of a process that crosses the security boundary of another process. While some of these events are benign, others can indicate an attempt to change the behavior of the target process by a malicious process.

Unlike scan-based security solutions, CB Response can expand detection beyond the moment of compromise with its robust endpoint sensor and access to the information provided by the CB Threat Intel.

CB Threat Intel provides three types of intelligence:

- **CB Threat Intel Reputation** – A cloud-based intelligence database that provides highly accurate and up-to-date insight into known-good, known-bad, and unproven software. It provides IT and security teams with actionable intelligence about the software installed in their enterprise. The capabilities of the reputation service are further enhanced by feeds from third party partners.
- **CB Response Threat Indicators** – Search for patterns of behavior or indicators of malicious behavior. Unlike signature-based detection, threat indicators can recognize distinct attack characteristics, based on the relationships between network traffic, binaries, processes loaded, and user accounts. CB Response also offers watchlists that are fully customizable saved searches that you can use to look for specific threat indicators.
- **Third Party Attack Classification** – Uses intelligence feeds from third-party sources to help you identify the type of malware and the threat actor group behind an attack. This enables security teams to have a better understanding of attacks so that they can respond more quickly and effectively. You can also leverage your own intelligence feeds to enhance response capabilities.

CB Response compares endpoint activity with the latest synchronization of CB Threat Intel feeds as it is reported. You can add intelligence feeds that you already have set up to give you zero-friction consumption of threat intelligence in CB Response, regardless of the source.

CB Response's sensor is lightweight and can be easily deployed on every endpoint, requiring little to no configuration. This enables endpoint security analysts and incident responders to deploy thousands of sensors across their environment to immediately answer key response questions.

CB Response's continuously-recorded sensor data is stored in a central server, which enables your team to see and understand the entire history of an attack, even if it deleted itself.

CB Response integrates with leading network security providers. This integration enables you to prioritize alerts that are detected on the network by correlating them with events that occurred on endpoints and servers. This enables you to fully investigate your entire enterprise instantly to accelerate detection, reduce dwell time, minimize scope, and immediately respond to and contain advanced threats.

You can use CB Response's APIs to customize or integrate with existing security technologies that you are using, and Security Information and Event Management systems (SIEMs).

System Architecture

This section provides a system architecture overview for both the cloud and on-premises versions of CB Response. In both systems, the CB Response server records *events* related to file changes, but copies of files and the data that changed are *not* recorded.

CB Response Cloud

The following diagram illustrates the components of a CB Response Cloud installation, which are:

- Sensors that can be installed on various endpoints such as laptops, desktops, servers, and point of sale (POS) machines.
- A cloud service that collects sensor data and makes it accessible with a web user interface or an API.
- The threat intelligence that includes the Carbon Black Threat Intel (CB Threat Intel) Reputation, CB Protection, and CB Response threat indicators, and third-party attack classification using CB Threat Intel partner feeds.

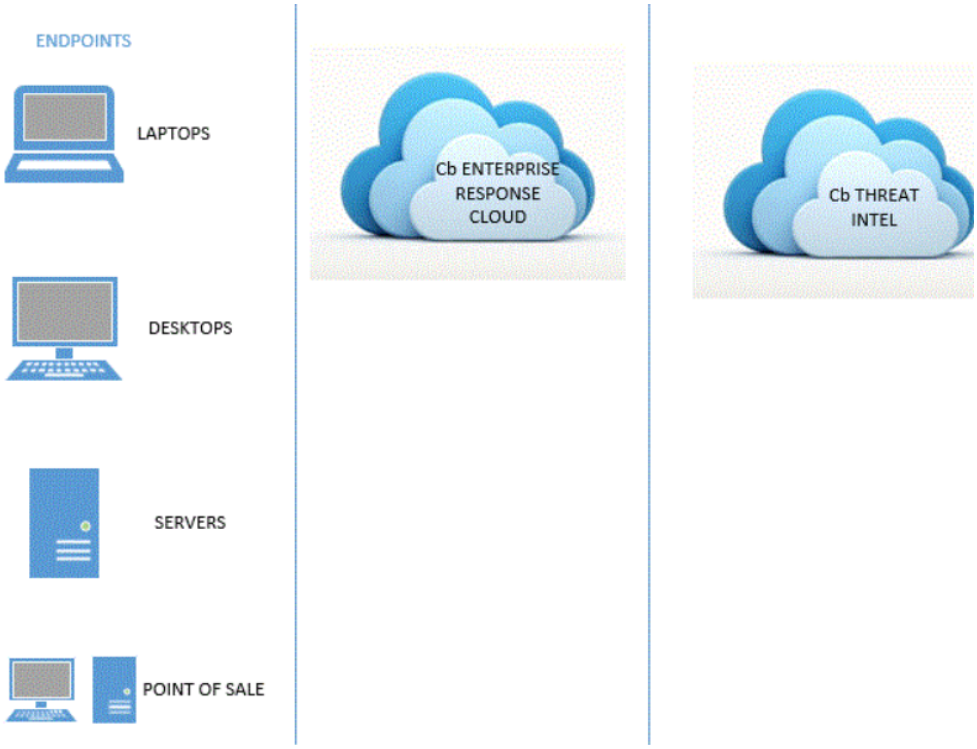

If your company is also using CB Protection, there is integration between it and CB Response. By leveraging CB Protection, you can contain advanced threats by globally blocking or banning them through CB Protection's customizable prevention techniques in the midst of a response.

CB Response On-Premises

A CB Response on-premises server software is installed on a Linux server. The following diagram illustrates the components of a CB Response installation, which are:

- Sensors that can be installed on various endpoints such as laptops, desktops, servers, and point of sale (POS) machines.
- A server that collects sensor data and makes it accessible with a web user interface or an API.

- The threat intelligence that includes the Carbon Black Threat Intel (CB Threat Intel) Reputation, CB Protection, and CB Response threat indicators, and third-party attack classification using CB Threat Intel partner feeds.

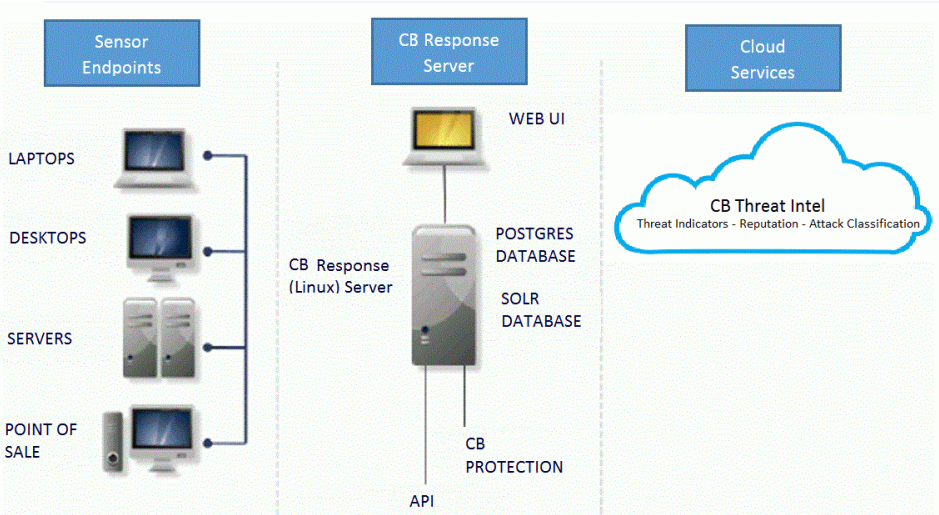

If your company is also using CB Protection, there is integration between it and CB Response. By leveraging CB Protection, you can contain advanced threats by globally blocking or banning them through CB Protection's customizable prevention techniques in the midst of a response.

Data Flow Diagrams

The following diagram illustrates the CB Response Cloud data flow:

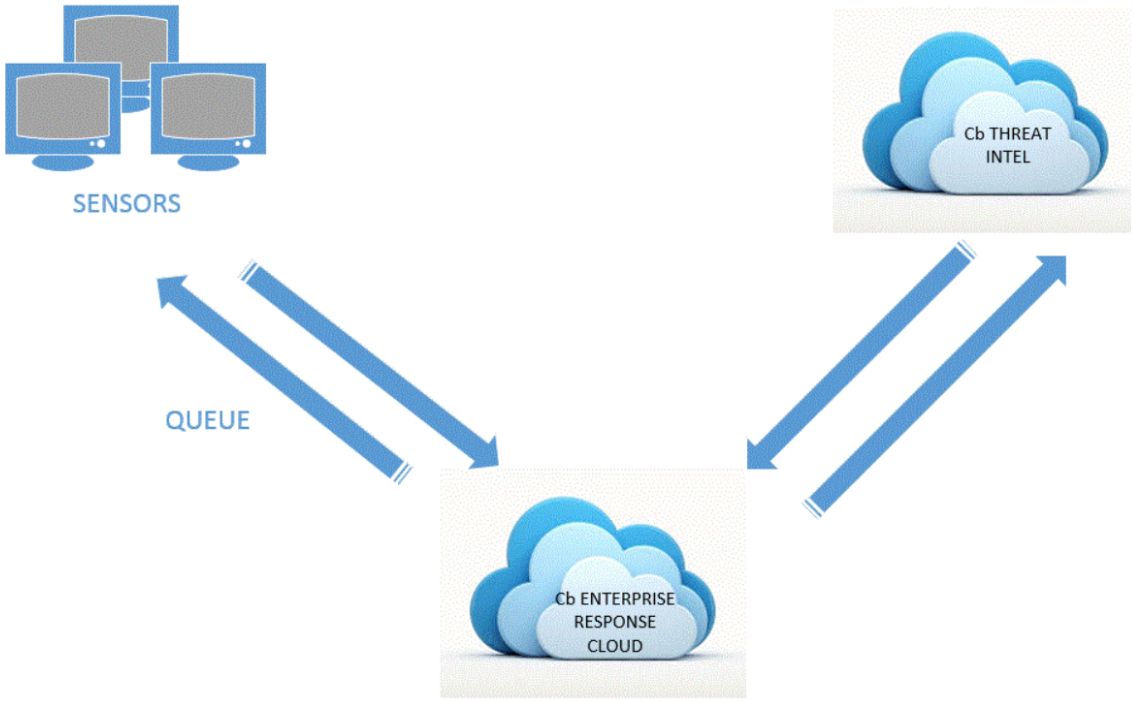

The following diagram illustrates the CB Response on-premises data flow:

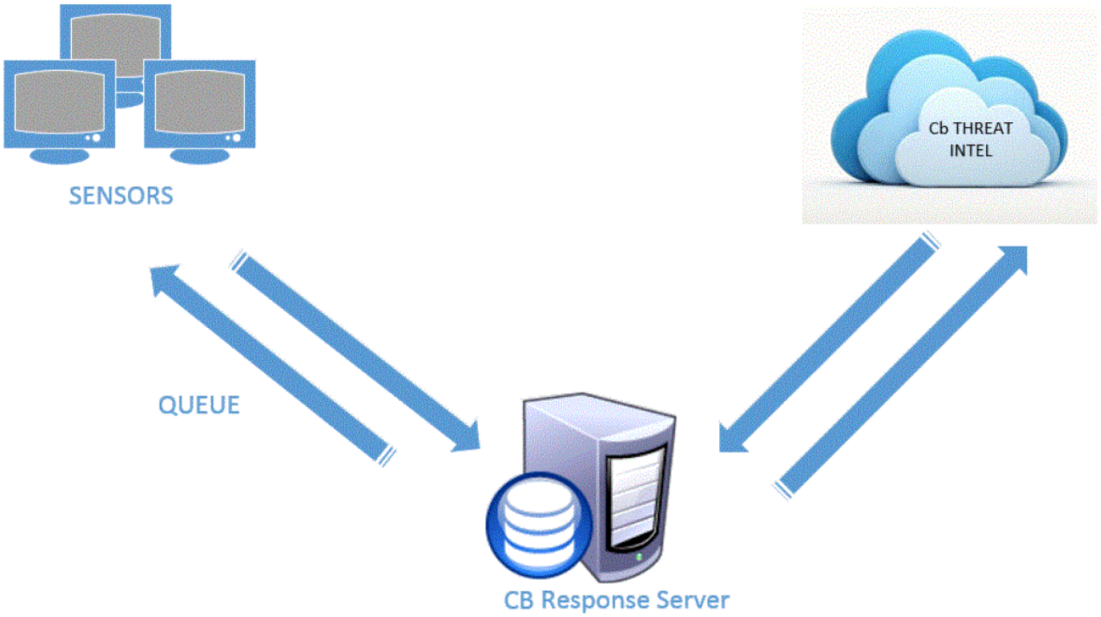

As soon as a sensor is installed, it begins buffering activity to report to the cloud service. This includes:

- Currently running processes that create events
- Binary executions
- File executions and modifications
- Network connections
- Registry modifications
- Cross-process events (events that cross the security boundaries of other processes)

Every few minutes, sensors check in with the cloud service, reporting what they have buffered, even if they are reporting that they have nothing buffered. When a sensor checks in, the cloud service responds, letting the sensor know when to send the data and how much data to send.

As the cloud service records data from sensors, the data is compared with the latest synchronization from any enabled CB Threat Intel feed partner. In most cases, incremental synchronizations occur hourly. Full synchronizations occur once every 24 hours by default.

Some CB Threat Intel feeds provide a list of all of the IOCs they track. Some feeds only include reports on files (identified through their MD5 or SHA-256 hashes) that are observed in your enterprise.

If you enable data sharing with the CB Threat Intel partners, CB Response pushes MD5 hashes that are observed by sensors and binaries originating from your enterprise to their cloud services. If there is a corresponding report or record, the feed is updated to include that information. If there is no corresponding third party-report, one is requested and when available, included in the feed.

When information about a specific binary is included in these feeds, the information remains there, even if the binary it is associated with is deleted from your endpoints and is no longer present in your environment.

The following table provides key additional information about data flows:

Data Flow	Description
Sensor to Server	• All communications are through HTTPS. • The TCP port is 443 by default, but is configurable. • Communications are always initiated from sensor to server (never from server to sensor). • By default, communications are mutually authenticated by statically pinned TLS certificates, both client and server. There is also an option to substitute user-provided certificates and use stricter validation. Sensors have the server's certificate embedded, and the server has all client certificates embedded. See "Managing Certificates for Server-Sensor Communication" on page 114 for more information. • All communications require a minimum of TLSv1+; only allow FIPS-compliant ciphers and use a 2048-bit Diffie Hellman key. • Sensor communication through a proxy is unsupported, unless the proxy is deployed in a transparent, in-line configuration. • Sensor communication is supported through transparent proxies. Due to certificate pinning, communication is not supported through traffic inspection proxies, or any other device that would affect SSL certificates. • The Windows sensor honors settings that are configured via a `proxy.pac` file. (This does not change the requirement that any proxy that is used must not modify SSL certificates or otherwise attempt to bypass the secure communications between sensor and server.) • Sensor communication through an TLS intercept/decryption device is not currently supported, even for in-line proxy configurations. • The server's sensor-facing interface can be configured in a DMZ to support endpoints outside the corporate LAN
Server to Alliance Server and CB Threat Intel	• All communications are explicitly opt-in. • All communications are HTTPS. • This connection is required for threat intelligence that is provided by CB Response. • TCP is 443 to api.alliance.carbonblack.com and threatintel.bit9.com. • Proxies are supported.
Server to yum Repository	• TCP is 443 for HTTPS to yum.distro.carbonblack.io. • TCP is 80 to a CentOS or RHEL.

CB Response Workflow Overview

Once sensors are installed and configured, your IT and security teams who are responsible for maintaining the health of your computer systems can perform basic tasks on a regular basis to ensure that there are no threats on any computer in your enterprise. Access to the CB Response user interface is via browser, although you can perform some functions through an API.

Note

Google Chrome is the only supported browser for this release. Although Firefox can be used, it will cause rendering issues on some pages and is not recommended. Other browsers should not be used for CB Response console access.

The basic workflow is continuous: you search for threats, analyze them, resolve then, and using the tools of your choice, prevent them from happening again. As you search, you can tag any items that seem unusual or that merit further investigation and then drill down further to find out more details about those items.

CB Response provides you with tools to help you detect and fix threats to your system. The following diagram shows the basic workflow that you use with CB Response:

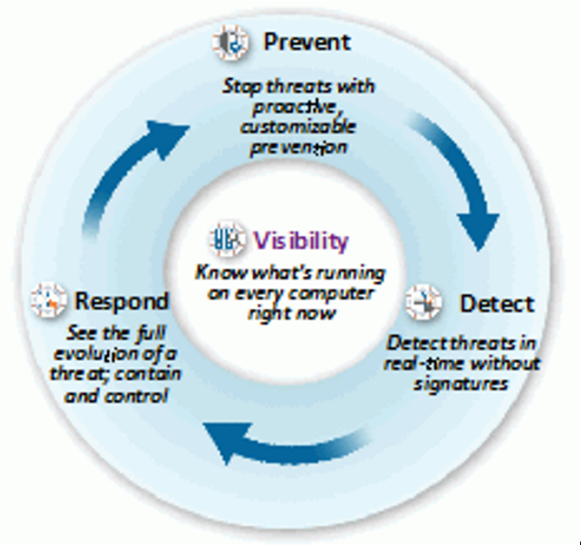

The following table shows how CB Response provides solutions to the problems you face.

Problem	Solution
What is the entry point of the threat?	Find out how the attacker got into your systems. Get oriented with visibility into everything that is running on every computer in your enterprise using the **Process Search** feature.
What did the attacker do?	Look deeper into suspicious processes and events to detect evidence of damage. Select processes that look suspicious and drill deeper using the **Process Analysis** feature.

Problem	Solution
How many machines were compromised?	Find out the scope of the damage by digging deeper into details about detected threats by using the **Process Details** and **Binary Details** pages. Set up **CB Threat Intel Feeds** and **Watchlists** by defining characteristics of interesting activity that you want to be notified about and receiving notifications as you need them. Create **Investigations** of suspicious processes to keep track of key events during a given response.
How do we respond to threats?	Find out how bad the threat is, and then determine how to respond to it by seeing its full evolution, containing the threat, and then controlling it.
How do we stop the threat from happening again?	Use the **Go Live** feature allows you to directly access content on endpoints that are running sensors which provide information. Set up **Watchlists** and **CB Threat Intel Feeds** that identify specific issues, and use the feeds and watchlists to perform continuous searches on your systems for immediate detection to help you stop the threat from happening again, and to ensure that you know of any new related activity.
How do we isolate threats?	You can isolate one or more Windows endpoints from the rest of your network and the Internet through the CB Response console. For more information, see “Isolating an Endpoint” on page 147.

Note

Access to CB Response features is determined by the permissions a logged-in user has. See “Managing User Accounts for On-Premise Servers” on page 51 and “Managing User Accounts for Cloud Servers” on page 65 for a description of how to create users with different permission levels.

CB Response APIs

CB Response Cloud includes extensive support for programmatic access to the underlying data and configuration through APIs. Documentation, example scripts, and a helper library for each of these libraries is available at https://developer.carbonblack.com.

Chapter 2

Getting Started

This chapter explains how to log in and out of CB Response, and includes instructions for using two-factor authentication for CB Response Cloud accounts. It also introduces the CB Response console controls that are available through the navigation bar and top menu, and summarizes the features that are accessible from those locations.

Sections

Topic	Page
Logging In (On-Premises)	38
Logging In and Configuring Two-Factor Authentication (Cloud Only)	38
Logging Out (Cloud and On-Premises)	44
CB Response Console Controls	45
Navigation Bar	45
Username Menu	47
EU Data Sharing Banner	49
Notifications	49
Help: User Guide and Customer Support	49

Logging In (On-Premises)

The CB Response console is a browser-based user interface for accessing the CB Response server and the information it collects from sensors and CB Threat Intel feeds. You log into the CB Response console from a supported web browser on any computer that has access to your server

Note

Google Chrome is the only supported browser for this release. Although Firefox can be used, it causes rendering issues on some pages and is not recommended. Other browsers should not be used for CB Response console access.

To log into the CB Response console:

1. From a supported web browser, enter the path to the CB Response server.
2. If your browser displays a warning about the certificate, you can safely ignore the warning and click through the remaining confirmation windows.

Note

To avoid future certificate warnings, accept the certificate permanently.

3. In the CB Response Login dialog box, enter your user name and password.
4. Click the **Login** button to display the CB Response **HUD** (Head Up Display) page.

Logging In and Configuring Two-Factor Authentication (Cloud Only)

This section explains how to:

- Log in to CB Response Cloud for the first time from an email invitation. See “Logging In for the First Time from an Email Invitation (Cloud)” on page 39.
- Configure two-factor authentication. See “To configure two-factor authentication:” on page 40
- Log in with or without two-factor authentication. See “Logging in After Initial Login (Cloud)” on page 43.

For information about using the CB Response User Management Console, see Chapter 4, “Managing User Accounts for Cloud Servers.”

Logging In for the First Time from an Email Invitation (Cloud)

If you have received an email inviting you to access a new CB Response Cloud, use the link in the email to either sign in with an existing account, or create a new account.

Note

The email link expires seven days after receipt.

To log into CB Response from an email invitation:

1. Click the link in your invitation email to open the CB Response **Login** dialog box.
2. Do one of the following:
   - If you already have an account, click **Sign In** and follow the procedures detailed in "Logging in After Initial Login (Cloud)" on page 43
   - To create a new account, enter values in the **Username**, **First**, **Last**, **Password**, and **Confirm Password** fields, and click **Sign up**.
3. If you are creating a new account, read the CB Response Cloud terms and conditions, and click **Accept**.

Note

This page only appears the first time you access the CB Response Cloud or when the terms and conditions are updated.

The **Two Factor Authentication** wizard opens where you can optionally configure two-factor authentication, which does the following:

- Adds a second authentication factor to your server
- Facilitates authentication management and security monitoring

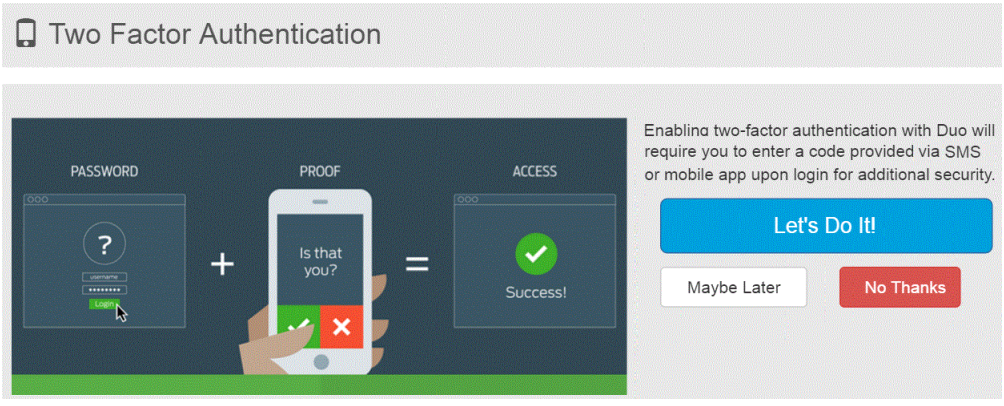

Two-factor authentication is available through Duo and requires that you download the Duo Mobile application on a device. (For more information, see https://duo.com).

If you change your mind later about using two-factor authentication, you can enable or disable it. (See "Enabling/Disabling Two-Factor Authentication (Cloud)" on page 44.)

Continue to the procedure in the next section to either decline or set up two-factor authentication.

Configuring Two-Factor Authentication (Cloud)

To configure two-factor authentication:

1. Log into the CB Response console.

   If you are logging in for the first time, or if you logged in previously and temporarily bypassed enrollment, the Two-Factor Authentication wizard appears.

2. On the first page of the Two Factor Authentication wizard, the following options are available:

   - **Maybe Later** – Bypass enrollment for now and proceed with logging in. You will have the option to configure two-factor authentication at your next login.
   - **No Thanks** – Do not enable two-factor authentication.
   - **Let's Do It!** – Enable two-factor authentication.

   If you decline to set up two-factor authentication, you are logged into the CB Response Cloud.

3. To set up two-factor authentication, select Let's Do It!, and then click **Start setup**.

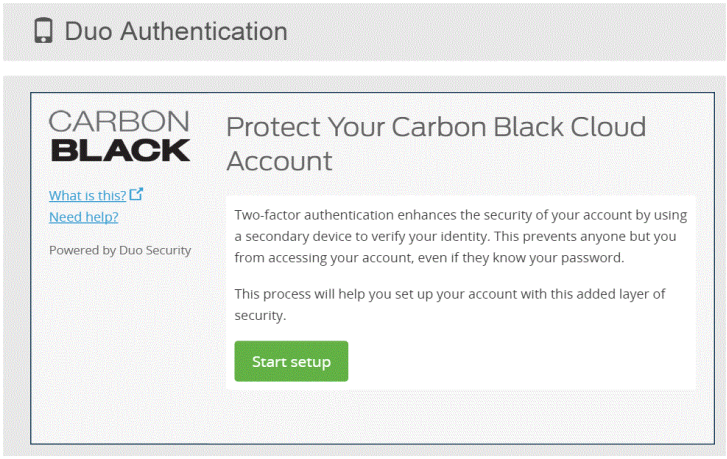

The **What type of device are you adding?** screen appears.

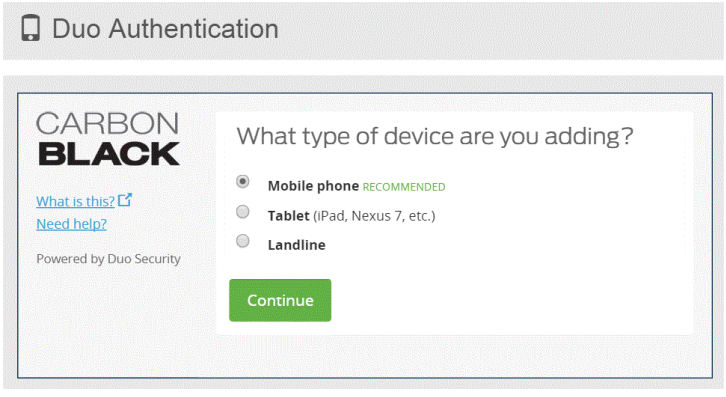

You can add any of several device types for two-factor authentication, including mobile phone devices, tablets, and landlines.

4. Select the type of device you are adding and click **Continue**.

5. Provide your phone number information and click Continue.
6. Select the type of phone device you are using and click Continue.

   This procedure uses an iPhone as an example, but you can add other types of devices.
7. Follow the instructions to install the Duo Mobile application on your device, and then click **I have Duo Mobile installed**.
8. On your phone, tap the **+** button.
9. With your phone, scan the bar code presented in the **Activate Duo Mobile** page:

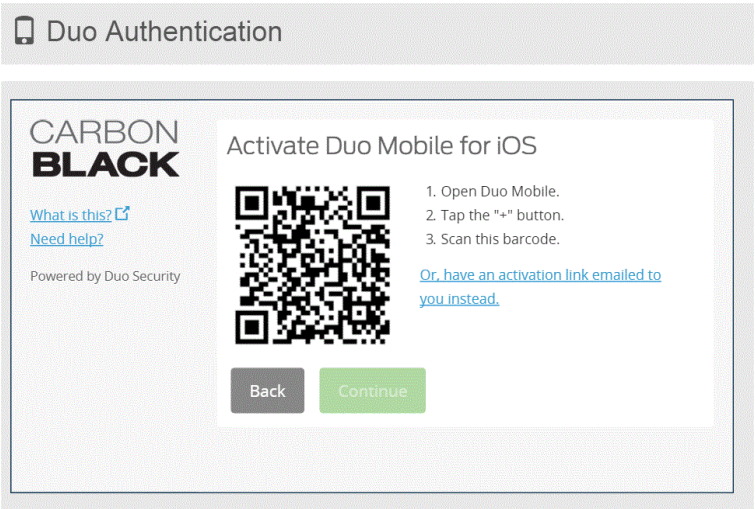

10. When the check mark appears on the bar code indicating success, click Continue.

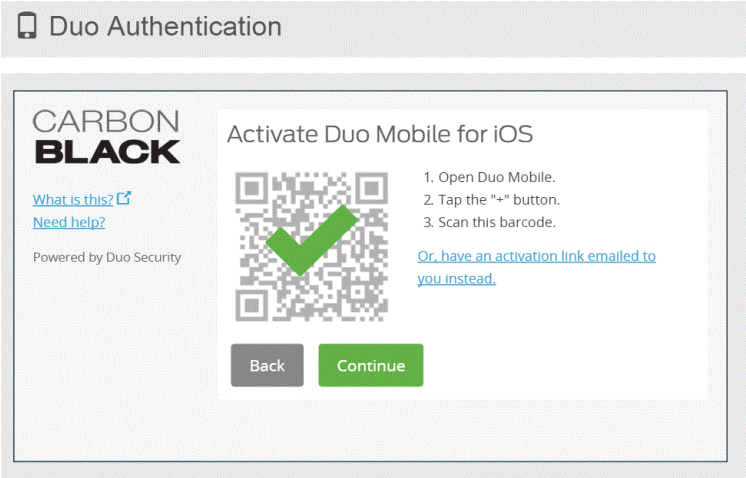

The My Settings & Devices page appears.

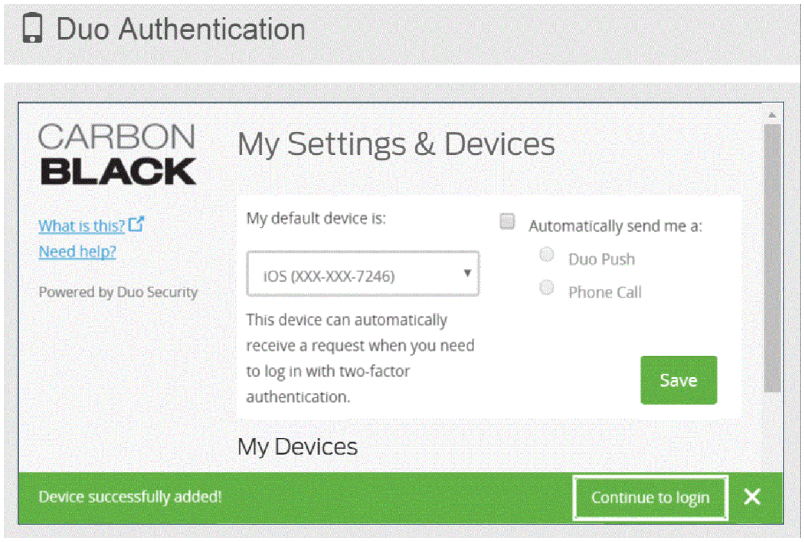

11. On the My Settings & Device page, make the following selections:
    a. From the **My default device is** drop-down list, select your default device. This is useful when you use multiple devices for two-factor authentication.
    b. Select the **Automatically send me a** check box and select either **Duo Push** or **Phone Call** as your preferred communication mode with Duo Mobile.
    c. Click **Save**.
    d. When your device is successfully added, click **Done** (you might need to scroll down to see the **Done** button).
12. On the **Choose an authentication method** page, select one of the following authentication methods:
    - **Send me a Push** – Select this recommended option to receive a Duo push notification to authenticate. Tap **Approve** on the Duo login request that is received on your phone.
    - **Call Me** – Select this option to receive a phone call to authenticate.
    - **Enter a Passcode** – Select this option to enter a Duo Mobile passcode to authenticate. Open the Duo Mobile application on your phone and click the key icon to generate a new passcode.

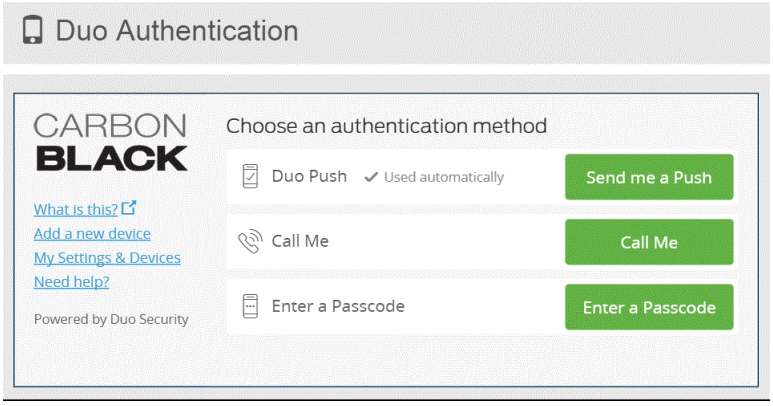

13. When authentication is successful, you are logged into the CB Response Cloud.

    The CB Response Head-Up Display (HUD) page appears.

Logging in After Initial Login (Cloud)

To log into the CB Response console after initial login:

1. In a supported web browser on a computer with access to your server, enter the path to the CB Response Cloud service and in the initial dialog, click **Login with CB Cloud**.

   The CB Response **Login dialog box** appears.

2. Do one of the following:
   - If **Username** and **Password** fields are pre-populated, click **Sign in**.
   - If the fields are not pre-populated, enter your **Username** and **Password** and click **Login**.

   The system responds in one of the following ways, depending on your two-factor authentication selection at initial login:
   - If you selected **No Thanks**, the CB Response **HUD** page appears.
   - If you selected **Maybe Later**, the **Two Factor Authentication** wizard opens, prompting you to enroll. You can either decline, or decide to configure two-factor authentication as described in the procedure "To configure two-factor authentication:" on page 40.
   - If you enabled two-factor authentication, the system contacts your configured device (mobile phone, tablet, or landline). Follow the prompts to authenticate.

   After you authenticate successfully, you are logged into the CB Response Configure Server / Users page.

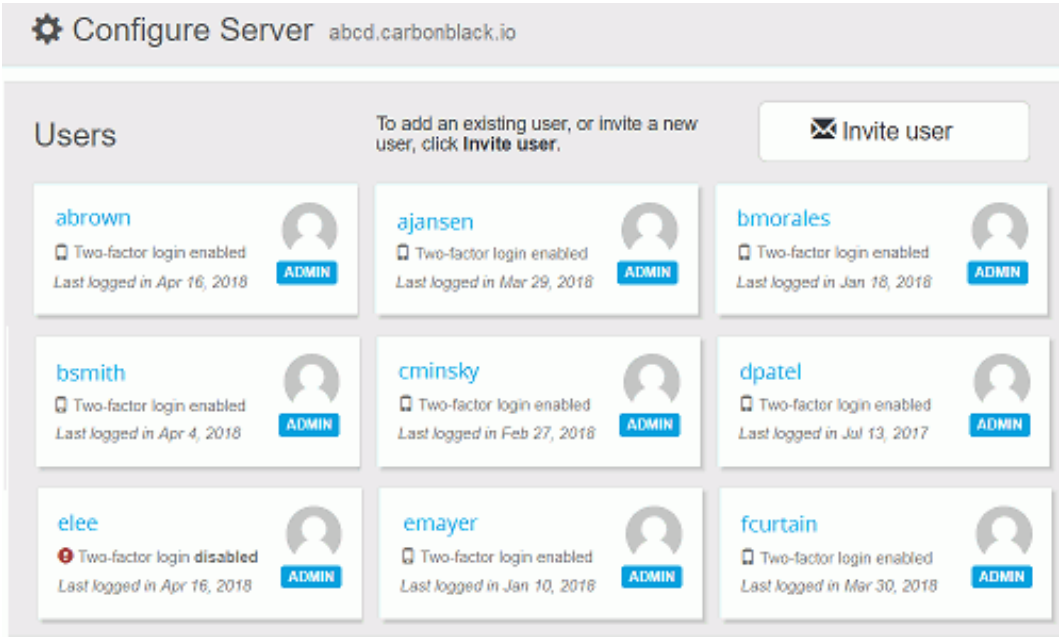

On the Configure Server / Users page, you can do the following:

- Click the Cloud link to go to the My Servers page, where you view a list of all CB Response Cloud servers to which you have authorized access. Click a server link to access the HUD page for that server.
- Click your user name to manage your CB Response Cloud user account details, such as change your password, and enable or disable two-factor authentication.
- Click Invite user to invite new or existing users.
- Click an existing user to manage their account and permissions. For more information, see Chapter 4, “Managing User Accounts for Cloud Servers.”

Enabling/Disabling Two-Factor Authentication (Cloud)

You can enable/disable two-factor authentication for your CB Response Cloud account at any time.

To enable or disable two-factor authentication:

1. From the menu bar, select ***Your Username*** **> Account**.
2. In the **Security** panel of the Account page, you can do one of the following:
   - If enabled, disable two-factor authentication by clicking **Disable Duo 2FA**.
   - If disabled, enable two-factor authentication by clicking **Enroll with Duo 2FA**.
3. Follow the prompts to complete the procedure to either enable or disable Duo 2FA.

   The **Security** panel of the Account page updates to display the changed status for **Two-Factor Authentication**.

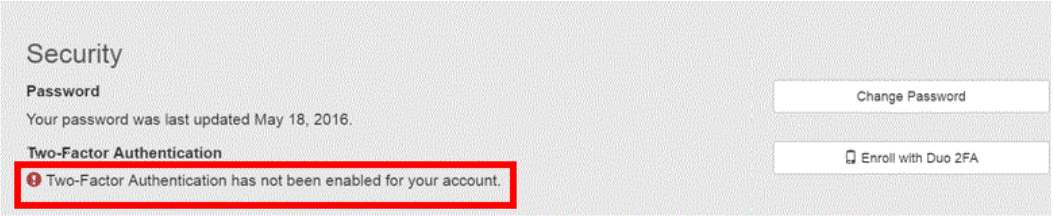

Logging Out (Cloud and On-Premises)

The top-right corner of CB Response console displays your user name. You can click to display a drop-down list from which you can:

- View your profile information
- Logout of CB Response
- For administrators only, view and modify Settings and Shared Settings.

To log out of the CB Response console:

1. Click your user name.

   The user drop-down list appears.
2. Select **Logout** to log out of the console.

CB Response Console Controls

Navigation Bar

In both on-premises and cloud instances, you use the CB Response navigation bar to access other console pages. The following examples show the options available for users with full Administrator/Global Administrator privileges – other users will see options that are appropriate to their privilege level. The **Teams** option appears only for cloud instances.

Links in the navigation bar are as follows:

Link	Description
CB Logo	Opens the Heads Up Display (HUD) page, which is a customizable page that provides a summary of alerts on hosts that report to your CB Response server. See Chapter 20, “Using the Head-Up Display Page”.
Threat Intelligence	Provides intelligence feeds. You can set up watchlists, incremental synchronizations, and full synchronizations with these feeds. You can also access information about process and binary matches found by each feed. See Chapter 15, “Threat Intelligence Feeds.”

Link	Description
Triage Alerts	Shows events that match queries that are defined by watchlists and indicators of compromise (IOCs) that are defined by feeds. The information provides criteria that is available to search for specific events. See “Managing Alerts on the Triage Alerts Page” on page 289.
Watchlists	Saved queries that are performed on process events and binary data stores. The queries contain lists you can use to track specific IOCs. See Chapter 18, “Watchlists.”
Process Search	Provides an overview of the sensor process data collection that is received from currently installed sensors. See Chapter 11, “Process Search and Analysis.”
Binary Search	Shows the metadata of binary files that have been executed. Binary file data is tracked at the moment of execution, and is identified by MD5 hash name. See Chapter 12, “Binary Search and Analysis.”
Go Live	This icon appears if you have enabled CBLR in **Username > Settings > Advanced Settings**. It opens a command line page that provides direct access to sensors. You can directly access content on endpoints that are running the sensors that are providing information. See Chapter 10, “Responding to Endpoint Incidents.”
Investigations	A collection of tagged process events that are products of search results that come from searching your networks and endpoints for threats. See Chapter 17, “Creating and Using Investigations.”
Sensors	Shows data for sensors and sensor groups. Sensor groups categorize sensors that share the same configuration. You can view, define and update sensors and sensor groups on this page. See Chapter 6, “Managing Sensors”.
Users	• (on-premises) Displays the User Management page. Tabs on this page allow CB Response administrators to add and configure new users who can log into the CB Response, view user activity, and create and manage teams of users. Non-administrator users can use this menu item to view their user profile details. See Chapter 3, “Managing User Accounts for On-Premise Servers”. • (cloud) Displays CB Response Cloud user accounts that are authorized to access the server. See Chapter 4, “Managing User Accounts for Cloud Servers”.
Teams	Cloud only. Teams are groups of users. Users can belong to several teams. This link goes to the Team Management page for your cloud instance. Tabs on this page allow CB Response administrators to configure users, view user activity, and create and manage teams of users. See Chapter 4, “Managing User Accounts for Cloud Servers”.
Event Forwarder	Displays only if enabled. This link opens the Event Forwarder Settings page, which lets you configure the Event Forwarder from within the CB Response Console. See Chapter 16, “Configuring the Event Forwarder”.

Link	Description
Server Dashboard	Shows server statistics such as sensor statistics and server communication status. See “Monitoring Sensor and Server Information” on page 91.
Banned Hashes	Opens the Manage Banned Hashes page, which shows process hashes for which a ban has been created. Banned processes are blocked from running on hosts that are managed by a CB Response sensor. See Chapter 10, “Responding to Endpoint Incidents.”

Username Menu

The top right corner of the console shows the name of the currently logged in user. A dropdown menu includes the following options:

Menu Choice	Description
My Profile	For all users, shows and allows editing of the current user’s first and last name and email address. Shows teams of which the user is a member. Provides access to dialogs for changing the user’s password and API token. For cloud users, the My profile page includes a link to the user’s cloud profile (Account) page, which is described in “Changing Security Settings, Email Address or Full Name” on page 71.
Sharing Settings	For administrators only, shows a page that allows administrators to determine whether to share different types of information with Carbon Black and its partners. See “Data Sharing Settings” on page 248.

Menu Choice	Description
Settings	For administrators only, this page includes several tabs that allow administrators to view and change settings that affect the operation of CB Response: • **Sites** – Provides a menu of sites and the ability to throttle sites by time and day of the week. • **E-Mail** – Lets you configure your own alert email server (recommended), use Carbon Black's email server, or not receive alert emails. • **License** – Shows the CB Response server's license and allows you to apply a new license. • **Server Nodes** – Shows all of the server nodes in your cluster, their Node ID, Name, Hostname (with domain) and full URL with port. • **Server Certificates** – Shows all of the sensor-server certificates that are available on the current server. It also shows the validation method that is being used for these certificates. See "Adding Certificates through the Console" **on page 122**. • **Ingress Filters** – Lets you define and view ingress filters. • **CB Protection Server** – Shows configuration information if this CB Response server is integrated with CB Protection. See "Integrating CB Response with CB Protection" in the *CB Response Integration Guide* for more information. • **Advanced Settings** – - **Process Search Settings** – Allows administrators to block certain process searches that could cause significant performance problems in CB Response. See "Process Search Settings in the Console" on page 180 for details. - **EU Data Sharing Banner** – Allows administrators to enable and disable the display of a red banner at the top of console pages that warns users to be cautious when sharing screenshots or other data. Note that this capability can be overriden in the `cb.conf` file by the `ShowGdprBanner` setting. See "EU Data Sharing Banner" on page 49 for details. - **CB Live Response** – Allows administrators to enable or disable Live Response, which opens a command interface for direct access to any connected host running the CB Response sensor. Note that this can be overriden in the cb.conf file by the CbLREnabled setting. This section also includes configuration settings that can be adjusted to help improve Live Response performance. See "Using Live Response" on page 150 for details of this feature.
Logout	Logs the current user out of the CB Response console.

EU Data Sharing Banner

CB Response displays information from all endpoints that report to a server through the sensor. Because some of this information can be sensitive, you might need to take extra steps to avoid exposing it in the wrong places.

As an extra precaution, CB Response provides administrators with the ability to display a red banner at the top of console pages that warns users to be cautious. This banner is enabled and disabled on the **Advanced Settings** tab of the Settings page.

Notifications

The **Notifications** menu at the top of the menu is to the immediate left of the name of the logged in user. It includes a count of new notifications and a dropdown menu that shows the number of each type of notification. Clicking any item on the menu takes you to the page that provides details for the item to which the notification refers.

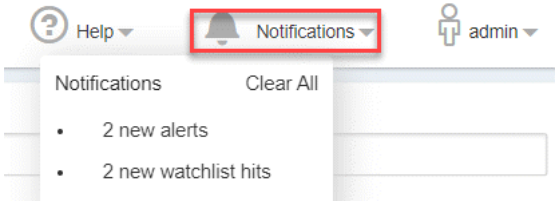

For example, if you click **2 new watchlist hits** in the preceding example, you are navigated to the Watchlists page. When you click into details for all of the new notifications, the counter resets to zero and the menu then displays **No new notifications**. You can also click **Clear All** to clear the menu without viewing notification details.

Help: User Guide and Customer Support

The **Help** menu in the top right area of the console provides access to two sources of assistance for answering questions about CB Response.

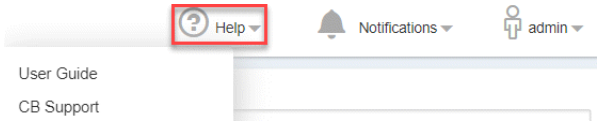

- Click **User Guide** to open a new browser window or tab showing the HTML version of the *CB Response User Guide* (this document). If it displays as a tab, you can drag the

tab off of the current browser to display the User Guide in its own window. The online version of the User Guide is compatible with Chrome browsers, preferably the latest release. You might be able to display help in other browsers but these are not supported and they might have issues that interfere with display or performance.

- The **CB Support** choice opens a new browser window or tab showing the VMware Carbon Black technical support page.

To display online documentation from the console:

1. In the top console menu, click the question mark button and choose **User Guide** from the menu. This opens the home page and table of contents for CB Response Help in a new window or tab. The controls on the help page vary depending on the size of the window, but all pages provide access to the table of contents and search features.
2. To view the table of contents if it is not visible, either expand the browser window to display the contents on the left, or click the **Table of Contents** button.
3. 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 subtopics. Click the down arrow icon to collapse the items below it.
4. 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.

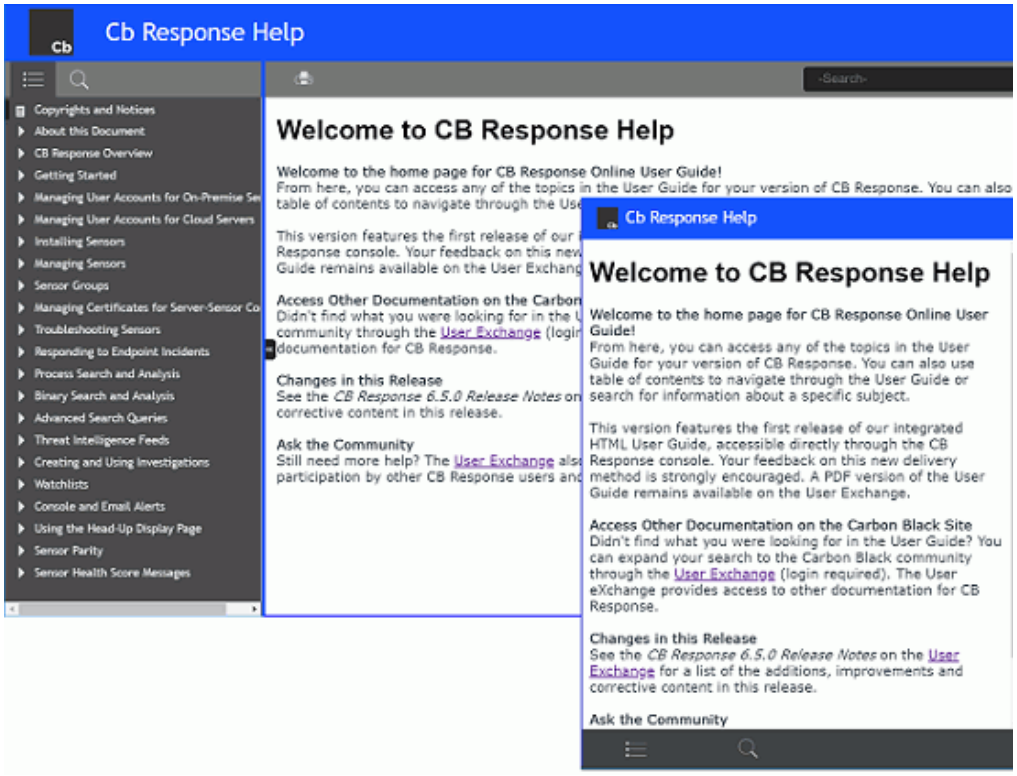

Chapter 3

Managing User Accounts for On-Premise Servers

This chapter explains how to manage user access to the CB Response console for servers that are installed on your premises. This includes managing teams to determine user roles and manage user accounts. For information on managing users for CB Response Cloud, see “Managing User Accounts for Cloud Servers” on page 65.

Sections

Topic	Page
Overview of User Management (On-Premises)	52
Managing User Access with Teams	53
Role-based Privileges for Teams	53
Adding Enhanced Permissions for Analysts	56
Modifying Teams	59
Creating User Accounts for On-Premise Servers	60
Changing Passwords	61
Resetting API Tokens	62
Deleting User Accounts	62
Viewing User Activity	63

Overview of User Management (On-Premises)

The CB Response console is the user interface for access to CB Response features. Each console user logs in to the system with a user name and password. Login accounts provide administrators, analysts, and others who use the console to access the CB Response features appropriate to their role, and also allow administrators to limit unnecessary access to features or sensors.

During the CB Response installation process, a default user account is created and assigned Global Administrator status. This user has full access to all sensors and all features. After you log in by using the default account, you can set up additional users, including other Global Administrators and users with other roles that can vary by sensor group.

The capabilities of a user are determined by the following factors:

- **Is the user a Global Administrator?** – A checkbox on the User Details page determines whether a user has administrator privileges (for all sensor groups and features). If a user is a Global Administrator, the following factors are not relevant.
- **What teams does the user belong to?** – The privileges of users who are not global administrators depend upon the *teams* they belong to. Global administrators can also be assigned to teams, although team membership doesn't affect them unless their administrator status is disabled.
- **What roles do team members have for each Sensor Group?** – Teams specify a *role*, which determines the level of privileges that their members have for each sensor group. There are three roles: **Analyst**, **Viewer**, and **No Access**. Teams can (and usually will) have different roles for different sensor groups. See "Role-based Privileges for Teams" on page 53.
- **What is the highest role for any team this user belongs to?** – Access to some features is not restricted by sensor group, but is still controlled by the roles that are assigned to a team. These features become available to a user if the user is on at least one team that has a high enough role for at least one sensor group.

  This helps control access to features that are not specific to sensor groups, but to which you might want to restrict access. For example, threat feeds, which are not specific to any sensor group, are an important tool for CB Response threat monitoring. If a user is an Analyst on any team, that user can take actions that are available on the Threat Intelligence Feeds page.
- **Is the user an Analyst with enhanced permissions?** – For Analysts, access to especially sensitive features (Live Response, sensor isolation, uninstalling sensors, file banning, and turning tamper detection on and off for a Sensor Group) is controlled by supplemental enhanced permissions. These are added on a per-user basis. See "Adding Enhanced Permissions for Analysts" on page 56.

A table with more specific details about team privileges is displayed in the "Managing User Access with Teams" on page 53.

Important

Creation and management of user accounts and teams is available only to Global Administrators in on-premise installations and by Administrators in cloud installations.

Managing User Access with Teams

If a CB Response user is a Global Administrator (on premises) or Administrator (cloud), that user has access to all functionality for all computers in all sensor groups. A default Global Administrator user is created during on-premises server installation; others can be created later.

For all other users, access to CB Response features is granted through membership on teams. Endpoints running the CB Response sensor are members of *groups*. Each team has a defined role in each sensor group, and this role defines what it can see and do with sensors and their information. In addition to determining privileges for each sensor group, team specifications control access to some features that are not group-specific.

During CB Response installation, a default sensor group (called **Default Group**) is created. You can put all sensors in the Default Group, but to use teams to limit access to certain sensors, we recommend that you create additional sensor groups. See "Sensor Groups" on page 102.

You might want one team to manage endpoints in one region and another team to manage endpoints in another region. Or, you might let all teams manage most endpoints but create a special team to manage the endpoints of your executive staff. You can also create teams that can view but not modify information and settings in CB Response.

If a user is assigned to multiple teams with permission to access the same sensor group and these teams have different rules, the user has the privileges of the highest role that is available from any of the teams.

Although you can assign a user to teams at any time, it is helpful to have teams set up before you create non-global-administrator users because the capabilities of most users are determined by their teams.

Role-based Privileges for Teams

The roles you can assign to a team for each sensor group are:

- **Analyst** – This role allows the user to monitor and respond to suspicious or malicious activity on endpoints in sensor groups for which it has the role.

  Analysts can be given additional, enhanced privileges on a per-user basis so that they can use special features: Live Response, isolation, hash banning, toggling tamper detection, and uninstalling the sensor. See "Adding Enhanced Permissions for Analysts" on page 56.

  Unless they are Global Administrators, Analysts do not have access to data or functions for managing the CB Response server itself, such as managing users and teams, viewing and changing server settings (including sharing settings), and viewing the server dashboard.

- **Viewer** – This role allows the user to access information that CB Response gathers, including suspicious or malicious activity, on endpoints in sensor groups for which it has the role.

  Unless they are Global Administrators, Viewers cannot access CB Live Response (Go Live), investigations, sensor isolation or file banning. They also cannot access CB Response server management functions.

- **No Access** – This role gives the user no access to information or management functions for the specified sensor groups. If the user does not have any higher role for any team, the only page available to them is My profile.

Some access control is applied on the page level – for example, certain pages are only visible to Global administrators or cloud Administrators. In other cases, access control determines the data that appears on a page and the actions that can be taken there. If users enter a URL for a page they do not have permission to view, they are redirected to the HUD page.

The following table provides more detail about privileges and access types that are available for each role.

Analyst & Viewer Access by Feature

Feature or Page	Permissions by Role
Server Dashboard	Only available to Global Administrator (on-premises) or Administrator (cloud)
Sensors	**Viewers:** Can view tables and details of sensors in sensor groups for which the user has Viewer access. **Analyst:** Can perform actions on a sensors in sensor groups for which the user has Analyst access. Additional enhanced user permissions are necessary for isolating and uninstalling sensors and using Live Response. Analysts can also move sensors between sensor groups if they are Analysts for both the source and destination sensor groups.
Sensor Groups	**Viewers:** Can view tables and details of sensor groups for which the user has Viewer access. **Analyst:** Can perform certain actions involving sensor groups for which the user has Analyst access: • Can toggle tamper detection if the user also has enhanced permissions for tamper detection. • Can toggle process banning if the user also has enhanced permissions for process banning. • Can edit other General, Sharing, Advanced, Event Collection, Upgrade Policy settings for the group. An Analyst cannot add or delete a sensor group.
Uninstall Sensors	**Viewers:** No Access **Analyst:** Can uninstall sensors from the console in sensor groups for which the user is an Analyst if the user also has the enhanced permission for uninstalling sensors.
Users, Teams and Activity Audit	Only available to Global Administrator (on-premises) or Administrator (cloud)
Tamper detection toggle	**Viewer:** No Access **Analyst:** Can turn tamper detection on and off for sensor groups for which the user is an Analyst if the user also has the enhanced permission for tamper detection.
HUD page	**Viewer:** Can view the page filtered to show alerts and sensors in sensor groups for which the user is a Viewer. **Analyst:** Can take action on alerts.

Feature or Page	Permissions by Role
Threat Intel Feeds	**Viewer:** No Access **Analyst:** Can view and modify the page, including enabling and disabling actions on hit (Email Me, Create Alert, or Log to Syslog).
Triage Alerts	**Viewer:** Can view all binary alerts, and can view other alerts in sensor groups for which the user is a Viewer. **Analyst:** Can view and take action on all binary alerts; can view and take action on other alerts in sensor groups for which the user is an Analyst.
Watchlists	**Viewer:** Can view watchlist results for binary searches and other searches involving sensor groups for which the user is a Viewer. **Analyst:** In addition to view access, can add, modify, and delete watchlists, and take actions including enabling and disabling email notification, log to Syslog, and alerts.
Process Search	**Viewer and Analyst:** Can view process search results for sensor groups for which the user has at least Viewer access.
Process Analysis	**Viewer:** Can view process analysis results for sensor groups for which the user has at least Viewer access. **Analyst:** Can take actions for processes in sensor groups for which the user is an Analyst if the user also has the enhanced permission for that action. Actions include Isolate host, Go Live, and Ban Hash.
Binary Search (results) **& Analysis** (details)	**Viewer:** Can view all binary search results on the Search Binaries page and also details about one binary (Binary Analysis), regardless of the sensor group of the binary instance. **Analyst:** ICan ban hashes in the search results if the user also has the enhanced permission to Ban hashes.
CB Live Response	**Viewer:** No Access. **Analyst:** Can use Live Response to access and take actions on the endpoints in sensor groups for which the user is an Analyst if the user also has the enhanced permission for Live Response.
Investigations	**Viewer:** Can view the Investigations page. Actions are limited to Export events to CSV and Export timeline to PNG. **Analyst:** Can create, delete, and modify investigations.
Isolation	**Viewer:** No Access. **Analyst:** Can isolate endpoints and restore them from isolation in sensor groups for which the user is an Analyst if the user also has the enhanced permission for isolating sensors.

Feature or Page	Permissions by Role
Banned Hashes	**Viewer:** No Access. **Analyst:** Can ban hashes and remove bans if the user has the enhanced permission for banning hashes. Not restricted by sensor group.
Notifications	**Viewer and Analyst:** All users can view notifications on the **Notifications** menu and receive notification emails.
User Name Menu	
Sharing Settings	Only available to Global Administrator (on-premises) or Administrator (cloud).
Settings	Only available to Global Administrator (on-premises) or Administrator (cloud).
Profile info	All users can view and edit their own profile.

Adding Enhanced Permissions for Analysts

The Analyst role allows access to features for monitoring and investigation of suspicious or malicious activity on endpoints. You might allow some Analysts to take certain actions to remediate threats or vulnerabilities. CB Response provides an interface for adding special permissions to Analysts on a per-user basis.

When enabled, these enhanced features allow a user to take action in sensor groups where the user is on a team with Analyst privileges:

Enhanced Permission	Description
Ban hashes	Can ban files by hash and remove bans. These bans are applied to all sensors.
Isolate sensor	Can isolate a sensor in that group from the network and restore the sensor from isolation.
Live Response	Can connect to and act on a sensor in that group using Live Response.
Tamper detection	Can disable and enable reporting of tamper events for all sensors in that group.
Uninstall sensors	Can use the console to uninstall a CB Response sensor in the group.

Notes

You can add enhanced Analyst permissions to any user, but these permissions are unnecessary for a Global Administrator or cloud Administrator and have no effect on users who are not on a team with the Analyst role in at least one sensor group.

The following procedures instruct you to navigate to **Users** from the navigation bar; for cloud administrators, navigate to **Teams** instead.

To provide enhanced Analyst permissions to a user:

1. On the navigation bar:
   - For an on-premises server, click **Users**.
   - For a cloud instance, click **Teams** and then click the **Users** tab.
2. Locate the user to whom to give enhanced permissions and click the **Edit user** button.

   If you are providing enhanced Analyst permissions for an on-premises user you have not created yet, use the **Add User** button and provide all of the necessary information as described in "Creating User Accounts for On-Premise Servers" on page 60.
3. In the **Enhance Analyst permissions** panel, check the box next to each permission to give this user.

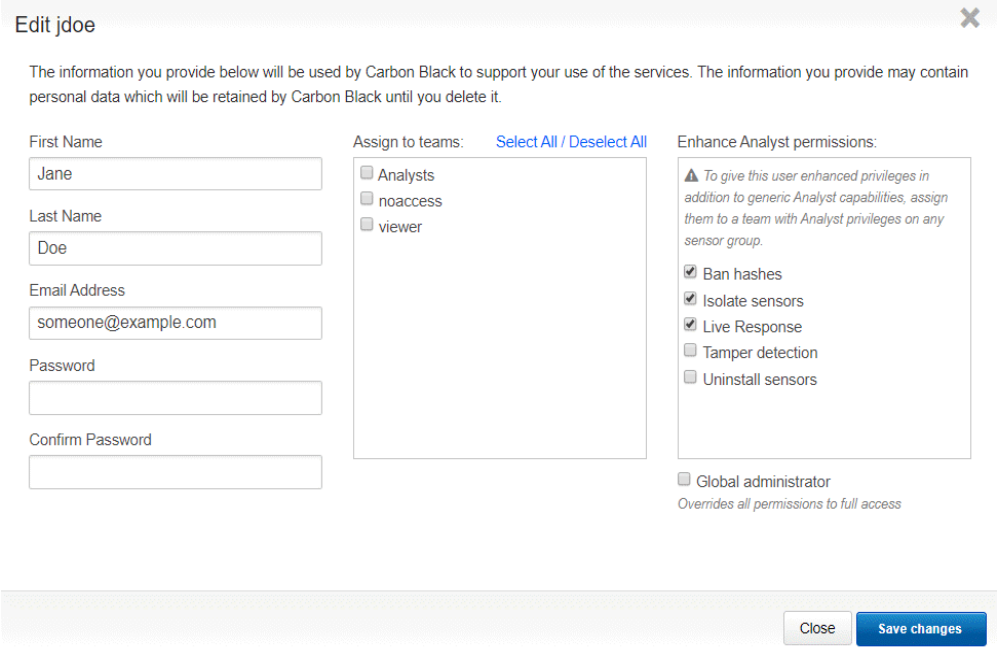

   If the user is not a member of an Analyst team, a gray triangle icon appears in the upper left of the **Enhance Analyst permission** panel. If the user is already a member of a team with Analyst permission for a Sensor Group, the icon is a green checkmark.
4. If necessary, add the user to a team with Analyst permission.
5. Click **Save changes**.

User/Team Permissions Example

The following scenario provides an example of how you can set up user accounts and teams. This is an simplified example — not a recommendation.

Suppose that a division of your company is based in Europe, with sites in France, Germany, and Italy. Also assume that all endpoints in these countries have sensors that are managed by one CB Response cluster.

- **Create Administrators**: You can make two Global Administrators in each country, so that they can set up their users and user teams. These Global Administrators can also monitor system performance and change settings that control the server. Although in this example, their primary responsibility is for the sensors in their own country, each Global Administrator can perform any CB Response activity on the any country's endpoints if necessary.
- **Create country-specific Analysts:** Create four additional users in each country that are assigned as Analysts to teams for the sensor groups that correspond to the endpoints in their own country. They can monitor the data from the sensors in these groups.
- **Enhance permissions for some Analysts:** For two of the Analysts in each country, add enhanced permissions that allow them to take actions that affect sensors.
- **Let Analysts be Viewers for other sensor groups:** So that Analysts can be aware of activities or trends that could affect all countries, give them the Viewer role for the sensor groups in countries in which they are not Analysts.
- **Create Viewer (only) users:** There might be a third, larger group of users that you assign to teams that make them Viewers for the sensor groups in their own country so they can monitor but not alter the sensors in those groups.

Creating Teams

To create teams:

1. On the navigation bar:
   - For an on-premises server, click **Users** and then click **Teams**.
   - For a cloud instance, click **Teams** and then click **Teams**.

2. Click the **Create Team** button to display the **Add Team Settings** page:

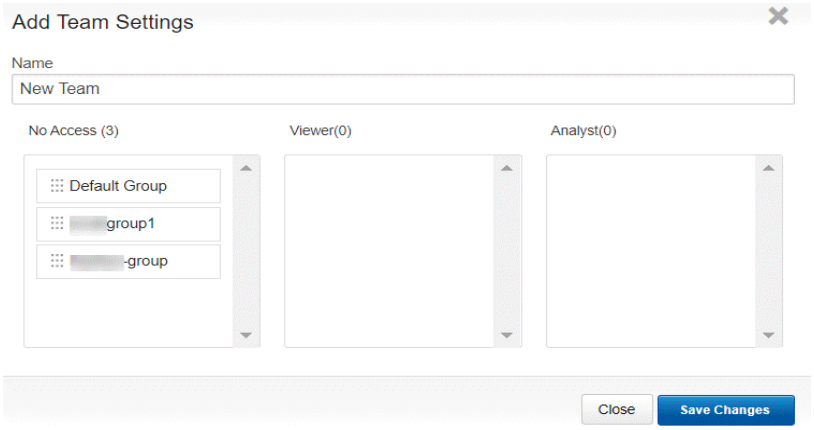

3. In the **Name** field, enter a name for the team.
4. Drag and drop the sensor groups with the appropriate permissions. For example, for this team to have no access to the sensor group named **Default Group**, you would drag the **Default Group** box to the **No Access** list.

   You could assign roles to users by adding them to teams that are set up with the type of privileges that are appropriate for the role (Analyst, Viewer, or No Access).
5. Click **Save Changes**.

Modifying Teams

To modify a CB Response team:

1. On the navigation bar:
   - For an on-premises server, click **Users** and then click **Teams**.
   - For a cloud instance, click **Teams** and then click **Teams**.
2. In the list of teams, click the **Edit** icon to the far right of the team name:
3. In the **Edit Team Settings** page, modify the team settings as needed and then click **Save Changes**.

Deleting Teams

When you delete a team, references to the team in user accounts are deleted, but the user accounts remain active.

To delete a CB Response team:

1. On the navigation bar:
   - For an on-premises server, click **Users** and then click **Teams**.
   - For a cloud instance, click **Teams** and then click **Teams**.
2. In the list of teams, click the **Delete (x)** icon to the far right of the team name.
3. Click **OK** to confirm the deletion.

Creating User Accounts for On-Premise Servers

Although you can assign a user to teams later, it is helpful to have teams set up before you create non-global-administrator users because the capabilities of most users are determine by their teams. See "Managing User Access with Teams" on page 53.

To create user accounts for on-premise servers, log in to the CB Response server console by using an account that has Global Administrator status. If no other users have been created yet, use the administrative account and password that were set up in the `cbinit` script during the server installation process.

To create an on-premises user account:

1. From a supported web browser, enter the path to your CB Response server.

   ```
   https://<your CB Response server address>/
   ```

2. Enter the username and password for a Global Administrator account.
3. On the navigation bar, click **Users**.
4. Click **Add User** in the top-right corner.

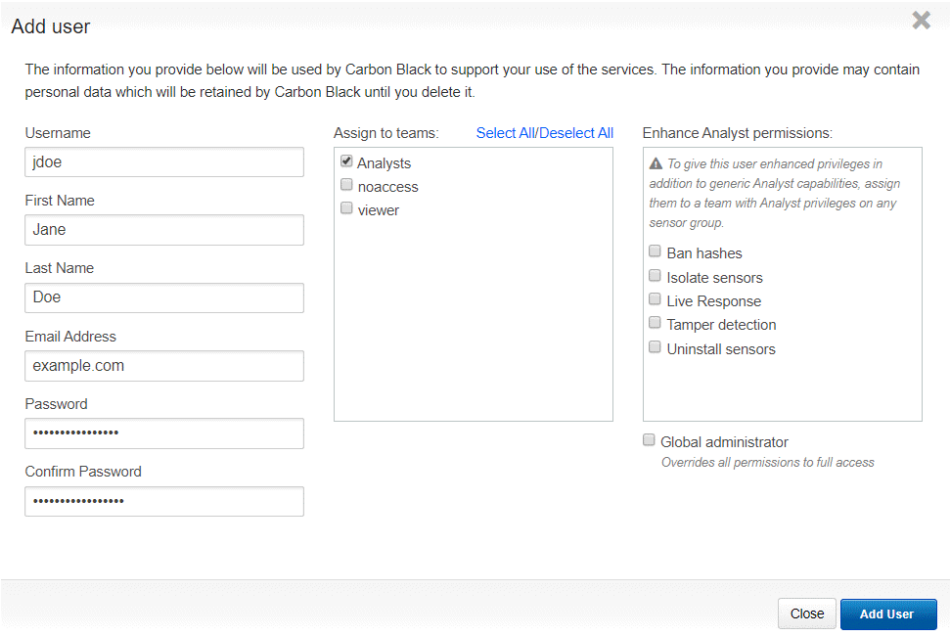

5. Enter the following information:

Field	Description
Username	Name that the user enters to login to the CB Response console. User names are case sensitive and restricted to standard Latin alphanumeric characters. Symbols and punctuation characters are not allowed. If you try to create a user account that contains an illegal character, the console displays a warning message.
First Name	First name of the user.

Field	Description
Last Name	Last name of the user.
Email address	Email address for the user.
Password	Password that authenticates this user. Enter any combination of letters, numbers, or special characters. Passwords are case sensitive. This field changes to **New Password** when you are editing existing accounts.
Confirm Password	Retype the password for confirmation.
Assign to teams	Select the teams to which the user will belong. The default team is Analysts. Users can belong to more than one team. See "Managing User Access with Teams" on page 53.
Enhance Analyst permissions	For a user that is an Analyst on any team, check one or more boxes to give the user permission to use additional features. See "Adding Enhanced Permissions for Analysts" on page 56.
Global Administrator	Check this box to give the user Global Administrator privileges. **Important:** A Global Administrator has full access to all CB Response features and data, including server management and response tools. This includes access to every endpoint with an active sensor, without needing to be assigned to teams.

6. Click **Save changes**.

Changing Passwords

Carbon Black recommends that new users change their passwords after logging in for the first time.

To change your password:

1. Login to the CB Response console.
2. Click ***Username*** **> My Profile**.
3. Click **Change Password**. Enter the current password, the new password, and then confirm the new password.
4. Click **Save changes**.

Note

Global Administrators can change the password of any user through the User Management page.

Resetting API Tokens

CB Response has RESTful APIs to create custom scripts for interactions. These are described at https://developer.carbonblack.com/reference/enterprise-response/.

An unique API token is assigned to each CB Response user. It serves as the key authentication mechanism when making calls to the APIs. Users can reset their own API token at any time.

> **Note**
>
> When a user's API token is reset, any affected custom scripts or integrations that use the API token must also be updated.

To reset the API token for a user account:

1. Login to the CB Response console with the account whose API token will be changed.
2. Click ***Username*** **> My Profile**.
3. Click **API Token**.
4. Click the **Reset API Token** button.

   A notification briefly appears notifying you that the API token has been reset.

Deleting User Accounts

A user account can be removed from the system when that user no longer needs access to the CB Response console or leaves the organization. Users who have Global Administrator privileges can delete any account except their own and the built-in administrator account. If the user with the deleted account belongs to a team, the user is automatically removed from the team when the account is deleted.

> **Note**
>
> A Global Administrator can delete any account except the one you are logged in as, including the administration account that was created during the server installation.

To delete a user account:

1. On the navigation bar, click **Users**.
2. Locate the user's name and click the **Delete (x)** icon to the far right of the user name. (Note that the delete icon next to the currently logged in user is grayed out,because this user cannot be deleted while logged in.)
3. Click **OK** to confirm the deletion.

   A popup message appears to report on the success or failure of this deletion. The Activity Audit for this user also shows the delete action.

Viewing User Activity

CB Response keeps an audit trail of user activity.

To view user activity:

From the navigation bar, click **Users** and then click **Activity Audit**.

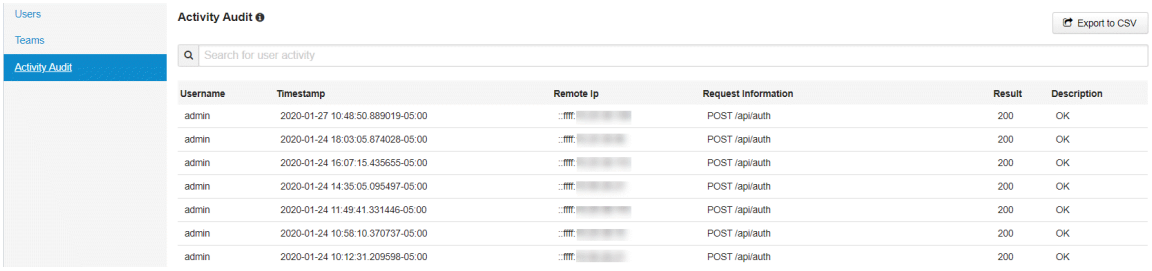

The following information is displayed:

Field	Description
Username	The user name of the user who accessed the console.
Timestamp	The date and time that the user logged in.
Remote IP	The IP address of the computer on which the user logged in.
Request Information	The request (POST, GET, DELETE, etc.) being sent to the server.
Result	The HTTP response code when the user accesses a resource. For example, a successful authentication shows an HTTP 200 code response. A request to access a resource to which the user does not have permission usually results in redirection to the HUD page or displays an HTTP 403 code.
Description	The HTTP response description. For example, an HTTP 200 response shows **OK**, while an HTTP 403 response shows a **Requires Authentication** response.

1. Click **Export to CSV** to export the activity results in a CSV format with the filename `UserActivity.csv`.

> **Note**
>
> If you have access to the CB Response server, you can directly view the log for user activity in the following file:
>
> `/var/log/cb/coreservices/debug.log`

User Activity API Audit Logging

You can enable API audit logging for a server by setting `EnableExtendedApiAuditLogging=True` in the `cb.conf` configuration file (see *CB Response cb.conf Guide*). In this case, CB Response logs all REST API requests from either the console or other sources (such as scripts). API audit log information is stored in the `/var/log/cb/audit/useractivity.log` file, and also appears as follows:

- In the **User Management** section of the CB Response console, under **Request Information** on the **Activity Audit** tab.
- In a CSV file downloaded from the **Activity Audit** tab, as in the following example:

```
2017-12-22 11:30:54:  username='bill' userid='1'
ip='::ffff:111.111.1.1' status='200' method='GET' path='/api/
v2/sensor'
2017-12-22 11:30:54:  username='bill' userid='1'
ip='::ffff:111.111.11.1' status='200' method='GET' path='/api/
v1/alert'
2017-12-22 11:30:55:  username='bill' userid='1'
ip='::ffff:111.111.11.1' status='200' method='GET' path='/api/
v1/detect/report/currentmonitoringstatus'
```

Chapter 4

Managing User Accounts for Cloud Servers

This chapter describes how to manage CB Response Cloud user accounts. In addition to the information in this chapter, the capabilities of cloud users are also affected by the user and team configurations that are described in “Managing User Access with Teams” on page 53.

For information on managing user accounts for an on-premises installation, see “Managing User Accounts for On-Premise Servers” on page 51.

Sections

Topic	Page
Overview of Cloud User Management	66
Creating Cloud User Accounts	67
Accessing Authorized Servers	70
User Account Lockout	70
Viewing and Modifying Cloud User Accounts	70
Changing Security Settings, Email Address or Full Name	71
Changing Administrator / User Status	72
Resetting API Tokens	74
Removing a User Account	74

Overview of Cloud User Management

CB Response Cloud users access their server console using a CB Response Cloud account. User accounts provide system management professionals, threat responders, and other console users the ability to access and manage CB Response features.

Cloud accounts are initiated when an administrator sends an email invitation to a new user, who can then respond to the invitation and create the account. From a cloud account, users can access one or more servers for which they have been authorized. In CB Response Cloud, there are not separate accounts created for each authorized server.

The capabilities of a cloud user are determined by the following factors:

- **Which cloud servers has the user authorized on?** – The administrator who sends out an account invitation is inviting the user both to create a cloud account (if they don't have one) and authorize that account for a particular server.
- **Is the user an Administrator?** – The administrator who sends out an account invitation determines whether the new user will have administrator privileges. This can be changed later. If a user is an Administrator, the next three factors are not relevant.
- **What teams does the user belong to?** – The privileges of users who are not administrators depend upon the teams to which they belong. Administrators can also be assigned to teams, although team membership doesn't affect them unless their administrator status is disabled.
- **What roles do team members have for each sensor group?** – Teams specify a **role**, which determines the level of privileges their members have, for each **sensor group**. There are three roles: **Analyst**, **Viewer**, and **No Access**. Teams can (and usually will) have different roles for different sensor groups.
- **What is the highest role for any team this user belongs to?** – Access to some features is not restricted by sensor group but is still controlled by the roles assigned to a team. These features become available to a user if the user is on at least one team that has a high enough role for at least one sensor group.

  This helps control access to features that are not specific to sensor groups but to which you can restrict access. For example, threat feeds, which are not specific to any sensor group, are an important tool for CB Response threat monitoring. If a user is an Analyst on any team, that user can take any of the actions available on the Threat Intelligence Feeds page.
- **Is the user an Analyst with enhanced permissions?** – For Analysts, access to sensitive features (Live Response, sensor isolation, uninstalling sensors, file banning, and turning tamper detection on and off for a sensor group) is controlled by supplemental enhanced permissions.

See "Managing User Access with Teams" on page 53 for information about creating and using teams.

Important

Creation and management of user accounts and teams is available only to Administrators in cloud installations.

Creating Cloud User Accounts

CB Response Cloud users are assigned to one of two classes when their account is created:

- **Administrator** – Administrators have full privileges on the CB Response Cloud server, including adding and removing other users.
- **User** – Users can access non-administrative functions of the CB Response Cloud server with the type of access determined by their team membership.

An administrator can authorize user access to a server in one of two ways:

- By inviting a new user through email.
- By inviting an existing CB Response Cloud user to become authorized on a new server

User invitations are created from the cloud users page for your cloud server(s).

Inviting a New or Existing User to Access a Cloud Server

This section explains how to invite a new user to create an account with access to a particular server, or authorize an existing CB Response Cloud user to access a server they don't currently have access to.

To invite a user to open a CB Response Cloud account or extend it to a new server:

1. In a browser, enter the URL for the CB Response Cloud server and log in as an Administrator.

   If no other administrator accounts have been created, use the administrator account that Carbon Black provided when you initiated your use of the CB Response cloud.

2. On the navigation bar, click **Users**.

   The **Users** page appears and shows the CB Response Cloud user accounts that are authorized to access the server. This page also shows users whose invitation has expired without their account being activated.

3. If you find that the user you want to invite is listed on the page but the box for user account is grayed out, a previous email invitation to register for this CB Response Cloud server has expired. You have three options:
   - Re-invite the user by double-clicking the account and clicking the **Resend Invite** button.
   - Remove the invitation (making any links in the email sent to this user candidate unusable) by double-clicking the account box and clicking the **Revoke Invite** button.
   - Leave the user's invitation in the expired state and decide what to do with it later.
4. If the user you want to invite does not appear on the page, click **Invite user**.

The **Invite user** dialog box appears.

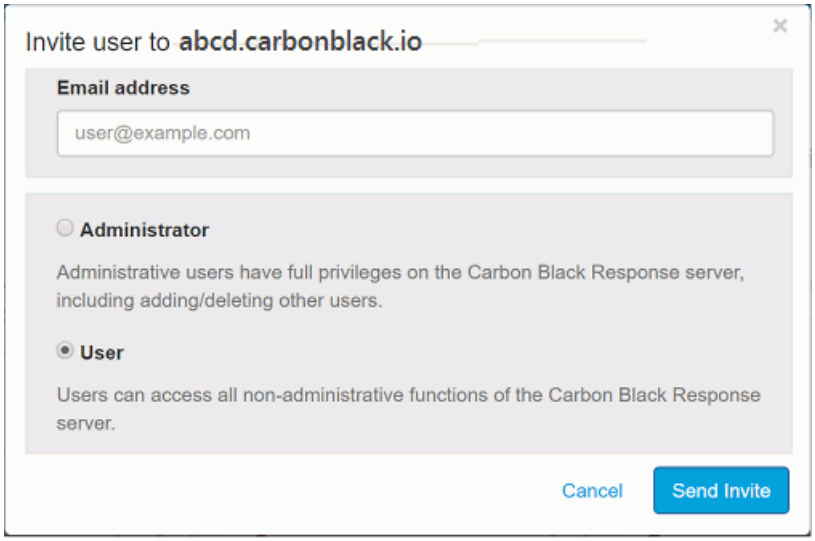

5. Type the email address where the invitation to the new user should be sent.
6. Select **Administrator** or **User**.

   Keep in mind that the Administrator role gives the user full privileges on this CB Response Cloud server. Administrators can add and delete other users.
7. Click **Send Invite**.

An email is sent to the user that contains a link to create a new account on the server. This link also authorizes an existing user who has access to a different server to log into this server by using the same account.

> **Note**
>
> The email invitation link expires after seven days of no activity.

Activating an Account from an Invitation

The invitation to activate a CB Response Cloud account arrives in an email that is sent to the address that is provided when an administrator creates the invitation.

To activate a new account from an invitation:

1. In the invitation email, click on the link to the CB Response Cloud.

   This opens an account creation dialog.

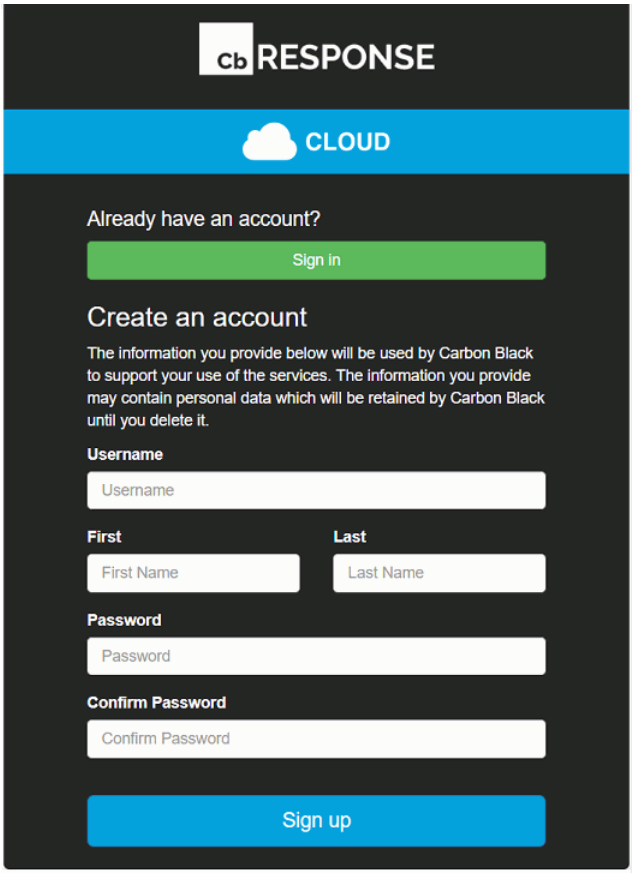

2. In the **Create an account** section, choose a username for your cloud account. User names are restricted to standard Latin alphanumeric characters without symbols and punctuation characters. Then enter your first and last names.
3. Click the **Sign up** button.

   You are immediately logged in to the HUD page for the server you were invited to access. Confirmation that the account was created is sent to the same email address that received the initial invitation.

To activate access to a new server from an existing account:

1. In the invitation email, click the link to the CB Response Cloud.

   This opens an account creation dialog.
2. Under the **Already have an account?** label, click the green **Sign In** button and log in with your existing account name and password.

   You are immediately logged in to the HUD page for the new server.

Accessing Authorized Servers

Using the **Cloud** option (in the User Management console menu), you can view the servers to which your account has access.

To view the CB Response Cloud servers to which you have access:

1. After logging in with your cloud account, click **Users** in the navigation bar.
2. If you have servers in different regions, choose the region (for example, U.S.) for which you want to view servers.
3. In the header area of the Users page, click the **Cloud** link.

   The **My Servers** page appears that lists the servers to which you have access.
4. Click a servername link to access the HUD page for that server.
5. Click the **Configure** button to manage users on that server.

User Account Lockout

To protect against brute force login attacks, CB Response Cloud locks a user account after seven consecutive, unsuccessful login tries in a period of 15 minutes.

An account can remain locked for up to 15 minutes after the lockout begins. Attempts to log in during the lockout period, with or without the correct credentials, have no effect.

Unlocking an Account

An account unlocks automatically after remaining locked for 15 minutes. However, a user can unlock an account before the lockout expires by clicking **Forgot your password?** and following prompts to reset the account password.

Note

While Carbon Black does not recommend using a group email address for a CB Response Cloud user account, keep in mind that for such an account:

- Any person in the group can lock the account with too many failed login attempts. In that case, none of the group members can log in during the lockout period.
- If someone unlocks the account by changing the password, all other group members must be informed of the password change.

Viewing and Modifying Cloud User Accounts

There are several places in the CB Response console in which CB Response Cloud user information can be viewed or changed:

- **Cloud: Users page** – The Users page for cloud accounts shows each cloud account holder, their last login, whether they have two-factor authentication enabled, and their top-level account status (User or Administrator). You can get to this page by clicking **Users** on the navigation bar.

  - Click the tile for any user to show the "User permissions for <*servername*>" page, which allows an administrator to change the user from an Administrator to User, or vice versa. See "Changing Administrator / User Status" on page 72.
  - An Administrator can also remove a user account via this page. See "Removing a User Account" on page 74.

- **Cloud: Account page** – On the Account page, individual CB Response Cloud users can manage their own account details, including:
  - Resetting their passwords
  - Enabling or disabling two-factor authorization
  - Changing the email address associated with the account
  - Editing their first and last names

  The Account page is accessible by clicking **View profile on carbonblack.io** from the My profile page on the server, or by selecting **Account** from the **User name** menu in the cloud view.

  See "Changing Security Settings, Email Address or Full Name" on page 71 for more details.
- **Server: Team Management Users page** – The Users view on the Team Management page shows a table of all users on the current server and their teams. The **View Details** button next to a user name opens an Edit <*user*> page, where you can add or delete the user from teams. See "Managing User Access with Teams" on page 53 for more details.
- **Server: My profile page** – For the currently logged-in user, the My profile page shows the user's teams, provides access to the API Token page where the API Token can be changed, and includes a **View profile on carbonblack.io** button that opens a new browser window that shows the Account page for this user on the CB Response Cloud.

  You can get to the My profile page through the **User name** menu in the top right of the CB Response console.

Changing Security Settings, Email Address or Full Name

Using the **Account** page, CB Response Cloud users can manage their account details, including resetting their passwords and enabling or disabling two-factor authorization. Users can also change the email address associated with an account and edit their first and last names.

To change cloud account details:

1. On the naviation bar, click **Users**.
2. In the top right of the My Servers page, click ***Username*** **> Account**.

Your cloud user **Account** page appears:

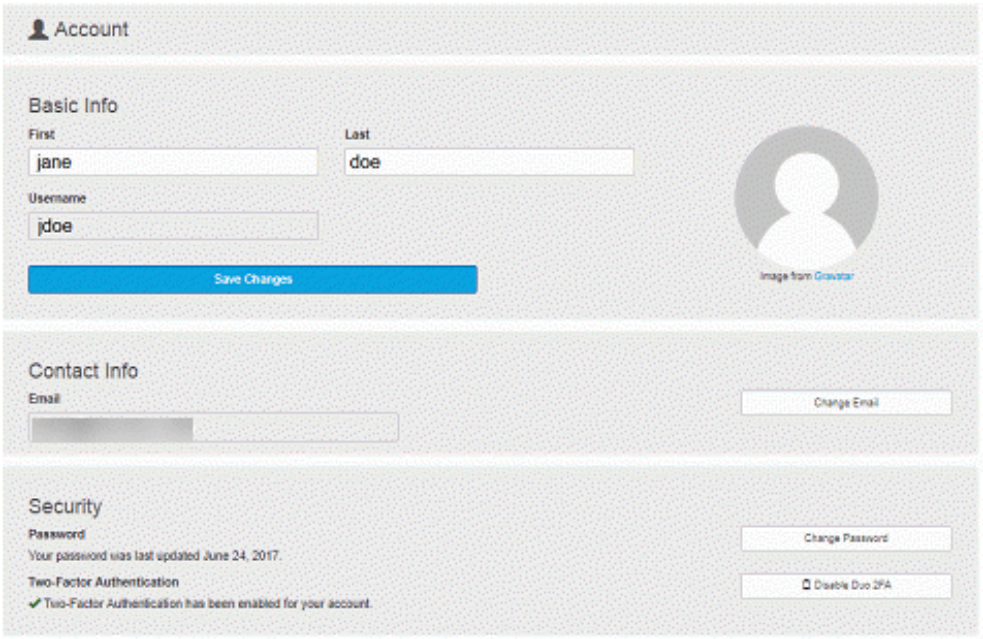

3. In the **Basic info** section, you can modify your first name, your last name, and you can upload an image. You cannot change your username.
4. Click **Save Changes**.
5. In the **Contact info** section, you can change the email address that is associated with your account by clicking **Change Email**.

   In the **Change Email Address** dialog box, enter the new email address and click **Change email**. You will receive an email notification to verify the new email address.
6. In the **Security** section, you can do the following:
   - **Change password** – Click to display the **Change Password** dialog box where you can change your password by entering it twice and clicking **Change password**.
   - **Enroll/Disable Duo 2FA** – Click this to enable or disable two-factor authentication. For more information on enabling two-factor authentication, see “Logging In and Configuring Two-Factor Authentication (Cloud Only)” on page 38.

Changing Administrator / User Status

You can change a user’s status to either an administrator or a non-administrative user in CB Response Cloud. Only administrators can perform this task.

To add or remove administrator status for a user:

1. Log into a CB Response Cloud server as an administrator.
2. On the navigation bar, click **Users**.

   The **Users** page appears and shows the CB Response Cloud user accounts that are already authorized to access the server.

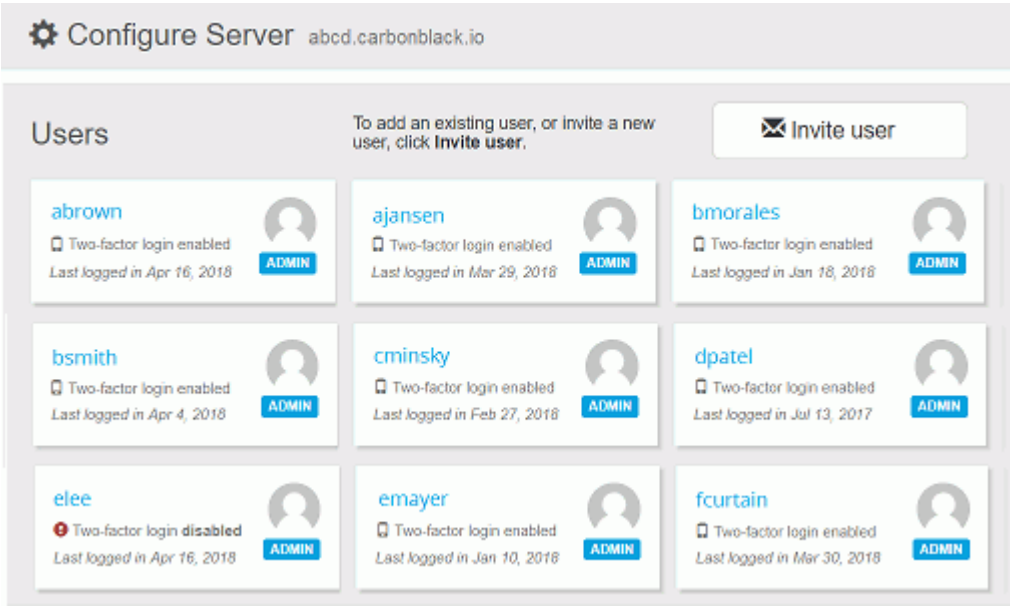

3. Click the user to modify.

   The **User permissions for <*server name*>** page appears for that user:

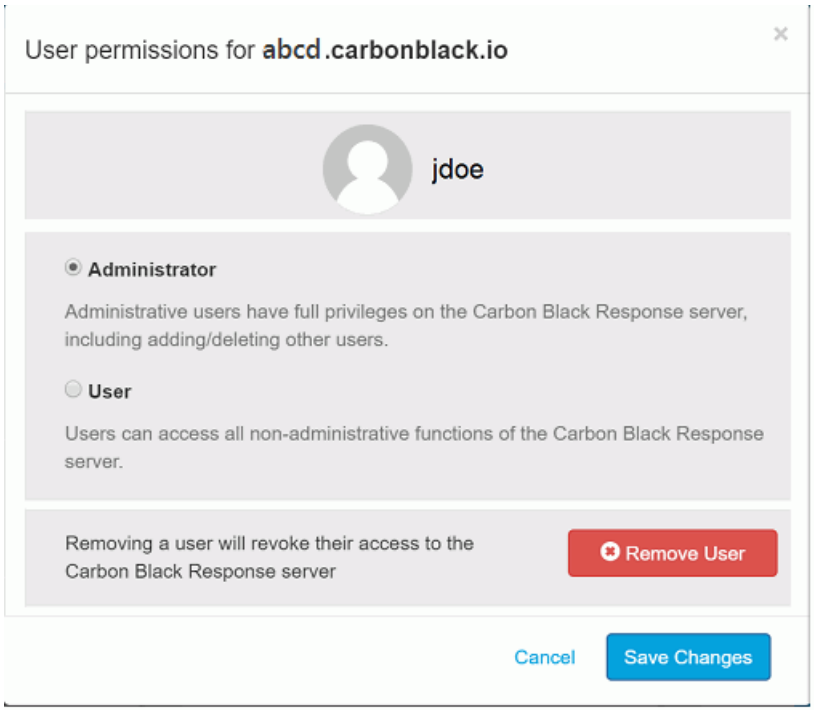

4. To change the user's status, select one of the following options:

- **Administrator** – Administrators are users with full privileges on the CB Response Cloud server, including adding/removing other users.
- **User** – Users that are not administrators can access all non-administrative functions of the CB Response Cloud server.

5. Click **Save Changes**. The system saves the change and returns you to the **Users** page.

Resetting API Tokens

CB Response has RESTful APIs that can be used to create custom scripts for interactions with its features. These are described at https://developer.carbonblack.com/reference/enterprise-response/.

A unique API token is assigned to each CB Response user. It serves as the key authentication mechanism for making calls to the APIs. A user can reset their own API token at any time.

> **Note**
>
> When a user's API token is reset, any affected custom scripts or integrations that use the API token must also be updated.

To reset the API token for a user account:

1. Login to the CB Response console with the account whose API token will be changed.
2. Click ***Username*** **> My Profile**.
3. In the **My Profile** window, select **API Token**.
4. Click the **Reset API Token** button.

   A notification appears briefly in the top-right corner of the console, notifying you that the API token has been reset.

Viewing User Activity

CB Response Cloud keeps an audit trail of user activity. See "Viewing User Activity" on page 63.

Removing a User Account

You can remove a user account from accessing the CB Response Cloud server, thereby terminating access for that account.

To remove a CB Response Cloud user account:

1. On the navigation bar, click **Users**.

   The **Users** page displays the CB Response Cloud user accounts that are authorized to access the server.

2. Click the user to remove.
   The **User permissions for <*server name*>** page appears for that user:

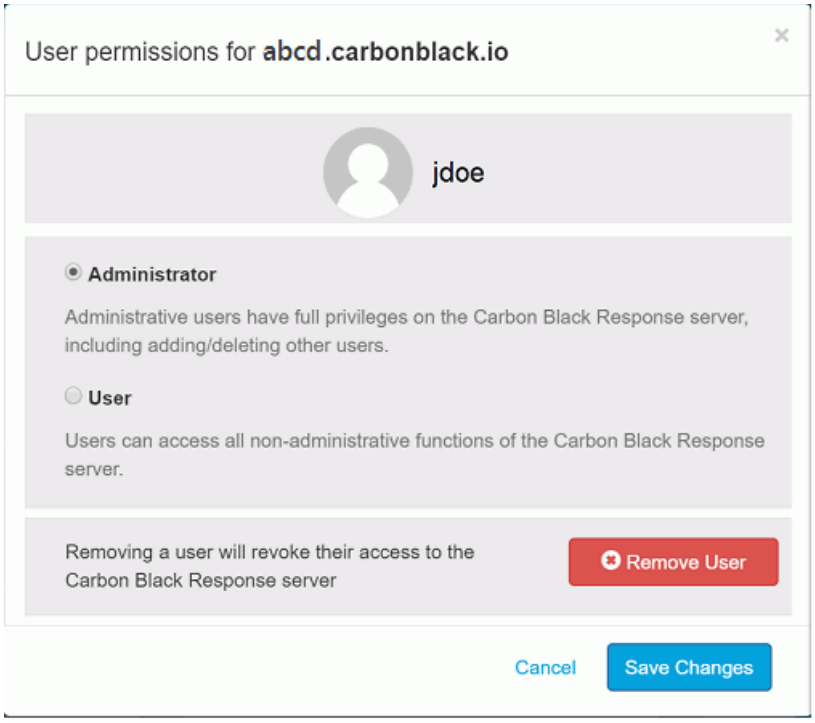

3. Click **Remove User.**
   When the user account has been removed, you are returned to the **Users** page.

Chapter 5

Installing Sensors

This chapter describes how to install and upgrade sensors on Windows, macOS, and Linux endpoints.

- See Chapter 6, 'Managing Sensors' for information on managing sensors.
- See Chapter 7, 'Sensor Groups' for information on managing sensor groups.
- See Chapter 8, 'Managing Certificates for Server-Sensor Communication' for information about certificate options.
- See Chapter 9, "Troubleshooting Sensors," for information on troubleshooting sensors.
- See Appendix A, "Sensor Parity," for information about which CB Response features are supported on sensors that are running in each operating system.

Sections

Topic	Page
Overview of Sensor Installation	77
Supported Operating Systems and Versions	77
Installing Sensors on Windows	77
Upgrading Sensors	83
Uninstalling Windows Sensors	79
Installing Sensors on macOS Systems	80
Upgrading Sensors on macOS	81
Uninstalling Sensors on macOS	81
Installing Sensors on Linux Systems	81
Upgrading Sensors on Linux	82
Uninstalling Sensors on Linux	82

Overview of Sensor Installation

CB Response provides lightweight sensors for installation on network endpoints, such as laptops, desktops and servers. You install a sensor on each endpoint in your enterprise. After installation, sensors gather event data on the endpoints and securely deliver it to the CB Response server for storage and indexing. You can use the default “legacy” certificate to secure communications or provide your own certificates, as described in Chapter 8, ‘Managing Certificates for Server-Sensor Communication’.

Sensor installers are accessible from the CB Response **Sensor Group** pages. Between server releases, there are often installers for newer sensor releases available on the Carbon Black User Exchange.

Supported Operating Systems and Versions

CB Response supports sensors for Windows, Mac, and Linux environments.

For the specific operating system versions supported for this release, see the following page on the Carbon Black User Exchange:

https://community.carbonblack.com/t5/Documentation-Downloads/CB-Response-sensors-amp-CB-Protection-agents/ta-p/33041

Installing Sensors on Windows

This section describes how to install CB Response Windows sensors.

> **Note**
>
> To install sensors on Windows systems, you must belong to the Local Administrators (or higher) permissions group.

There are two ways to install Windows sensors:

- **Windows Standalone EXE** – Installs a sensor onto a single host. This option is useful for bringing a new host online in your network.
- **Windows MSI for GPO Installation** – Deploys sensors to multiple hosts over the network using Microsoft's Group Policy Objects (GPO). This option is also appropriate for deploying sensors remotely with third-party software deployment applications using standard **Msiexec** commands.

  For information about Windows MSI for GPO, see https://technet.microsoft.com/en-us/library/hh147307(v=ws.10).aspx.

To have access to the Download Sensor Installer controls for a sensor group, you must be **one** of the following:

- a user that is a member of a team that has either the Viewer or Analyst role for the sensor group
- for on-premise installations, a Global Administrator
- for cloud installations, an Administrator

To install sensors on Windows endpoints:

1. On the navigation bar, click **Sensors**.
2. In the **Groups** panel, select the sensor group you want the new sensor to join.

   The **Download Sensor Installer** drop-down list appears at the top of the Sensors page.
3. From the **Download Sensor Installer** list, select one of the following to download a ZIP file that contains the package installer:
   - To install a single sensor only, select either **Windows Standalone EXE - Latest Version** or **Windows Standalone EXE - <*specific version*>**.
   - To install one or more sensors, select **Windows MSI for GPO Installation - Latest Version** or **Windows MSI for GPO Installation - <*specific version*>**.

     where <*specific version*> is the sensor version specified in the group's upgrade policy for automatic updates. This is useful for bringing on a new host and installing the same sensor version installed group-wide. See "Upgrade Policy Settings" on page 112 for information about sensor group upgrade policies.
4. For a standalone installation do the following:
   a. In the operating system, copy the downloaded `<install package name>.zip` file to the Windows endpoint (XP SP3 or higher, either 32- or 64-bit).
   b. Extract the contents of the `<install package name>.zip` file to a temporary folder. Do not skip this step.
   c. Double-click to run the file `CarbonBlackClientSetup.exe`, and then follow the installation prompts.
5. For the Windows MSI for GPO installation, follow the instructions in the `GPO_README.txt` file, which is included in the downloaded `<install package name>.zip` file.

After the installation is complete, the Windows sensor is installed and running. The Sensors page shows the sensor as registered and checking into the CB Response server.

HTTP Proxy Support in Windows Sensors

For Windows sensors versions prior to 6.2.3, proxies are supported in the following way:

- In the `Sensorsettings.ini` file, set a configuration string `"Proxy=server:port"`. *Server* can be the host name or IP address of a proxy server, and *port* is an optional numerical port value.
- There is no server-side user interface to set this configuration option. The server cannot inform the sensor of this configuration via the `Protobuf` message.

For Windows sensor version 6.2.3 and above, proxies are supported in the following ways:

- Set the proxy configuration string in `Sensorsettings.ini` to the following values (case sensitive):
  - `"Proxy=@wpad"`: this instructs the sensor to use the WPAD protocol to autodetect the proxy settings. If autodetection fails, a direct connection is used. This is done for every request that the sensor needs to make.
  - `"Proxy=@pacurl:URL"`, where URL is the URL of a PAC file; for example, `https://server.example.com/example.pac`. This instructs the sensor to download the PAC file at this URL, and to use the proxy that it configures for the

current networking conditions and request. If the PAC file cannot be found, a direct connection is used. This is done for every request that the sensor makes.

- "`Proxy=server:port`": This setting was supported in previous sensor versions; however, the behavior is enhanced. If communications fail with a DNS resolution failure, or fails to connect to the server, a direct connection is attempted. Note that an HTTP failure (status code other than 200 OK) does not trigger a direct connection.

- There is no server-side user interface to set this configuration option. You must edit the `Sensorsettings.ini` file before you install the sensor.
- Alternatively, for already installed sensors, you can set the registry value:"`HKLM\SOFTWARE\CarbonBlack\config\Proxy`" to the desired string (type `REG_SZ`). For example, run the following command as an admin from PowerShell or a command prompt window to set the proxy configuration to use WPAD:
  "`reg add HKLM\SOFTWARE\CarbonBlack\config /v Proxy /t REG_SZ /d @wpad`"
  It is recommended to use either "`@wpad`"or "`@pacurl:URL`", depending on your configuration.

Uninstalling Windows Sensors

To uninstall Windows sensors, you can either use the CB Response console and follow the instructions in "Uninstalling Sensors via the Console" on page 83, or you can manually uninstall them by using the following procedure.

In Windows, you must be in the Local Administrators (or higher) permissions group to manually uninstall the sensor.

To manually uninstall Windows sensors:

- Either:
  - Launch the uninstall file in the `%windir%/CarbonBlack/` directory, or
  - Navigate to **Control Panel > Add/Remove Programs** and use the Windows application uninstall feature.

Installing Sensors on macOS Systems

This section describes how to install CB Response macOS sensors.

Note

To install a sensors on macOS systems, you must have access to an administrator account on that system.

To have access to the Download Sensor Installer controls for a sensor group, you must be one of the following:

- a user that is a member of a team that has either the Viewer or Analyst role for the sensor group
- for on-premise installations, a Global Administrator
- for cloud installations, an Administrator

To install sensors on macOS endpoints:

1. In the navigation bar of the CB Response console, click **Sensors** to display the Sensors page.
2. In the **Groups** panel, select the sensor group for installing the sensor package.
3. From the **Download Sensor Installer** list at the top of the Sensors page, select **OSX Standalone PKG**.

   The sensor package file is downloaded to your system.

4. In the operating system, do the following:
   a. Copy the `<install package name>.zip` sensor installation package to the macOS endpoint.
   b. Extract the `<install package name>.zip` file to a temporary folder. Do not skip this step.
   c. From the extracted `.zip` file, double-click to run the `.pkg` file, and then follow the installation prompts. You can also launch the `.pkg` file by using a silent installer. For example:

      ```
      installer -pkg <install package name>.pkg -target /
      ```

      This installs the macOS sensor using the configuration that is provided in the `sensorsettings.ini` file.

After the installation is complete, the macOS sensor is installed and running. The Sensors page shows the sensor as registered and checking into the CB Response server.

Upgrading Sensors on macOS

To upgrade sensors on macOS, see the instructions in “Upgrading Sensors” on page 83. You must have access to an administrator account on the macOS system to perform the upgrade.

Uninstalling Sensors on macOS

To uninstall macOS sensors, you can either use the CB Response console and follow the instructions in “Uninstalling Sensors via the Console” on page 83, or you can manually uninstall them by using the following procedure.

To manually uninstall a macOS (or OS X) sensor, you must have access to an administrator account or be assigned to an Analyst team with uninstall sensor privileges for the sensor group that the sensor is in.

To manually uninstall macOS sensors:

- On the macOS endpoint where the sensor is installed, run the following command:

  ```
  /Applications/CarbonBlack/sensoruninst.sh
  ```

After this process is complete, the endpoint stops reporting events and binaries to the CB Response server and all the caching information for logs is deleted.

Installing Sensors on Linux Systems

This section describes the steps to install the CB Response Linux sensor.

You must have the following requirements before installing the sensor:

- CB Response server version 5.0 or higher
- OpenSSL version 1.0.1 or higher

> **Note**
>
> To install sensors on Linux systems, you must be a root user or have “sudoer” permissions and run the installer with “sudo”.

To have access to the Download Sensor Installer controls for a sensor group, you must be one of the following:

- a user that is a member of a team that has either the Viewer or Analyst role for the sensor group
- for on-premise installations, a Global Administrator
- for cloud installations, an Administrator

To install sensors on Linux endpoints:

1. In the navigation bar of the CB Response console, click **Sensors**.
2. In the **Groups** panel, select the sensor group for which you want to install the sensor package.

3. From the **Download Sensor Installer** drop-down list, select **Linux Standalone RPM**.

   The sensor package file is downloaded to your system.

4. In the operating system, do the following:

   a. Copy the `<install package name>.tar.gz` sensor installation package to the Linux endpoint.

   b. Untar the `<install package name>.tar.gz` file to a temporary folder. Do not skip this step.

      For example, at a command prompt and from the directory where the file is installed, run this command:

      ```
      tar -zxvf <install file name>.tar.gz
      ```

   c. From the extracted `.tar.gz` file, run the `.sh` file and then follow the installation prompts.

      This installs the Linux sensor using the configuration provided in the `sensorsettings.ini` file.

After this process is complete, the Linux sensor is installed and running. The Sensors page shows the sensor as registered and checking into the CB Response server.

Upgrading Sensors on Linux

To upgrade sensors on Linux, follow the instructions in “Upgrading Sensors” on page 83.

> **Note**
>
> To upgrade sensors on Linux systems, you must be a root user or have “sudoer” permissions and run the installer with “sudo”.

Uninstalling Sensors on Linux

To uninstall Linux sensors, you can either use the CB Response console and follow the instructions in “Uninstalling Sensors via the Console” on page 83, or you can manually uninstall them by following the steps described here.

To manually uninstall a Linux sensor, you must be a root user or have “sudoer” permissions and run the installer with “sudo”.

To manually uninstall Linux sensors:

- On the Linux endpoint where the sensor is installed, run the following command:

  ```
  /opt/cbsensor/sensoruninstall.sh
  ```

When this process is complete, the endpoint stops reporting events and binaries to the CB Response server.

Upgrading Sensors

A new release of CB Response server can include a new sensor version. Check the server release notes to confirm if a new sensor version is available. Decide if you want to deploy the updated sensor immediately to existing sensor installations, or only install it where there has not been a sensor before.

Important

Carbon Black strongly recommends that you upgrade your sensors as soon as possible when a new version is available. However, you should not upgrade all sensors at once if you have a large number of sensors due to potential performance issues.

If you want to use automatic upgrades, consider gradually enabling automatic upgrades one sensor group at a time.

Each sensor group has an upgrade policy that determines how and when the sensors in the group are updated, and to what version. You set the upgrade policy for a sensor group in the **Create or Edit Group** panel of the Sensors page.

Upgrade policy options are as follows:

- Manually update sensors at the time of your choice using the **Download Sensor Installer** menu.
- Automatically upgrade sensors to the latest version.
- Update sensors to a specific version.

See “Upgrade Policy Settings” on page 112 for a description of the upgrade policy options for a sensor group.

For information on the latest sensors, visit the Carbon Black User eXchange.

To upgrade a sensor using the CB Response console, you must be **one** of the following:

- a user that has the Analyst role for the sensor group for the endpoint being acted upon
- for on-premise installations, a Global Administrator
- for cloud installations, an Administrator

Uninstalling Sensors via the Console

For all OS platforms, you can uninstall sensors using the CB Response console. After you uninstall sensors, they will stop reporting events and binaries from the endpoints on which they are installed to the CB Response server.

To uninstall a sensor using the CB Response console, you must be one of the following:

- a user that has the enhanced Analyst permission for uninstalling sensors and is a member of a team that has the Analyst role for the sensor group for the endpoint being acted upon
- for on-premise installations, a Global Administrator
- for cloud installations, an Administrator

To uninstall sensors using the console (all platforms):

1. In the navigation bar of the CB Response console, click **Sensors**.
2. In the **Groups** panel, select the sensor group with the sensor to uninstall.
3. In the sensors list, select the checkbox next to the sensor(s) to uninstall.
4. From the **Actions** drop-down list select **Uninstall**.
5. In the **Uninstall Sensors Confirmation** dialog box, click **Okay** to confirm the uninstall action.

   The sensor(s) are uninstalled.

Note

The sensor receives the uninstall request the next time it checks in with the server, which can be anytime between 30 seconds to several minutes, depending on the number of active sensors and the server load. Uninstalled sensors do not appear in sensor and host lists, unless the **Show Uninstalled Sensors** check box is selected.

Obtaining New Sensor Installation Packages

Periodically, Carbon Black releases new sensor versions either standalone or with a version of CB Response server. When you install or upgrade the server, you can choose to load the latest sensor installers, or to install or upgrade the server version only.

The sensor installers are downloaded to the sensor installation directory, either the default of `/usr/share/cb/coreservices/installers`, or a custom location specified by SensorInstallerDir|Osx|Linux in the `cb.conf` file. (See the *CB Response Server Configuration Guide* for details.)

Apart from a server installation or upgrade, you can download any new sensor installers manually from the Carbon Black yum repo (as described in the release announcement on the Carbon Black User eXchange).

After the installation packages are in the sensor installation directory, they can be made available in the following places in the CB Response console UI:

- The **Download Sensor Installer** drop-down list on the Sensors page when a group is selected.
- The sensor versions available for upgrades (either automatically or manually) according to the upgrade policy for a sensor group. For information about Upgrade Policy settings, see “Upgrade Policy Settings” on page 112.

The installer packages are made available through the UI in the following ways:

- At startup through coreservices
- By running this command:

  ```
  /usr/share/cb/cbcheck sensor-builds --update
  ```

Chapter 6

Managing Sensors

This chapter describes how sensors work, the information they provide, and how to search for and monitor sensors.

- See “Managing User Access with Teams” on page 53 for information about the roles and permissions that are required to view and modify sensors and their information.
- See Chapter 5, “Installing Sensors,” for information on installing, upgrading, and uninstalling sensors.
- See Chapter 7, “Sensor Groups,” for information on managing sensor groups.
- See Chapter 8, ‘Managing Certificates for Server-Sensor Communication’ for information about certificate options.
- See Chapter 9, “Troubleshooting Sensors,” for information on troubleshooting sensors.
- See Appendix A, “Sensor Parity,” for information on features supported on the sensor operating systems.

Sections

Topic	Page
Overview of Sensor Management	89
Monitoring Sensor Status and Activity	87
Monitoring Sensor and Server Information	91
Viewing Sensor Details	93

Overview of Sensor Management

Installed sensors gather event data on host computers (endpoints) and securely deliver the data to the CB Response server for storage and indexing. This enables your team to see and understand the history of an attack, even if the attacker deleted artifacts of its presence.

A sensor checks in with the CB Response server every five minutes to report the activity it detects. The server responds and notifies the sensor about how much data to send. To aid in detecting IOCs, the server compares the data it records from sensors with the latest data synchronized from threat intelligence feed partners you have enabled.

Each sensor belongs to a sensor group that defines the configuration and security characteristics for the sensor. For example, sensor groups define the upgrade policy and types of event information that sensors in the group collect. One sensor group can contain many sensors, but a single sensor can only belong to one sensor group. See Chapter 7, “Sensor Groups,” for more information.

To secure communication between sensors and the server, CB Response uses HTTPS and TLS. You can use the default server certificate or add your own server certificates and assign different certificates to different sensor groups. See Chapter 8, ‘Managing Certificates for Server-Sensor Communication’ for details.

Collected Data Types

Sensors collect information about the following data types:

- Currently running parent and child processes
- (OS X and Linux only) Fork and posix_exec processes
- Modules loaded by processes
- Processes blocked as the result of a CB Response hash ban
- Binaries
- File executions
- File modifications
- Network connections
- (Windows only) Registry modifications
- (Windows only) Cross-processes (an occurrence of a process that crosses the security boundary of another process)
- (Windows only) Enhanced Mitigation Experience Toolkit (EMET) events and configuration

Incident-Response Features

To help you manage sensors and work with the information they capture, CB Response provides incident-response features that provide the following capabilities:

- Directly respond to a threat detected on an endpoint through a command interface
- Isolate an endpoint with a suspicious process or threat
- Ban process hashes to prevent known malware from running in the future
- Set watchlists to monitor suspicious activity on endpoints

For information on these incident-response features, see Chapter 10, “Responding to Endpoint Incidents,” and Chapter 18, “Watchlists.”

Monitoring Sensor Status and Activity

The CB Response console provides multiple views into sensor activity on your endpoints.

- On the HUD page, the sensors panel gives a snapshot of sensor health, status, and activity.
- On the Sensors page, you can search for sensors and manage sensor groups.
- From anywhere in the console (Process Search or Watchlists pages, for example), you can click a hostname to get detailed information about a particular sensor.

The Sensors Page

The Sensors page in the CB Response console provides information about sensors and their host computers.

To access the Sensors page:

1. On the navigation bar, click **Sensors**.

   The Sensors page is organized as follows:

   - Computers (endpoints, or hosts) that have installed or uninstalled CB Response sensors appear in the right (Sensors) panel.
   - Sensor groups in which the sensors are included appear in the left (Groups) panel.

2. To change the list of displayed sensors, do any of the following:
   - To include uninstalled sensors in the list of sensors, check the **Show uninstalled sensors** box above the sensors table.
   - To view all sensors in a particular sensor group, click the name of the group in the **Groups** panel. The group name appears at the top of the **Sensors** panel.
   - To view all sensors, click **All Sensors** at the top of the **Groups** panel.
   - To search for one or more sensors, see “Searching for Sensors” on page 89.
3. When the list of sensors cannot fit on a single page, use controls at the bottom of the page to navigate multiple pages:

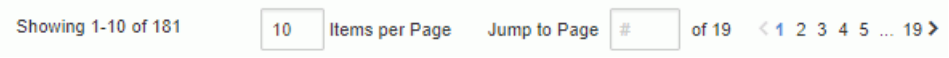

   - Enter the number of **Items per Page**.
   - Enter a number to **Jump to a Page**.
   - Click the forward and back arrows to navigate pages sequentially.
   - Click a number between the arrows to go to a specific page.
4. The following information appears for all sensors in the list:

Field	Description
Computer Name	The hostname corresponding to the computer on which the sensor is installed.
Domain Name	The registered DNS name for the IP address of the computer on which the sensor is installed.
IP Address	The IP address of the computer on which the sensor is installed.
Status	Describes the status of sensor connectivity as follows: • **Online** – Sensor communicated with the CB Response server within the previous expected check-in interval. • **Offline** – Sensor did not communicate with the CB Response server for more than a five-minute period after the expected check-in interval provided during the previous check-in. If known, offline status might include one of the following reasons: - **Offline (Suspended)** – Sensor detected that an OS-level suspend operation occurred before the sensor went offline. - **Offline (Restarting)** – Sensor detected that an OS-level restart operation occurred before the sensor went offline. - **Offline (Isolate configured)** – Sensor is offline and marked for isolation upon next check-in. - **Offline Uninstalled** – Appears when an uninstall was requested for an offline sensor. If a sensor is being uninstalled, one of the following status descriptions might appear: • **Uninstall uninstalled** – Requested uninstall operation has completed, and the sensor was successfully uninstalled. • **Uninstall pending uninstalled** – Uninstall operation was requested but has not yet completed.
Activity	The time that updated data is expected from the sensor; for example, “Was expected 2 seconds ago” or “Last seen about 3 months ago.”
OS Version	The operating system version of the computer on which the sensor is installed.
Server Certificate	The server certificate being used to secure communications with this sensor.
Node Id	In a clustered environment, the server ID to which a sensor sends data. For a standalone instance, the value is 0 (zero).

Field	Description
Sensor Version	Version of the currently installed CB Response sensor.

Searching for Sensors

On the Sensors page, you can search for sensors using either the **Search** box or a search based on filtered criteria.

To search for a sensor:

1. On the Sensors page, do one of the following:

   - In the **Search** box, enter characters in the name of the endpoints to find. Searching commences incrementally as you type and is case-insensitive.

     Search results include endpoints with sensors installed that have a name that matches the search string; if **Show uninstalled sensors** is selected, matching endpoints that have both installed and uninstalled sensors are included.
   - Click **Filter** and select any of the following criteria:

Filter Criteria	Description
Sensor Version	Installed version of a sensor.
Last Checkin Time	Timespan in which a sensor last checked into the CB Response server (last hour, last day, last week, and so on).
Node ID	In a clustered environment, the server ID to which a sensor sends data. For a standalone instance, the value is 0 (zero).
Feature Support	One of the following features a sensor reports as supporting: • **Live Response** – CBLR • **Isolation** – The sensor can be isolated. • **2nd Gen Modloads** – (macOS only) Binary modules the sensor reports as being loaded by a process.

     The list of sensors updates dynamically according to the filters selected. Search results include computers with installed, or with **Show uninstalled sensors** selected, uninstalled sensors that match the search criteria specified by one or more selected filters.
2. To clear all filters and search-box criteria and reset the Sensors page to an unfiltered list of sensors, click **Reset Filters**.

Exporting Sensor Data

From the Sensors page, you can download detailed sensor data to a CSV file.

To export sensor data from the Sensors page:

- From the **Export** drop-down list, click one of the following options:
  - **Export All** – Downloads data either for all installed sensors, or with **Show uninstalled sensors** selected, for both installed and uninstalled sensors on endpoints in your environment.
  - **Export Visible** – Downloads data only for sensors that are visible on the current page. For example, if there are 40 sensors in a list, and the list displays 20 items per page, the CSV file only contains data for the 20 sensors that are currently visible.

Sensor Actions

On the Sensors page, you can select sensors by selecting the check boxes next to the sensor names. Use the **Actions** drop-down list to perform the following actions on one or more selected sensors:

- **Sync** – Forces the sensor to send all the data that it has collected to the CB Response server immediately, ignoring any bandwidth throttles that might be configured.
- **Restart** – Restarts the sensor process.
- **Move to group** – Moves the sensor to another sensor group.
- **Uninstall** – Uninstalls the sensor from the host computer.
- **Isolate** – Isolates a computer from the rest of the network, leaving only the connections necessary for the CB Response server to access its sensor. The console UI provides the following cues for an isolated host:
  - On the Sensors page, the word **Isolated** appears in the **Status** column.
  - On the Sensor Details page, the message “This host has been isolated from the rest of the network” appears at the top; **Remove isolation** is the only option on the **Actions** list.

  For more information about the isolation feature, see “Isolating an Endpoint” on page 147.
- **Remove isolation** – From an isolated state, rejoins an endpoint to the network so that it can resume sending data to the CB Response server.

Monitoring Sensor and Server Information

The Server Dashboard provides an overview of the following sensor and server details:

- Sensor statistics
- Server communication status
- License information

This section describes how to view this information in the Server Dashboard, and descriptions of the details that display there.

> **Note**
>
> The Server Dashboard is available only to Global Administrators for on-premises installations and Administrators for cloud installations.

To view server and sensor information in the Server Dashboard:

1. On the navigation bar, click **Server Dashboard**.

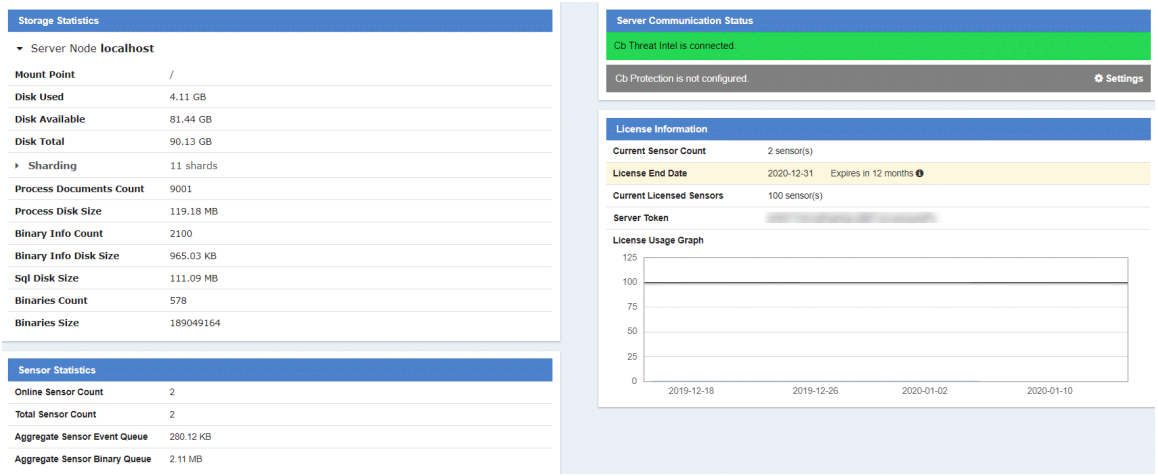

2. Review information in the **Storage Statistic** panel:

Field	Description
Mount Point	The mount point for the CB Response server data directory.
Disk Used	The amount of the disk space used for storage.
Disk Available	The amount of the disk space still available for storage.
Disk Total	The total amount of disk space.
Sharding	The number of shards on the disk. Expand to see associated ID, size, document count, and max document count.
Process Documents Count	The number of process documents uploaded to the database on the server. This is the same number as the total number of processes on the **Process Search page**.

Field	Description
Process Disk Size	The disk space used by the process documents.
Binary Info Count	The number of binaries that are seen by the sensor. This is the same number as the total number of binaries on the **Binary Search page**.
Binary Info Disk Size	The total number of bytes of binary information that is uploaded to the server.
Sql Disk Size	The psql database disk utilization.
Binaries Count	The number of binaries stored on the CB Response server.
Binaries Size	The total number of bytes of binaries stored on the CB Response server.

3. Review information in the **Sensor Statistics** panel:

Field	Description
Online Sensor Count	The number of sensors that are detected as being online by the CB Response server.
Total Sensor Count	The total number of sensors installed and registered for this server.
Aggregate Sensor Event Queue	The total size of queued events needing to be pushed to the CB Response server for all online sensors.
Aggregate Sensor Binary Queue	The total size of queued binaries needing to be pushed to the CB Response server for all online sensors.

4. Review information in the **Server Communication Status** panel:

Field	Description
CB Threat Intel is connected	Shows whether or not communication between the CB Response server and the CB Threat Intel has been established.
CB Protection is not configured	Shows whether or not communication between the CB Response server and a CB Protection server has been established.

5. Review the information in the **License Information** panel:

Field	Description
Current Sensor Count	Total number of unique online sensors in the last 24 hours (active).
License End Date	The date when the license terminates.
Current Licensed Sensors	The total number of sensors on the current license.

Field	Description
Server Token	The token for the CB Response server. This token is primarily used for CB Response support purposes.
License Usage Graph	A weekly depiction of how many sensors are active for the license.

Viewing Sensor Details

The Sensor Details page provides detailed information about each sensor.

To access the Sensor Details page:

- Do one of the following:
  - From a Process Search page, click the right arrow at the end of a row to open the Analysis Preview page.
  - From the Process Search page, HUD Sensors widget, or Sensors page, click the name of the endpoint.

Sensor Details Heading and Options

The heading in the **Sensor Details** panel displays the following information and options:

- An icon that represents the operating system of the host computer.
- The name of the host computer.
- Status of the sensor. This can be **Online**, **Offline**, **Uninstalled**, and whether the sensor is restarting, isolated, or syncing. If CB Response is not running or if there is a communication problem, the status is **Offline**.
- Click **Go Live** to open an interactive live session on the host computer to execute commands in real time to help isolate or eradicate a threat. For more information about **Go Live**, see “Using Live Response” on page 150.
- **Isolate / Remove Isolation**: Isolates a host computer from the rest of the network, leaving only the connections to access the CB Response server. If the sensor is isolated, this button displays **Remove isolation**. For more information about this feature, see “Isolating an Endpoint” on page 147.
- The **Actions** menu provides the following options:

- **Sync** – Forces the sensor to immediately send all the data that it has collected to the CB Response server, ignoring any configured bandwidth throttles.
- **Restart** – Restarts the sensor process.
- **Move to group** – Moves the sensor to another sensor group.
- **Uninstall** – Uninstalls the sensor from the host computer.

Summary

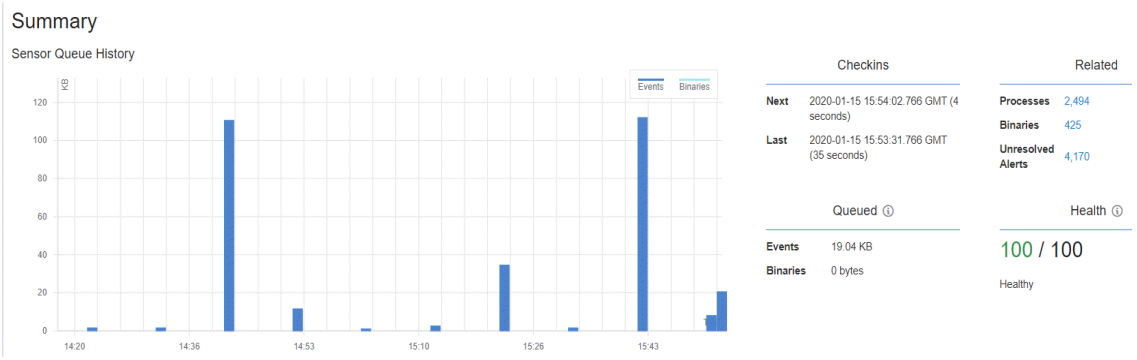

The **Summary** panel displays the following information about the sensor's health and activity:

Item	Description
Sensor Queue History	A graph that displays the recent history of activity in the sensor's queue. Click **View Details** to view this data in tabular form.
Checkins	Shows the **Next** and **Last** times (GMT) that the sensor checked in with the CB Response server.
Related	Processes, binaries, and alerts that are associated with the host.
Queued	Events and binaries that the sensor has not yet sent to the CB Response server. The maximum queue size can be managed in Sensor Group Settings.
Health	Sensor health score, and a health message from the sensor if there is one. See Appendix B, "Sensor Health Score Messages".

Vitals and Configuration

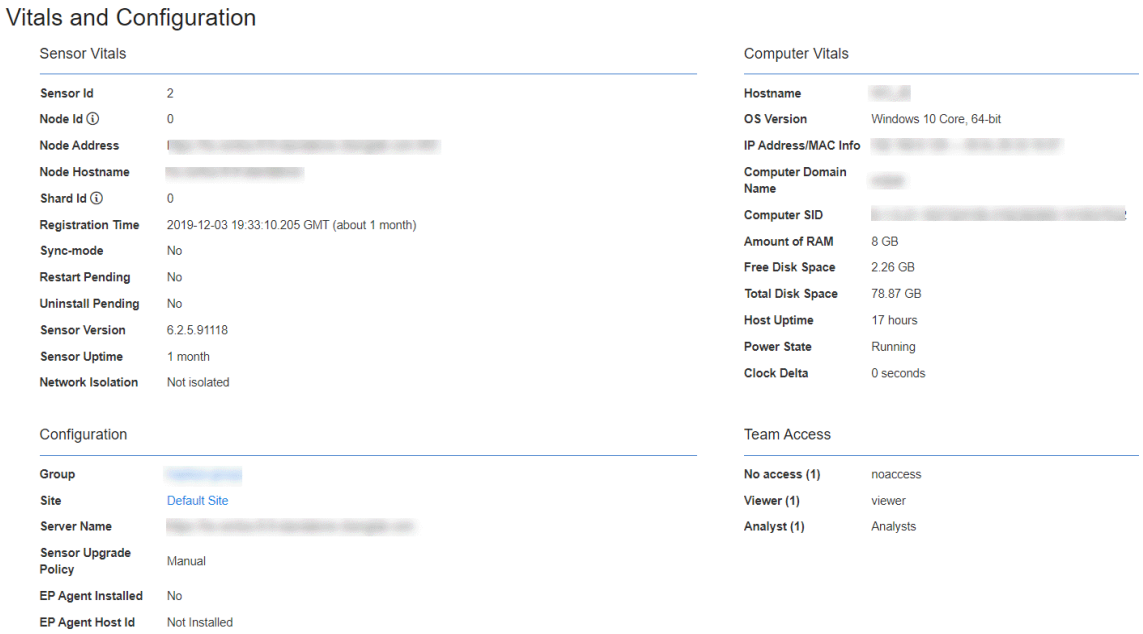

Sensor Vitals

The **Sensor Vitals** section of the **Vitals and Configuration** panel displays the following information:

Field	Description
Sensor Id	The internal CB Response sensor GUID of the host computer.
Node Id	The server ID to which the sensor sends data in a clustered environment.
Node Address	The address of the server to which the sensor sends data in a clustered environment.
Node Hostname	The host name of the server to which the sensor sends data in a clustered environment.
Shard Id	The shard ID to which the sensor submits event and binary metadata.
Registration Time	The date and time that the sensor registered (the start time of the sensor) with the CB Response server.
Sync-mode	Shows if the sensor is currently synchronizing with the server.
Restart Pending	Shows if the sensor host is restarting.
Uninstall Pending	Shows if the sensor is being uninstalled from the host.
Sensor Version	The current version of the sensor.

Field	Description
Sensor Uptime	The duration of time that the sensor has been actively running.
Network Isolation	Shows if the host is isolated from the rest of your network and the Internet. For information about isolating hosts, see "Isolating an Endpoint" on page 147.

Computer Vitals

The **Computer Vitals** section of the **Vitals and Configuration** panel shows the following details about the host:

Field	Description
Hostname	Hostname of the endpoint on which the sensor is installed.
OS Version	Version of the operating system on the host.
IP Address/MAC Info	IP and MAC address of the host.
Computer Domain Name	Name of the endpoint on which the sensor is installed. This can be the same name as the hostname.
Computer SID	Unique security identifier for the computer.
Amount of RAM	Amount of available RAM.
Free Disk Space	Amount of free disk space.
Total Disk Space	Amount of total disk space.
Host Uptime	Time that the host computer has been running since the last sensor boot. If the sensor has not rebooted since its installation, this value reflects the install time.
Power State	Indicates whether the host endpoint is running.
Clock Delta	The difference between the sensor clock and the server clock. If the delta is greater than 5 seconds, an alert is displayed.

Configuration

The **Configuration** section of the **Vitals and Configuration** panel shows the following sensor information:

Field	Description
Group	Sensor group to which this sensor belongs.
Site	Site to which the sensor group that contains this sensor belongs. For information about assigning sensor groups to sites, see "General Settings" on page 104.
Server Name	URL of the CB Response server.
Sensor Upgrade Policy	Upgrade policy for this sensor. See "Upgrade Policy Settings" on page 112 for details of sensor group upgrade policies.

Field	Description
EP Agent Installed	Indicates whether the CB Protection agent is installed.
EP Agent Host Id	If the CB Protection agent is installed, displays the CB Protection agent host Id.

Team Access

The **Team Access** section of the **Vitals and Configuration** panel shows the teams that have access to this sensor and the permissions that team members have. This information is defined in **Sensors > Edit Settings > Permissions**.

See "Permissions Settings" on page 109.

Recent Activity

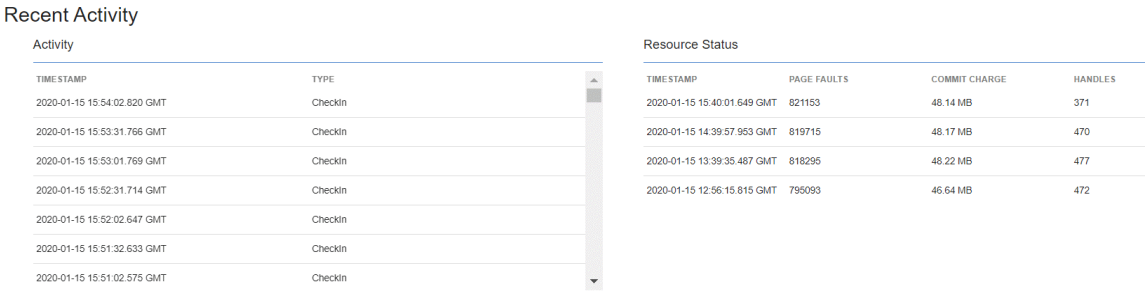

Activity

The **Activity** section of the **Recent Activity** panel shows the activity types that the sensor has been engaged in, and the date and time of each activity. This data is applicable for the duration that the sensor has been up and running.

Resource Status

The **Resource Status** section of the **Recent Activity** panel shows information for tracking internal performance metrics for sensors. If the performance of a sensor is degrading, the values in this table can help diagnose the cause. The panel shows the following details:

Field	Description
Timestamp	The date and time in one-hour intervals for which the sensor resource status is tracked.
Page Faults	The number of page faults that occurred on the date and time in the `Timestamp` field.
Commit Charge	The total memory used by all applications on the host endpoint, including memory that has been temporarily paged to disk at the date and time that is displayed in the `Timestamp` field.
Handles	The number of handles in use at the date and time displayed in the `Timestamp` field.

Diagnostics

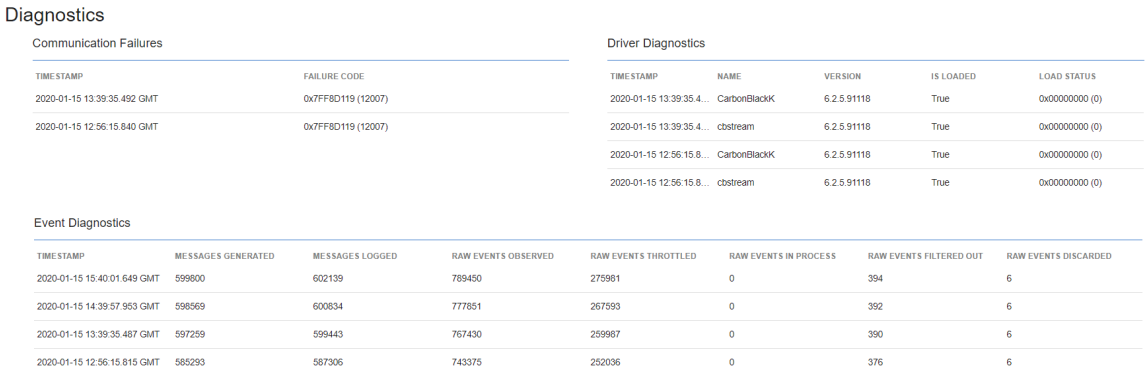
Diagnostics

Communication Failures

TIMESTAMP	FAILURE CODE
2020-01-15 13:39:35.492 GMT	0x7FF8D119 (12007)
2020-01-15 12:56:15.840 GMT	0x7FF8D119 (12007)

Driver Diagnostics

TIMESTAMP	NAME	VERSION	IS LOADED	LOAD STATUS
2020-01-15 13:39:35.4...	CarbonBlackK	6.2.5.91118	True	0x00000000 (0)
2020-01-15 13:39:35.4...	cbstream	6.2.5.91118	True	0x00000000 (0)
2020-01-15 12:56:15.8...	CarbonBlackK	6.2.5.91118	True	0x00000000 (0)
2020-01-15 12:56:15.8...	cbstream	6.2.5.91118	True	0x00000000 (0)

Event Diagnostics

TIMESTAMP	MESSAGES GENERATED	MESSAGES LOGGED	RAW EVENTS OBSERVED	RAW EVENTS THROTTLED	RAW EVENTS IN PROCESS	RAW EVENTS FILTERED OUT	RAW EVENTS DISCARDED
2020-01-15 15:40:01.649 GMT	599800	602139	789450	275981	0	394	6
2020-01-15 14:39:57.953 GMT	598569	600834	777851	267593	0	392	6
2020-01-15 13:39:35.487 GMT	597259	599443	767430	259987	0	390	6
2020-01-15 12:56:15.815 GMT	585293	587306	743375	252036	0	376	6

Communication Failures

The **Communication Failures** section of the **Diagnostics** panel shows the timestamp and failure code of communication failures between the sensor and the server.

You can locate the correct failure code and cross reference it with the information provided at https://curl.haxx.se/libcurl/c/libcurl-errors.html. For example, if you see error code 0x80c80013, locate “13” on this page.

Driver Diagnostics

The **Driver Diagnostics** section of the **Diagnostics** panel shows diagnostic information about the sensor driver.

CB Response OS X sensors have the following components:

- **CbSystemProxy** – A core kernel driver that improves interoperability with third-party products. When the OS X sensor is uninstalled, the next two kernel drivers are immediately removed and unloaded. The core kernel driver remains until the system reboots. Immediately unloading the core kernel driver can cause system instability if other products (typically security) are running in the system that integrate in the same way as CB Response.
- **CbOsxSensorProcmon** – A kernel driver to capture all other events.
- **CbOsxSensorNetmon** – A kernel driver to capture network events.
- **CbOsxSensorService** – A user-mode service to communicate with the CB Response server.

CB Response Windows sensors have the following components:

- **CoreDriver**
  - For Windows XP/2003/Vista/2008 (Vista server version), the driver binary name is `carbonblackk.sys`.
  - For Windows 7 and later, the binary name is `cbk7.sys`.
  - In all cases, the core driver is a mini-filter driver with the service name `carbonblackk`.
  - The core driver captures all events except for network connection events and passes all events, except tamper events, to the user-mode service.
  - The core driver attempts to directly send Tamper events to the CB Response server. If this fails, then the core driver attempts to send the Tamper events to the user-mode service.
- **Network Filter Driver**
  - For Windows XP/2003, the network filter driver is a Transport Driver Interface (TDI) filter driver with the binary name `cbtdiflt.sys` and service name `cbtdiflt`.
  - For Windows Vista and later, the network filter driver is a Windows Filter Platform (WFP) driver with the binary name `cbstream.sys` and service name `cbstream`.
  - The network filter driver is responsible for collecting network connection events and implementing the network isolation feature of the Windows sensor.
- **User-mode Service**
  - The sensor uses a user-mode service with the binary name `cb.exe` and service name `CarbonBlack`.
  - This service communicates with the core and network filter drivers to gather and process events from the kernel and send those to the server.

CB Response Linux sensors have these components:

- **Kernel Module –** This module does the following:
  - Uses a binary named `cbsensor.ko.<kernel version>` where the `<kernel version>` is a currently supported kernel.
  - Captures all system events and makes them available to the user mode daemon to process.
  - Exposes performance statistics in the `/proc/cb` directory.
- **User Mode Daemon** – This user-mode daemon uses a binary named `cbdaemon`. This service communicates with the kernel module to gather and process events to be sent to the server.

The sensor starts recording activity as soon as the core driver is loaded. It queues up the activity for the user mode service to receive as soon as it starts. This occurs early in the sensor boot process.

The network driver is loaded after the core driver, but it also starts recording as soon as it is loaded, and it also queues events for the user mode service.

These kernel driver components usually work in sync with each other, but it is possible for the sensor to be communicating with the server while one of the drivers is inoperable.

The **Driver Diagnostics** section of the **Diagnostics** panel shows the following information about the status of these drivers:

Field	Description
Timestamp	The date and time that the driver was loaded.
Name	The name of the driver.
Version	The version of the driver.
Is Loaded	Shows whether the driver is loaded (true or false).
Load Status	The load status of the driver.

Reducing the Impact of Netconn Data Collection (Windows)

On systems that have a large number of network connections (for example, DHCP/DNS servers, domain controllers, build servers, etc.), netconn data collection by the CB Response sensor can cause significant CPU utilization by the Carbon Black service. If this is an issue but you want to continue collecting netconn data, Windows sensors beginning with v6.1.4 let you disable the DNS name resolution in data collection for network connections, thereby reducing the amount of netconn traffic on these systems. This is done by configuring the following Windows registry key:

```
[HKEY_LOCAL_MACHINE\SOFTWARE\CarbonBlack\config]
"DisableNetConnNameResolution"=dword:00000001
```

Event Diagnostics

The **Event Diagnostics** panel shows the date and timestamp (in GMT) of sensor events, together with the number of each of the following event elements:

- Messages Generated
- Messages Logged
- Raw Events Observed
- Raw Events Throttled
- Raw Events in Process
- Raw Events Filtered Out
- Raw Events Discarded

Status History

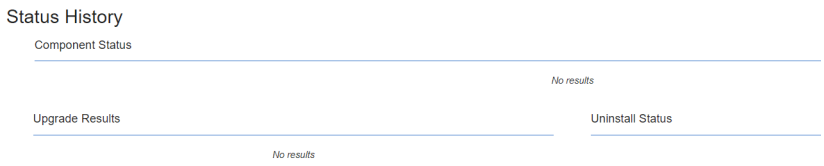

Component Status

The **Component Status** section in the **Status History** panel shows information about threads and the sensor component.

Upgrade Status

The **Upgrade Status** section of the **Status History** panel shows details about the sensor upgrade process.

Uninstall Status

The **Uninstall Status** section of the **Status History** panel shows data if the sensor has been recently uninstalled.

Chapter 7

Sensor Groups

This chapter describes creating, moving, editing, and deleting sensor groups.

Sections

Topic	Page
Overview of Sensor Groups	103
Create or Edit a Sensor Group	103
Moving Sensors to Another Group	112
Deleting Sensor Groups	113

Overview of Sensor Groups

CB Response sensors are lightweight data gatherers installed on network endpoints (such as laptops, desktops, and servers). They gather event data on the endpoints and securely deliver it to the CB Response server for storage and indexing. Each sensor is associated with a sensor group that defines its configuration and security characteristics. One sensor group can contain many sensors, but a single sensor can only belong to one sensor group.

Sensor groups can be based on your security and organizational requirements. For example, you might base sensor groups on functional groups (such as marketing, customer service, or IT) or location.

If you move sensors from one sensor group to another, the sensors will receive security settings from the new group the next time they check back into the server. In most cases, you do not have to re-install the sensors when you move them.

For more information:

- See Chapter 5, “Installing Sensors” for information on installing sensors.
- See Chapter 6, “Managing Sensors” for information on managing sensors.
- See Chapter 8, ‘Managing Certificates for Server-Sensor Communication’ for information about certificate options.
- See Appendix A, “Sensor Parity,” for information on supported operating systems. This appendix indicates whether or not a supported configuration is available on a sensor and configurable on a sensor group.

Create or Edit a Sensor Group

You can create sensor groups before or after you install sensors.

To create or edit a sensor group:

1. On the navigation bar, click **Sensors**.
2. In the **Groups** panel of the Sensors page, do one of the following:
   - To create a new group, click **NEW** at the top of the **Groups** panel.

     The **Create Group** panel appears to the right of the **Groups** panel.
   - To edit an existing group, do one of the following:

     -Select a group and at the top of the Sensors page click **Edit**.

     -Next to a group name, click the gear icon(⚙).

     The **Edit Group** panel appears.

   In the **Create Group** or **Edit Group** panel, sensor group settings are organized in sections that you can expand or collapse.

> **Note**
>
> To quickly open or close all sections of sensor group settings at one time, click **Display All Sections** or **Collapse All Sections** at the top right of the **Create Group** or **Edit Group** panel.

3. Complete the settings in the following categories:
   - **General** – See “General Settings” on page 104.
   - **Sharing** – See “Sharing Settings” on page 105.
   - **Advanced** – See “Advanced Settings” on page 107.
   - **Permissions** – See “Permissions Settings” on page 109. Note that if you do not specify otherwise, all teams are set up with the **No Access** role for the new sensor group.
   - **Event Collection** – See “Event Collection Settings” on page 109.
   - **Isolation Exclusions** – See “Isolation Exclusions” on page 149.
   - **Exclusions** – See “Exclusion Settings (OS X/macOS only)” on page 110.
   - **Upgrade Policy** – See “Upgrade Policy Settings” on page 112.

   Tip

   While viewing settings for one sensor group, you can switch to display settings for a different group by clicking the gear icon next to the other group. The **Edit Group** panel refreshes to show the settings for the newly selected group.

4. When you finish configuring the sensor group settings, do one of the following:
   - Click **Create Group** to create a new group.
   - Click **Save Group** to save edits to an existing group.

     Sensor group changes take effect the next time the sensors report to the CB Response server.

   Note

   If any errors are introduced in the CB Response server URL (in the **General** section of the Edit Group page), you will lose communication with deployed sensors.

General Settings

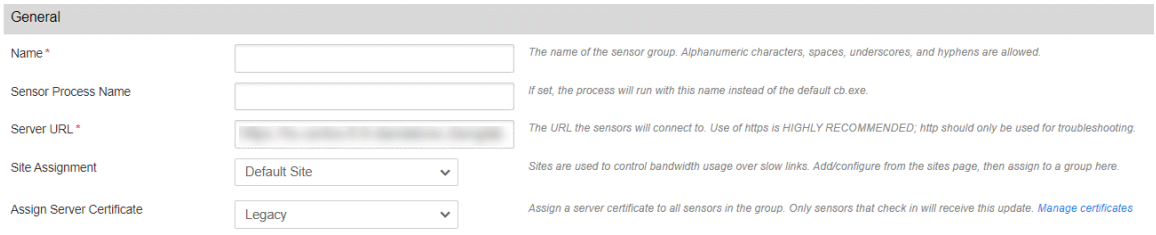

The **General** section includes the following settings:

Setting	Description
Name	The name of the sensor group; alphanumeric characters only.
Sensor Process Name	(Optional, Windows-only) An alternate name for the sensor group process. The default name of the process is `cb.exe`. For example, you can change the default name if Operations Security (OPSEC) policies require sensors to run with a non-standard or obfuscated executable name. If you change the name of the sensor process, the process will run with this name instead of the default `cb.exe`. This will not change the Windows service display name, but it will change the name of the actual executable that is run.
Server URL	The URL that the sensor group uses to communicate with the CB Response server. This URL is the same one that is used to log into the CB Response server, prefixed with *sensors*. Use HTTPS and specify the secure port in the URL.
Site Assignment	Select a site to assign to this sensor group. You can use site definitions to define throttle settings to manage bandwidth for groups of computers. These settings are applied per site, not per sensor group. If bandwidth is an issue for this group of sensors, create or configure a site with the appropriate bandwidth settings in ***Username* > Settings > Sites**, and then assign the site to this sensor group by selecting the site in this field. **Note:** Modifying bandwidth settings for sites requires Global Administrator status for on-premise installations and Administrator status for the cloud. Additional information about site throttling is available in the *Carbon Black Enterprise Response - Operating Environment Requirements* on the User eXchange.
Assign Server Certificate	Assign a server certificate to all sensors in the group. Only sensors that check in receive this update. This field includes a *Manage certificates* link that goes to the **Server Certificates** tab of the Settings page. See Chapter 8, 'Managing Certificates for Server-Sensor Communication'.

Sharing Settings

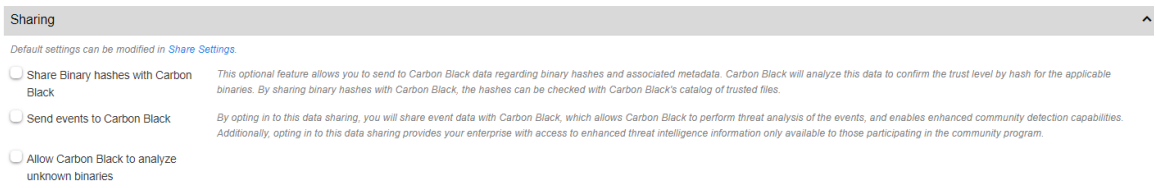

The **Sharing** section includes the following settings:

Setting	Description
Share Binary hashes with Carbon Black	Select this option to be notified of any binary that is flagged by Carbon Black Collective Defense Cloud. For more information about this choice, do the following from the **Sharing** section of the **Create** or **Edit Group** page: **1.** Click **Share Settings** to open the Sharing page for the CB Response server. **2.** On the Sharing page, scroll to **Endpoint Activity Sharing**. **3.** In the **Carbon Black** column next to **Binary Hashes & Metadata**, click the current setting (**Enabled**, **Disabled**, or **Partial**) for a description.
Send events to Carbon Black	Select this option to: • Allow advanced analysis of your aggregated process execution events by the Carbon Black Threat Research Team. • Give your enterprise access to enhanced CB Threat Intel information that is only available to those who participate in the community program. The Carbon Black Threat Research Team receives process events (as shown on the Process Analysis page) for more detailed analysis of the behavior of a process as it executes at the customer site. For more information on this page, see “Process Search and Analysis” on page 172.
Allow Carbon Black to analyze unknown binaries	Select this option to get advanced analysis of binary content from Carbon Black's Threat Research Team. By sharing binaries with Carbon Black, our researchers will perform advanced static analysis of your binaries to alert you to suspicious activity. Select this option to detect new variants of known malware. For more information about this choice, do the following: **1.** Click **Share Settings**. **2.** Scroll to **Endpoint Activity Sharing**. **3.** Click the current setting (**Enabled**, **Disabled**, or **Partial**) for a full description.

Default settings for the sensor group sharing settings are defined on the global Sharing page, which you can access in the following ways:

- Click the **Share Settings** link in the sensor group **Sharing** section.
- Click ***Username*** **> Sharing Settings**.

For more information, see “Data Sharing Settings” on page 248.

Advanced Settings

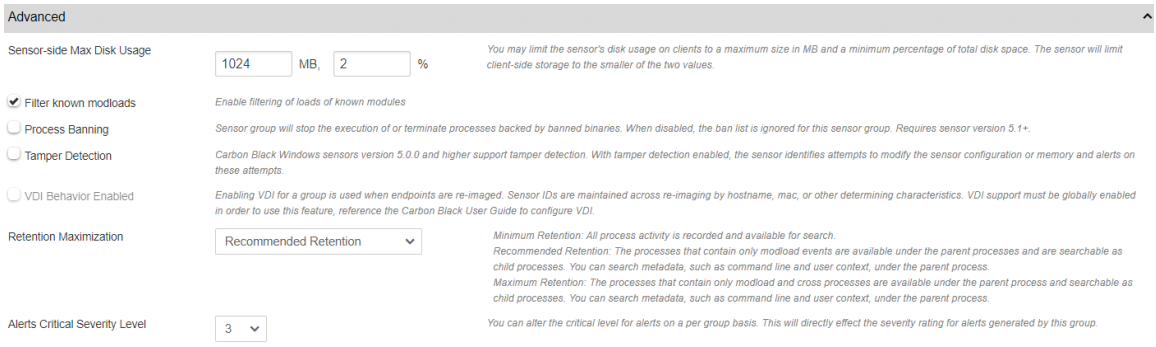

The **Advanced** section includes the following settings:

Setting	Description
Sensor-side Max Disk Usage	Contains two options to limit sensor disk consumption on clients either by raw available space (in megabytes) or percentage of the total space available. The sensors will limit the amount of space they use on clients based on the smaller of these two values: • In the **MB** field, enter the maximum available space on the client (between 2 and 10240 megabytes). • In the **%** field, enter the maximum percentage (between 2% and 25%) of total disk space on the client.
Filter known modloads (Windows and OS X only)	When selected, CB Response will not report the module load events of known good Windows and OS X modules that reside on the operating system. This helps reduce the number of known good events that are reported to the server.
Process Banning	When selected, this setting enables process hash bans in this group. By default, this setting is disabled and process hash bans prevent banned processes from running. For more information, see “Banning Process Hashes” on page 162.
Tamper Detection (Windows only)	When selected, the sensor identifies when attempts are made to modify the sensor's binaries, disk artifacts, or configuration. To change this setting you must be one of the following: a Global Administrator (on premises), an Administrator (cloud), or a user who is an Analyst for this sensor group and who also has enhanced permission for isolating sensors.
VDI Behavior Enabled	When selected, this setting enables Virtual Desktop Infrastructure (VDI) for sensors on virtual machines. Use VDI when endpoints that are virtual machines are re-imaged. Sensor IDs are maintained across re-imaging by hostname, MAC, or other determining characteristics. **Note**: VDI support must be globally enabled to use this feature. For more information, see “Server VDI Support” in the *CB Response Integration Guide*.

Setting	Description
Retention Maximization	These settings change how sensor process data that contains only modload processes or only modload and cross processes is recorded on the server. Minimum Retention makes this data more easily searchable but leaves a bigger footprint and can lead to a reduction in data retention time. Recommended and Maximum Retention consolidate data under parent processes, reducing the data footprint and helping increase the retention time. Data consolidated in this way is still searchable, as child processes. • **Minimum Retention** – All process activity is recorded and available for search. • **Recommended Retention** – The processes that contain only modload events are available under the parent processes and are searchable as child processes. You can search metadata, such as command line and user context, under the parent process. • **Maximum Retention** – The processes that contain only modload and cross processes are available under the parent processes and are searchable as child processes. You can search metadata, such as command line and user context, under the parent process. **Note:** Recommended and Maximum Retention can result in false positives in the results of cmdline searches. See "Retention Maximization and cmdline Searches" on page 231. **Note:** This setting was called **Data Suppression Level** in pre-6.5 versions of CB Response.
Alerts Critical Severity Level	Select a value from the menu to alter the critical level for alerts on a per-sensor-group basis. This directly effects the severity rating for alerts generated by this sensor group. On the Triage Alerts page, the severity score of an alert (located in the **Severity** column of the results table) is determined by three components: • Feed rating • Threat intelligence report score • Sensor criticality. For example, server sensors can have a higher criticality than engineering workstations. If two sensor groups have different alert criticalities, and they receive alerts from the same feed and for the same report, the sensor group that has the higher alert criticality will have a higher severity score on the Triage Alerts page, and servers in that group will appear at the top of the queue. For more information about alerts, see Chapter 19, "Console and Email Alerts". For more information about threat intelligence feed scores, see "Threat Intelligence Feeds" on page 244.

Permissions Settings

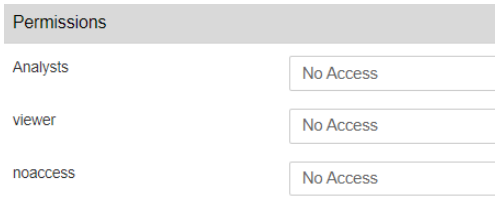

In the **Permissions** section, you can define user team permissions for sensors groups.

Available permission levels are as follows:

- **Analyst** – Users can configure the sensor host and group details.
- **Viewer** – Users can view the data collected from hosts in this sensor group. Users cannot make any configuration changes to this group or hosts that belong to it.
- **No Access** – When users in a team try to access or view details on a host in this sensor group, the system generates the following HTTP 405 response: "The method you are using to access the file is not allowed."

For information about user teams and access levels, see "Managing User Access with Teams" on page 53.

Event Collection Settings

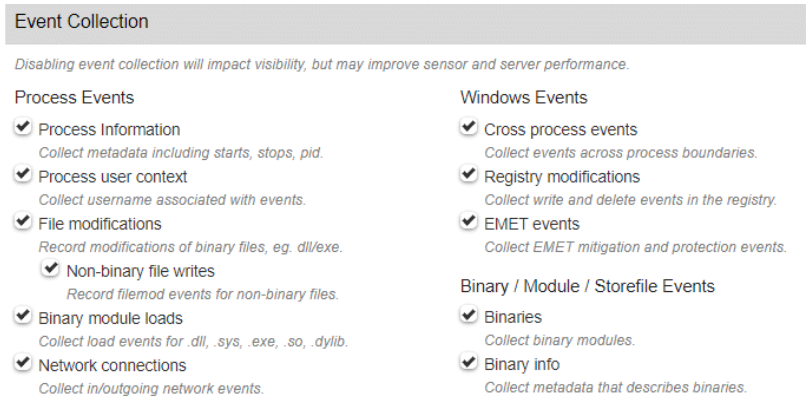

In the **Event Collection** section, you can define which types of events to record for the sensors in this group by selecting/deselecting the event types listed. Disabling event collection impacts visibility, but can improve sensor and server performance.

Most of the Event Collection options are self-explanatory, except for the following:

- **Process user context** – Enables the sensor to record the user name that is associated with each running process. This associates endpoint activity with the operating system user account.
- **Cross process events** – Enables the sensor to record instances when a process crosses the security boundary of another process. Although some of these events are benign, others might indicate an attempt to change the behavior of the target process by a malicious process.

Certain limitations exist on the cross process events that are reported by the sensor:

- Parent processes that create cross process events to their children are not reported.
- Cross process events that are part of the normal OS behaviors are ignored. For example, no cross process events are recorded for the Windows process `csrss.exe`.
- Cross process events are not reported for OS X and Linux sensors.
- Cross process, open process, and open thread events are not supported on Windows XP and Windows 2003.

Exclusion Settings (OS X/macOS only)

Through an addition to the `cb.conf` file, an **Exclusions** section can be added to the **Create Group** or **Edit Group** panel on the Sensors page. This **Exclusions** section lets you define paths on OS X/macOS systems and customize event collection at those paths to improve performance or eliminate unnecessary data. For example, you can specify that actions coming from one group of paths do not collect network connections or non-binary file writes. You can create another exclusion for a different set of paths that collects everything except cross-process events.

To add Exclusion settings to the sensor group panel on the Sensors page:

1. On the CB Response server, open `/etc/cb/cb.conf` for editing.
2. Add the following setting and value to the `cb.conf` file; consider including a comment to remind you of the purpose of the setting (and its current limitation to macOS):

   ```
   EventExclusionsEnabled=True
   ```
3. Save the `cb.conf` file.
4. You must stop and restart the server or cluster to make the new setting effective:
   - For a standalone server:

     ```
     sudo service cb-enterprise restart
     ```
   - For clusters:

     ```
     sudo cbcluster stop
     (...wait for all the nodes to shut down, and then...)
     sudo cbcluster start
     ```

Creating Exclusions

You can specify exclusions when you create a sensor group, or add them to an existing sensor group. The following procedure assumes that the group already exists.

To create an OS X/macOS event collection exclusion for a sensor group:

1. On the navigation bar, click **Sensors**.
2. In the **Groups** panel of the Sensors page, click the gear icon(⚙) next to the sensor group for which to create exclusions.
3. Click the **Exclusions** bar and click the **Add Exclusion** button.

   The Exclusion configuration fields are exposed.

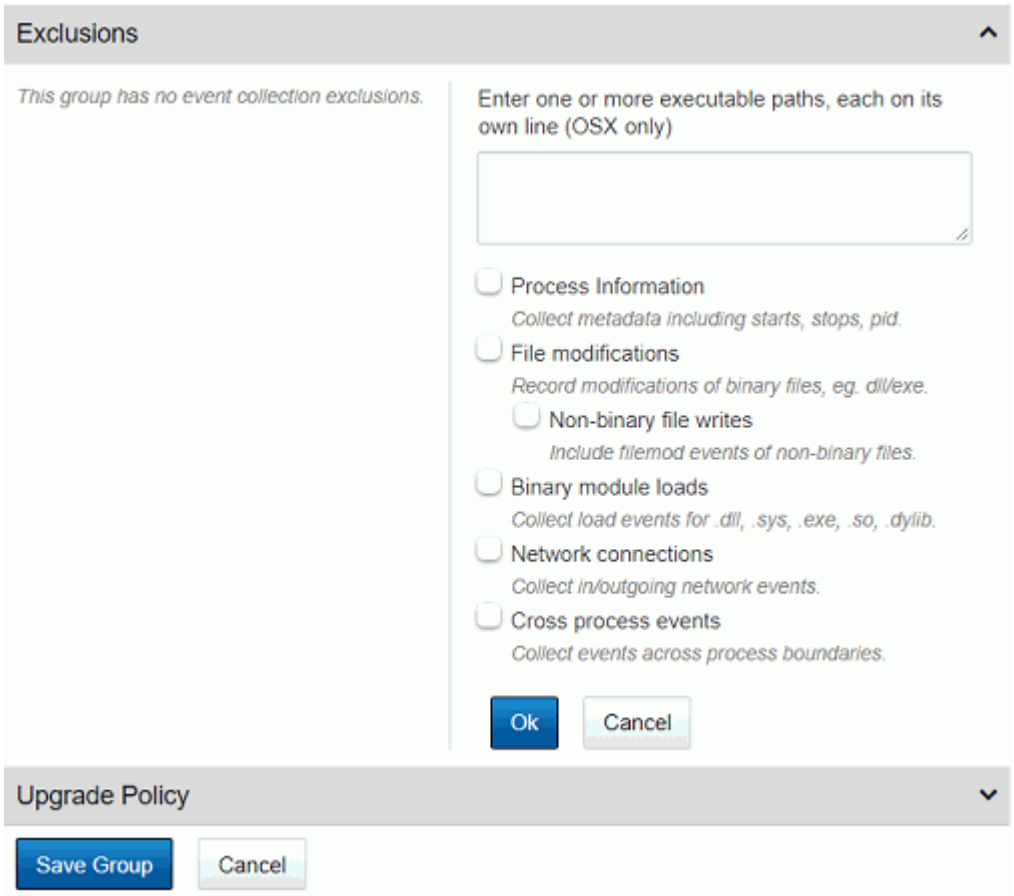

4. Enter the path(s) to affect with this exclusion in the textbox in the upper right corner of the panel. Put each path on a new line. You must use complete paths without wildcards.
5. Check the box next to each type of information to *not* collect for the specified paths. Click **Ok**.

   The exclusions are saved and displayed in the panel. You can edit or delete any exclusion.

6. When you have finished creating exclusions, click the **Save Group** button.

Upgrade Policy Settings

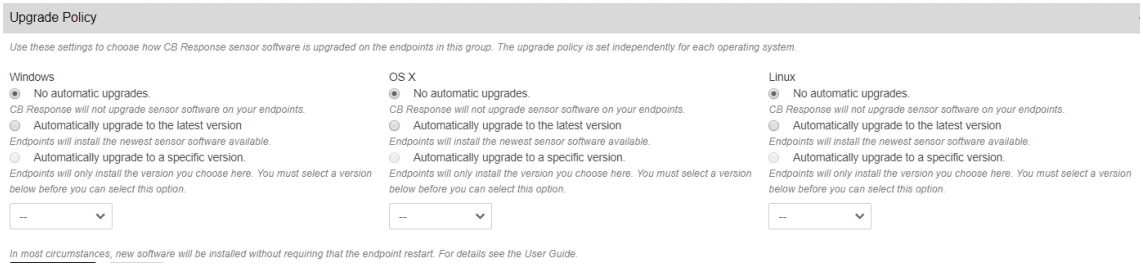

The **Upgrade Policy** section lets you set a policy to upgrade installed sensors in the sensor group.

Upgrade policy options are as follows:

- **No automatic updates** – Manually decide when to upgrade sensors.
- **Automatically install the latest version** – Automatically upgrades the sensors to the latest version.
- **Automatically install a specific version** – Installs a specific version for all sensors in a group. This maintains all sensors at the selected version. Select a version number using the drop-down list. Selecting the upgrade policy of a specific version is useful when sensor versions must be tested.

Moving Sensors to Another Group

After you create sensor groups, you can add sensors to them. By default, sensors are installed into the Default Group. On the Sensors page, you can select the group that contains the sensors to add, and then move those sensors from their original group to the new group.

To move sensors to a new sensor group:

1. On the navigation bar, click **Sensors**.
2. In the **Groups** panel, click the sensor group that contains the sensors to move.
3. In the **Sensors** panel, select the sensors to move.
4. Click **Actions > Move to group**.
5. From the drop-down list, select the sensor group to which to move the selected sensors and click **Okay**.

   The selected sensors are removed from the former sensor group list and appear in the new sensor group list.

Note

If you have set up custom server certificates and strict certificate validation, and you have assigned different certificates to different sensor groups, moving a sensor to another group could affect connectivity. See Chapter 8, ‘Managing Certificates for Server-Sensor Communication’.

Deleting Sensor Groups

You can delete sensor groups on the Sensors page. When you delete a sensor group, the teams for which you defined permissions no longer have access to sensors that belong to the sensor group.

To delete sensor groups:

1. On the navigation bar, click **Sensors**.
2. In the **Groups** panel, click the sensor group to delete.
3. Click **Delete Group** at the top of the Sensors page.

   A confirmation message appears indicating that any sensors remaining in this sensor group will be moved to the **Default Group**.
4. Click **OK**.

Chapter 8

Managing Certificates for Server-Sensor Communication

This chapter describes how CB Response uses HTTPS and TLS to secure communication and two-way authorization between endpoints and the server. It also details certificate management features, including the ability to add your own server certificates, assign different certificates to different sensor groups, and opt for stricter certificate validation.

Sections

Topic	Page
TLS Server Certificate Management Overview	115
Server-Sensor Certificate Requirements	116
How CB Response Supports Multiple Certificates	117
Managing Certificates on the Server	120
Viewing Certificate Information in the Console	120
Substituting a Legacy Certificate during Server Installation	121
Adding Certificates through the Console	122
Choosing a Validation Option	122
Changing the Expiration Notification Period	123
Deleting Certificates	124
Upgrades from Previous Server Releases	124
Assigning Certificates to Sensor Groups	125
Sensor Support for Certificate Management	126

TLS Server Certificate Management Overview

CB Response uses the HTTPS and TLS (formerly SSL) protocols to secure communication and two-way authorization between endpoints and the server so that the endpoint communicates only with the Response server that it trusts, and the server only communicates with trusted endpoints.

Prior to server version 6.4.0, CB Response established the trust between endpoints and the server by using “certificate pinning,” which is an out-of-band, reliable and secure trust mechanism. The server built the endpoint installer packages, and those came pre-initialized with the server identity (the public portion of server’s TLS certificate). The CB Response server acted as its own root certificate authority (CA), which allowed it to issue client-side certificates that the endpoints could use. This feature is still available and is the default option for securing server to sensor communications.

If you are satisfied with the security that is provided by the certificate generated by your CB Response Server and do not have any special compliance requirements, you may continue to use the standard certificate and validation method, which relies on certificate pinning only. Past and current sensors will continue to support this method.

Beginning with CB Response Server 6.4.0, you can instead choose to provide certificates signed by your organization. In addition, you can use different server certificates to authenticate the connections between the CB Response Server and different sensor groups, thereby reducing the exposure to a compromised server certificate. You can also add stricter validation methods to certificate pinning so that if a server certificate used by a sensor has expired or fails to meet other operating-system-specific criteria, server-sensor communication is disabled.

See “Sensor Support for Certificate Management” on page 126 for information about the sensor versions that support certificate management on each operating system.

In a cluster environment, master and minion servers use the same certificates. If you add your own certificates to the master, they are automatically propagated to the minions within a few seconds (unless there are connection issues). No server restart is required. The required format for user-provided certificates allows them to be seamlessly used in a clustered environment.

In addition, CB Response provides new certificate visibility features that can be useful for user-provided and CB Response “legacy” certificates.

Notes

- Currently, you can use certificates signed by *your own* certificate authority but use of a certificate that requires validation by a *third-party CA* is not supported.
- The certificate management features described here apply only to server-sensor communications. They are not used for managing other CB Response interactions, such as the connection between the console user interface and the server.

Certificate Management Feature Summary

- **Add and delete certificates** – You can add new certificates and delete certificates from your server.
- **View certificate inventory** – A table lists all server certificates that are available on the current server, how many sensors are using each one, and additional certificate information.
- **Choose validation method** – You can use standard certificate “pinning” validation, which only requires that sensors have a certificate matching the server, or you can add stricter validation methods. A certificate that uses standard validation continues to allow sensor and server to communicate even after it expires. but strict validation disables communication after expiration.
- **Be notified of expiring certificates** – When a certificate is close to its expiration date, an alert banner can be displayed at the top of each console page. You can set the number of days in advance you want to be warned, or turn off warnings. Deleting the expired certificate eliminates the notification.
- **Assign and change certificates by sensor group or apply one to all sensor groups** – If you have more than one certificate available, you can choose the certificate that is assigned to secure server communications for each sensor group. You can also apply one certificate to all sensor groups. This can be done for both the initial certificate assignment and to assign a new certificates — for example if a certificate is ready to expire.
- **View the certificate for a sensor** – The Sensors page shows the server certificate that was used for the last successful check-in for each sensor.
- **Control access to certificate features** – Because of their security implications, certificate management features require Global Administrator privileges on the server.

Server-Sensor Certificate Requirements

Whether added during server installation or later through the console, server certificates that are used for sensor communications must meet the following requirements:

- The files you provide must be valid certificate and key files (that is, they must be recognized as a certificate/key pair by the OpenSSL library).
- Certificate files must be in unencrypted ASCII PEM format – this includes both the certificate file and the key file.
- The certificate must have valid dates when uploaded – that is, its "not valid before" date should be in the past and its "not valid after" date should be in the future.
- Certificates must have two distinct SAN DNS entries to address the CB Response cluster scenario where sensors must resolve master and minion virtual addresses to different IP addresses or FQDNs. This is required for every server cert, even in standalone configurations, so that certificates remain valid if a standalone instance is upgraded to a cluster. The second SAN field is a single virtual address used for all minions, but it is mapped to a different IP address or FQDN hostname as needed by the sensor itself.
- SAN DNS entries must meet the standards for hostname formatting. Allowed characters include the hyphen and alphanumeric characters (a to z and 0 to 9). Invalid SAN DNS entries can silently fail.

- The CN field is not used for validation of new certificates because it has been deprecated. Sensors perform their own local resolution of virtual names to real Server addresses, so no additional DNS entries are required.
- No duplicate SAN entries are allowed in any active certificates – if a duplicate entry is found, the upload will not be allowed.

The following example shows how to set up the SAN portion of the certificate if you wanted to upload two certificates. The first SAN.DNS entry is used for the master and the second is used for the minions.

Certificate A

```
CN=<something>
SAN.DNS.1=virtual-a.master
SAN.DNS.2=virtual-a.minion
```

Certificate B

```
CN=<something>
SAN.DNS.1=virtual-b.master
SAN.DNS.2=virtual-b.minion
```

How CB Response Supports Multiple Certificates

The CB Response Server uses virtual server names to allow multiple active routes to the server using the same real address and port. Each virtual name and route uses a different certificate, as depicted in the schematics below. This implementation is done via runtime server blocks in NGINX configuration files.

Virtual server names are parsed from a SAN.DNS entry in the certificate so that each certificate can be validated from the sensor's perspective.

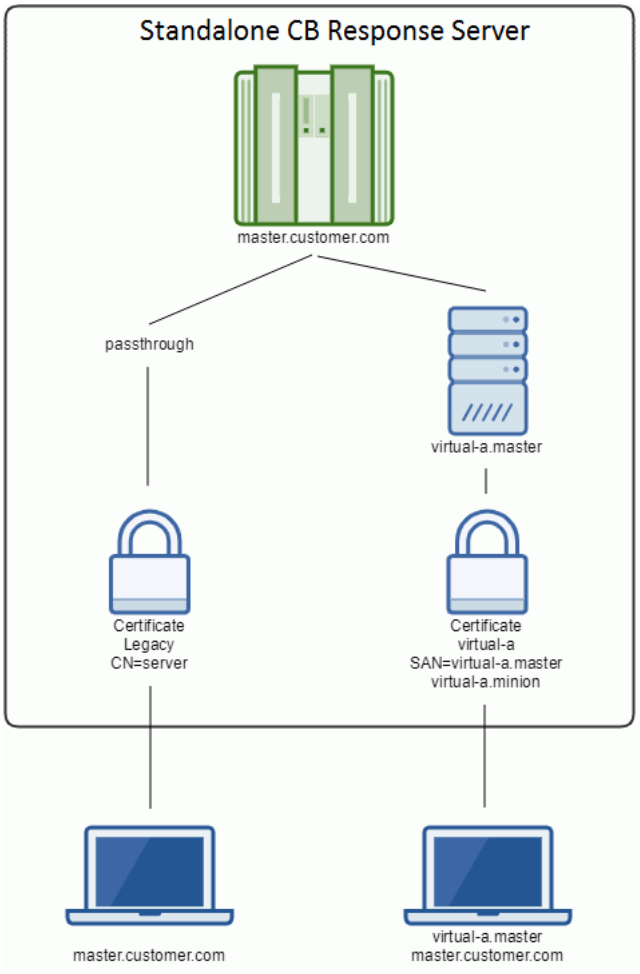

Sensors use the Server Name Indication (SNI) extension to the TLS protocol handshake to access a specific route and certificate. The Legacy certificate remains available without any SNI indications so that older sensor versions are able to use it to access the server.

Sensors confirm the resolution of virtual names to addresses internally; for example, resolution from "virtual-a.master" in the preceding certificate example to the actual "master.customer.com" server address. That means that virtual server addresses do not need to be added to external DNS servers.

On an upgrade to CB Response 6.4.0, the previously used self-signed certificate is named "Legacy" and would be served via the default server access route (using no virtual servers), which is supported for all sensor versions and configurations. New server installations also include a "Legacy" certificate.

Note

If you are using a Reverse Proxy, you must manually configure your Reverse Proxy to match the SNI configuration in your server environment. Contact Carbon Black Support for additional details.

Using Multiple Active Certificates in a Cluster

The following schematic describes some details of the clustered server set up with TLS certificate management. Although the legacy route is not depicted in this schematic for readability, it does exist and works in a cluster just as it does in a standalone scenario.

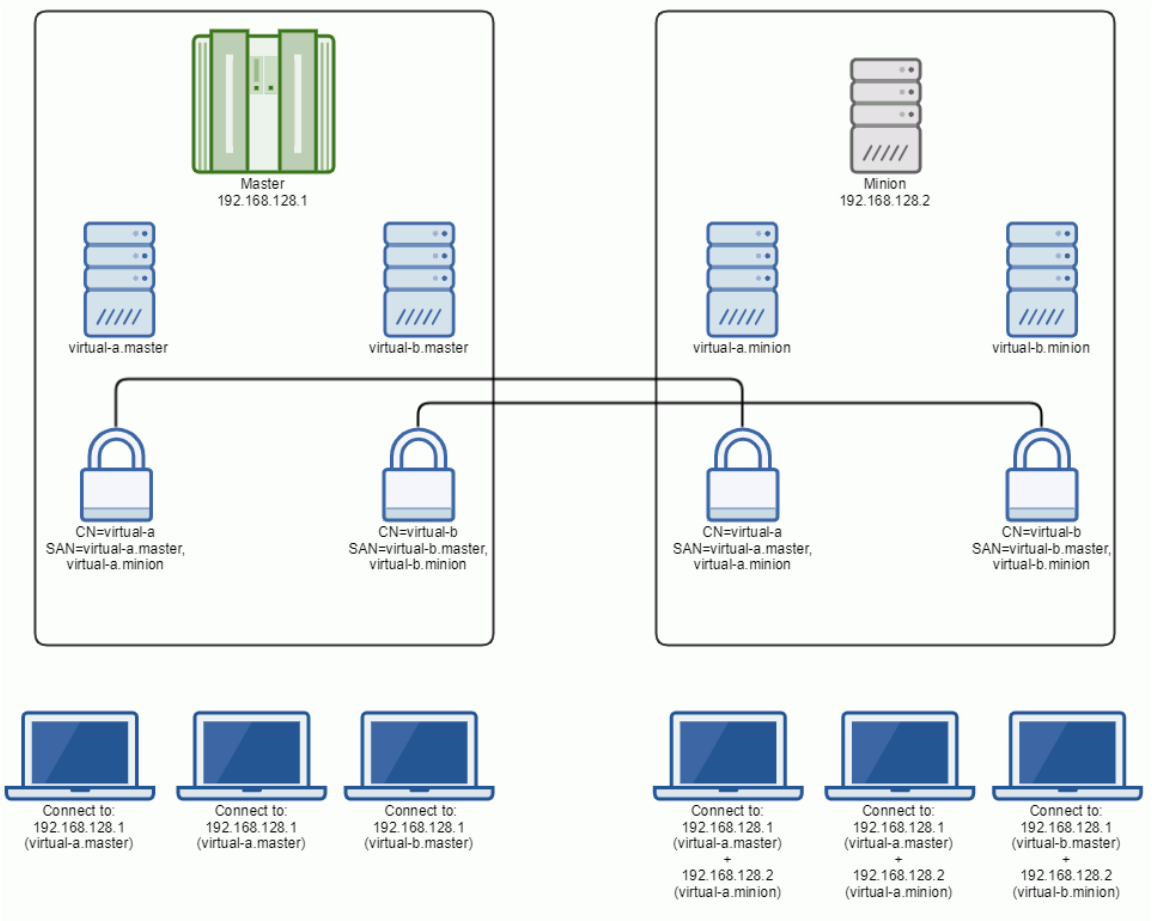

When TLS certificate management is used in a cluster environment:

- Added TLS certificates are automatically copied to minions.
- Virtual server names are replicated on minions.
- Sensors expect the same server certificate to be present on master and the minion.
- Server certificates support the switch from standalone to cluster or adding new nodes without having to re-issue certificates. This is the reason why the certificates must have two distinct SAN DNS entries.
- Sensors need to distinguish only two different virtual addresses to communicate to the servers (master and minion). Because resolution of those virtual names happens internally, and virtual minion addresses can get mapped to different real addresses, there is no need for SAN DNS entries for each individual minion in a cluster.

Managing Certificates on the Server

This section describes the tasks you perform to use the CB Response certificate management features.

You have two opportunities to use server certificates other than the default legacy certificate:

- During server initial installation and configuration, you can substitute your own certificate for the one that would be created by default. See “Substituting a Legacy Certificate during Server Installation” on page 121.
- After the server is installed and configured, you can add certificates through the console. You can do this whether or not you supplied a new legacy certificate during installation. See “Adding Certificates through the Console” on page 122.

When you have the certificates you intend to use in place, you can:

- Choose the validation methods that sensors use for certificates. See “Choosing a Validation Option” on page 122.
- Specify the certificate to use for each sensor group or specify the certificate to use for all sensor groups. See “Assigning Certificates to Sensor Groups” on page 125.

You can add certificates, change validation method, and change certificates assigned to sensor groups later, but implementing an initial certificate configuration as soon as possible may be more efficient and prevent disruptions in server-sensor communication.

Viewing Certificate Information in the Console

Certificate information appears in several places in the CB Response console:

- The Sensors page includes a column showing the certificate for each sensor.
- The Edit Group page for a sensor group shows the certificate assigned to that group.
- The Server Certificates page shows all of the sensor-server certificates that are available on the current server. It also shows the validation method that is being used for these certificates. See “Choosing a Validation Option” on page 122.

To view the available certificates on a server:

1. Click ***Username*** **> Settings**.
2. Click **Server Certificates**.

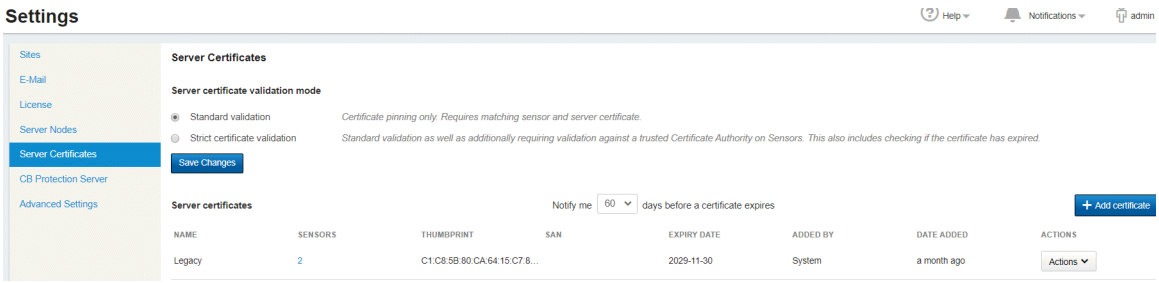

Substituting a Legacy Certificate during Server Installation

When you install a new Cb Response Server, the `cbinit` configuration program you run after installation installs a legacy certificate for use with the standard pinning validation method. By default, this is a certificate that the server produces. As an alternative to the default legacy certificate, you can substitute your own certificate during the server installation process. In either case, the certificate will be named “Legacy” where certificates appear in the console, and it will be protected from deletion.

Important

Certificates and key files added in this way must meet the requirements described in “Server-Sensor Certificate Requirements” on page 116.

When you substitute your own certificate using `cbinit`, CB Response runs tests to confirm that the certificate is valid for this use. If the certificate passes the test, it is used for this server. If it is not valid, the default legacy certificate is used, an error message will appear, and the certificate import failure will be logged to `/var/log/cb/cli`. The `cbinit` process still continues if the substitution fails by using the default certificate instead of the one you tried to substitute.

Note

This procedure is for substituting your certificate for the single, legacy certificate only. If you intend to use more than just the legacy certificate, use the console interface for any additional certificates you need. See “Adding Certificates through the Console” on page 122 for details.

To upload a custom “legacy” certificate during server installation:

1. Prepare the certificate you want to use and place it and its key file in an accessible location on the system hosting the CB Response Server (the master in a clustered environment).
2. Enter the yum install command for installing the correct server version and wait for that process to complete. See the *CB Response Server / Cluster Management Guide* if you need additional instructions for installation.
3. When the installation completes, run the following command, providing the arguments and file paths to the certificate file and the key file as shown here:

   ```
   cd /usr/share/cb
   sudo cbinit --server-cert-file=<certpath> --server-cert-key=<keypath>
   ```

4. If the certificate and key files pass all tests, they become the default server certificate and key, and are copied into the server as `/etc/cb/certs/cb-server.crt` and `/etc/cb/certs/cb-server.key`.

Adding Certificates through the Console

You can add certificates to the CB Response Server through the console to secure server-sensor communications.

Important

Certificates and key files added in this way must meet the requirements described in “Server-Sensor Certificate Requirements” on page 116.

To add a new certificate to a server through the console:

1. Click ***Username*** **> Settings**.
2. Click **Server Certificates** and click the **Add certificate** button.
3. In the **Add certificate** dialog, provide a unique name for the certificate to identify its purpose (use 50 or fewer alphanumeric characters without spaces).

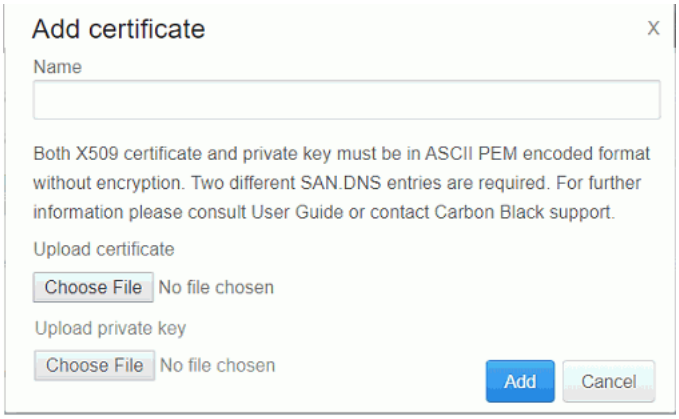

4. Under *Upload certificate*, click **Choose File** and provide the path to a certificate file that meets the requirements that are described in “Server-Sensor Certificate Requirements” on page 116.
5. Under *Upload private key*, click **Choose File** and provide the path to the ASCII PEM-encoded, unencrypted key file for this certificate.
6. When you have entered all required information, click the **Add** button in the dialog.

   If it passes all tests, the new certificate is listed in the table on the Server Certificates page and is available for use by sensors.

Choosing a Validation Option

You can choose one of two validation methods that sensors use for the server certificates that are used to secure server-sensor communication. The validation method can be set through the following console method or by providing a value in the `cb.conf` file for `CbServerSSLCertStrictCheck`, in which case it cannot be changed in the console.

If the standard validation method (certificate pinning only) is used, certificate expiration does not interrupt server-sensor communication, although an expiration warning will appear if this is configured. The only requirement is that the server and sensor certificates match.

If strict certificate validation is used, the requirements of standard validation must still be met, but additional checks are done on the sensor side. A certificate that has expired or fails any other validation requirements causes server-sensor communication to be disabled. See "Sensor Support for Certificate Management" on page 126 for the validation requirements on different sensor platforms.

Caution

Do not enable strict validation if you are using the legacy certificate created during CB Response server installation. Using strict validation for this or any other certificate that cannot pass validation will disable communication between the sensor and server on some sensors that support the certificate management features, and may require uninstalling and reinstalling sensors.

To change the validation method for server certificates:

1. Click ***Username*** **> Settings**.
2. Click **Server Certificates**.
3. Two radio buttons/options appear under Server certificate validation mode:
   - **Standard validation** – Sensors will only require that their certificate matches the server certificate when connecting.
   - **Strict certificate validation** – Sensors will require that a matching certificate is valid on the host machine when connecting. This also includes checking if the certificate has expired.

   If the button for the method you want to use is not selected, click it.
4. Click the **Save changes** button, and click **Confirm**.

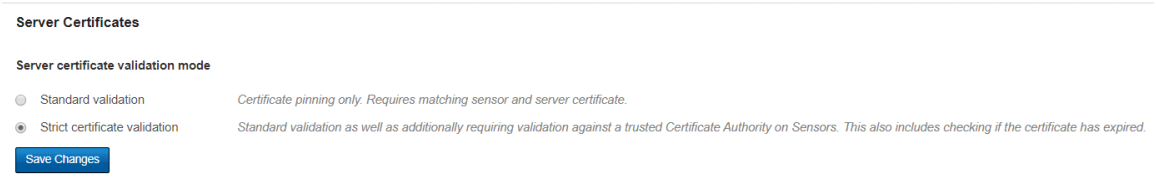

The change will be propagated to all sensors that support TLS server certificate management during their next checkin.

Caution: As the confirmation dialog states, changing validation method can disable communication between sensor and server. Make sure that you have configured certificates properly before changing this setting, especially if you are changing to strict validation.

Changing the Expiration Notification Period

You can configure the CB Response server to display a warning banner when any of its server certificates is ready to expire. The values available for this are: 0, 15, 30, 60 or 90 days. If the value provided is 0 (zero), no warning appears.

This value can also be set through the console as described below or by providing a value in `CbServerCertWarnBeforeExpirationDays` in the `cb.conf` file, in which case it cannot be changed in the console.

If enabled, the warning displays for any expiring certificate listed on the Server Certificates page, even one not used by any sensor groups. If you see warnings for a certificate you are not using and will not use later, delete the certificate to prevent unneeded warnings.

To change the notification period for an expiring certificate:

1. Click ***Username*** **> Settings**.
2. Click **Server Certificates**.
3. In the **Notify me** menu above the table, select the number of days in advance to be warned about expiration of any certificates (whether or not they are being used by any sensors).

Deleting Certificates

You might want to remove a certificate so that it cannot be used (for example, if it has expired or has been compromised). You can remove any certificate except the following:

- You cannot delete a certificate that is currently in use by a sensor group.
- You cannot delete the legacy certificate that is created during server installation.

To delete a certificate from a server:

1. Click ***Username*** **> Settings**.
2. Click **Server Certificates**.
3. Check that the certificate to delete does not have any sensors using it. If the certificate is not in use, click **Actions > Delete** for that certificate.
4. In the confirmation dialog, click **Delete**.

   Caution: After you confirm the deletion of a server certificate, any sensors that were using the certificate can no longer communicate with the server. There is no Undo for this action. Although you cannot delete a certificate that is being used by a Sensor Group, it is possible that an offline sensor could miss a change of certificate for its group, and come back online configured to use a certificate that has been deleted.

Upgrades from Previous Server Releases

When you upgrade to CB Response Server version 6.4.0, the previously used certificate appears in the server certificates table – that is, the certificate called “Legacy”. Unless you change it, standard validation (certificate pinning) remains in effect. This allows the server and sensors to communicate as before the upgrade.

After the upgrade is complete, you can implement a different certificate management strategy if you choose. Subsequent server upgrades maintain whatever certificates you have in place at the time of the upgrade.

Assigning Certificates to Sensor Groups

If new or different certificates are assigned to any sensor group, the change of certificates is made for each sensor the next time it checks in with the server. In addition to using the newly assigned certificate on all subsequent communications with the server, the sensor also stores certificate details locally for use on sensor restarts.

During a change of certificates, the server accepts connections from the sensors utilizing either of two server certificates: the certificate being replaced or the new certificate. Sensor-server communication is not interrupted by certificate replacement. After the connection is successfully established using the new certificate, the old certificate is overwritten and is no longer available for use by the sensor.

If the sensor cannot connect with the new certificate, it reverts to the previous sensor certificate.

For older sensor versions that do not support certificate swaps, the legacy certificate remains in place, regardless of a global or per-sensor-group certificate change. Consider reviewing which sensors support certificate management features before assigning certificates to a group. See “Sensor Support for Certificate Management” on page 126.

In clustered environments, certificate changes are automatically propagated to all servers within a matter of seconds, without requiring a restart.

Assigning different certificates to different sensor groups

Your organization might consist of multiple sites or groups whose endpoints are mapped to different sensor groups. You can use different server certificates to authenticate the connections between the CB Response Server and different sensor groups, thereby reducing the exposure to a compromised server certificate. This lets you manage certificate expiration on a per sensor group basis.

You can also use the per-sensor-group assignment of certificates to gradually change a certificate, even if you want to use the same certificate for all sensors. After you see successful server-sensor communications for one group, you can assign the certificate to all sensors or continue assigning the new certificate on a per sensor group basis.

To change the server certificate for one sensor group:

1. In the navigation bar, click **Sensors**.
2. In the left panel, click the name of the sensor group whose certificate you want to change. Click the **Edit** button.
3. In the **General** panel of the Edit Group page, click the **Assign Server Certificate** dropdown menu to choose the certificate to use for this group.

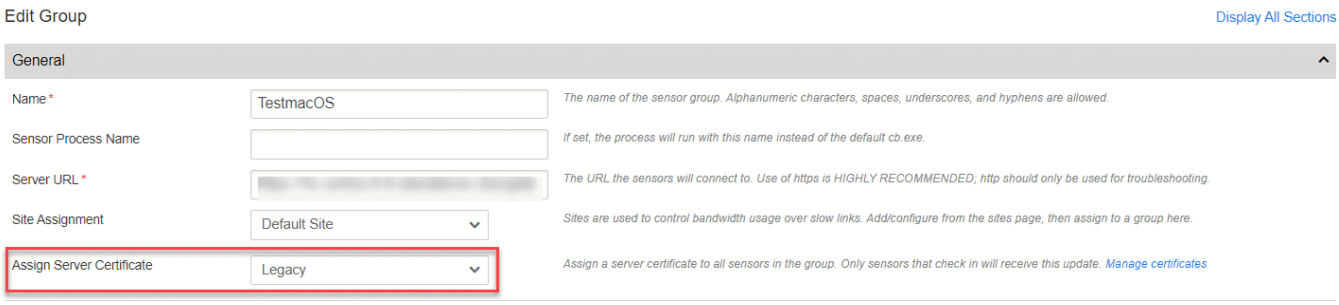

4. Click the **Save Group** button at the bottom of the page.

Assigning a new certificate to all sensor groups

You might need to use your own custom certificate, capable of strict validation, for communication between the CB Response Server and all sensors. If you do not need different certificates for different sensor groups, CB Response provides a single-click method to assign one new certificate to all groups.

> **Note**
>
> Before assigning to all sensor groups, the recommended best practice is to validate certificate connectivity on at least one active sensor group first.

To apply one certificate to all sensor groups:

1. Click ***Username*** **> Settings**.
2. Click **Server Certificates**.
3. Click the **Actions** dropdown menu and click **Assign to all sensor groups**.

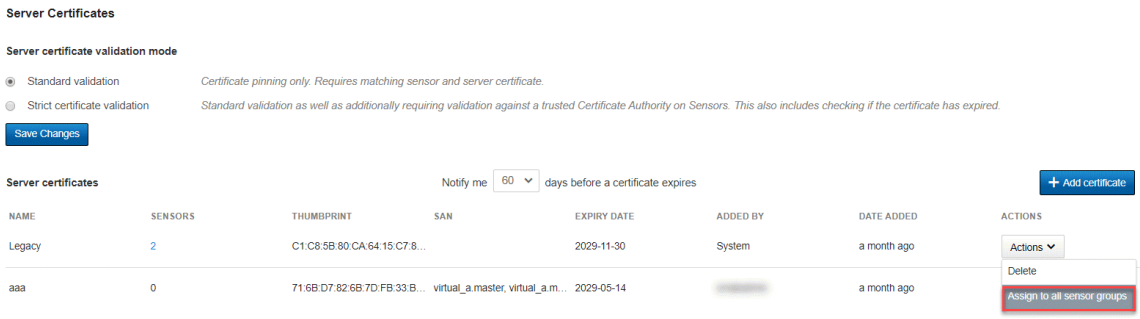

4. Click **Confirm**.

Sensor Support for Certificate Management

Certificate management features are available on CB Response Server versions 6.4.0 and later. How those features affect sensors depends on the sensor version and the OS platform of the sensor. Other than expiration warnings, sensors that don't support TLS certificate management are unaffected by any of the new certificate management settings.

Sensors that do not support certificate swaps continue using the legacy certificate provided by the server, regardless of the certificate assigned to their sensor group.

If you select Standard validation, the only requirement for a valid connection is that there is an exact hash match between the certificate on the sensor and the certificate on the server. If you select Strict validation, the exact hash match is still required, plus additional validation criteria that varies by platform. The following table shows the different validation criteria that are available for the sensor versions on each platform.

The following list shows the sensors that are included with CB Response Server 6.4.0 and their support for certificate management:

- **Windows sensor 6.2.3** – This and later sensors support certificate management and handles strict validation. See “Special Requirement for Windows Sensors” on page 127.

  Windows XP and Windows Server 2003 do not support TLS certificate swap, regardless of the CB Response Sensor version.
- **OS X (macOS) sensor 6.2.5** – This and later sensors support the new certificate management features and handles strict validation as shown in the table below.
- **Linux sensor** – As of the version 7.0.0 server release, Linux sensors do not support certificate management but continue to use the default “Legacy” certificate. Monitor the Carbon Black User Exchange for any news about Linux sensors that support the new features.

Strict validation mode requirements by sensor platform		
Requirement	**OS X Sensor 6.2.5+**	**Windows Sensor 6.2.3+**
Exact certificate match (certificate pinning)	Yes	Yes
Expiration date	Yes	Yes
Certificate validation chain	-	Yes
Hostname matches (SAN=)	-	Yes
Writable host file	-	Yes
Revocation check	-	-
Key Usage is Server Auth (1.3.6.1.5.5.7.3.1)	-	Yes

Special Requirement for Windows Sensors

Certificate swapping on an endpoint running the Windows sensor requires that the sensor is able to update the system `hosts` file. This is a text file that maps IP addresses to hostnames. The file is located at: `C:\Windows\System32\drivers\etc\hosts`.

To be sure the host file can be updated successfully:

- **Check AV Exceptions** -- The CB Response sensor service must be allowed to open and edit the `hosts` file. By default, it has that permission since it is running as administrator. However, other security products (typically anti-virus products or other monitoring tools) must not block the CB Response sensor from accessing the file. If necessary, add exclusions to other security products to allow the CB Response Windows sensor to access the `hosts` file. Failure to do so can result in loss of communications between sensors and server.
- **Save the File in ASCII (Windows Sensor 6.2.3 and 6.2.4)** -- For Windows sensor releases through version 6.2.4, the `hosts` file is assumed to be in ASCII encoding. If the sensor modifies an instance of the file that was saved with non-standard encoding, the file can become unreadable.

  If it has been saved in a different format, resave the file in ASCII. For example, in the Windows Notepad application, choose **Save As**… and then select **ANSI** as the encoding.

See the release notes for post-6.2.4 sensors to determine whether this requirement still applies.

Upgrading to Sensors that Allow Certificate Management

To use the certificate management features of CB Response and upgrade your sensors to a version that is compatible with certificate management, the best practice is to upgrade the sensors first and let the upgrades complete before applying a custom certificate to them. This reduces the possibility of communication issues due to a mismatch between the server certificate and the sensor during the upgrade. After the sensors are updated, you can apply the custom certificate.

Important

If a sensor group is assigned a custom certificate, sensors in that group that support custom certificates cannot be downgraded to sensor versions that do not support custom certificates. Attempts at such a downgrade fail and log an error in the sensorservices debug log.

Chapter 9

Troubleshooting Sensors

This chapter describes ways to troubleshoot sensors. Additional useful troubleshooting information appears in Appendix B, “Sensor Health Score Messages”.

Sections

Topic	Page
Troubleshooting Windows Sensor Installations	130
Troubleshooting Linux Sensor Installations	133
Troubleshooting OSX Sensor Installations	139
Diagnostic Uploads Utility	140

Troubleshooting Windows Sensor Installations

This section describes places to check to troubleshoot errors that could have occurred during Windows sensor installation:

- Confirm that the `%WINDIR%\CarbonBlack\` directory exists. CB Response should be installed in this directory.
- Confirm that the `%WINDIR%\CarbonBlack\InstallLogs\` directory contains installation logs. Review the latest log file for errors.
- Confirm that the current sensor log exists and review it for errors: `%WINDIR%\CarbonBlack\Sensor.log`
- Confirm the settings in the registry key at `HKLM\Software\CarbonBlack\Config`. A typical configuration looks like the following:

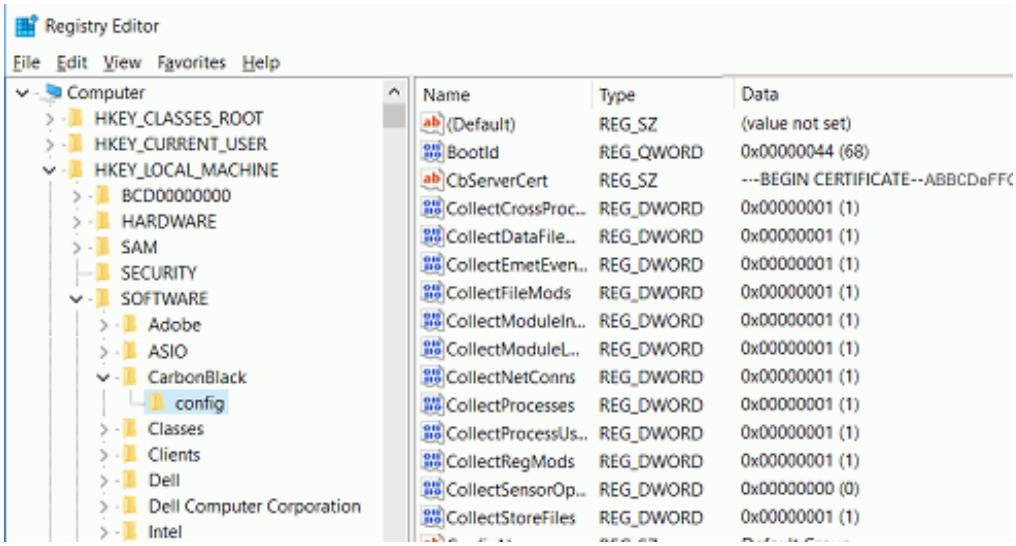

Using Control Codes to Generate Logs of Diagnostic Data

You can use sensor control codes to get diagnostic information.

To issue a sensor control request to the sensor:

1. At a command line prompt, run this command:

   ```
   sc control carbonblack <CONTROLCODE>
   ```

2. Use one of the following codes:

   a. `200` – Initiates a connection attempt to the CB Response server. In most cases, this is a near-immediate connection attempt. Exceptions are during sensor startup and shutdown, or if any outstanding connection or connection attempts to the server are in progress. For example, if an event log or other data is currently being uploaded to the server, or if an attempt to connect to the server is in progress, the attempt does not occur until after the current operation is complete.

   b. `201` – Initiates a dump of diagnostic data to the `%WINDIR%\CarbonBlack\Diagnostics\\` directory. The `201` control code generates the following logs:

Log	Description
EventConverter.log	The internal memory state for event conversion.
EventLogger.log	Top-level event logging statistics.
MachineStatistics.log	General system, process, and kernel statistics.
ModuleInfo.log	Internal module statistics.
NetConnEvents.log	Network event logging statistics. **Note:** To reduce netconn traffic for systems that have a large number of network connections, see "Reducing the Impact of Netconn Data Collection (Windows)" on page 100.
RawEventStats.log	Internal statistics for the conversion of raw events (that were generated by the core sensor driver) to event messages that are stored on the CB Response server.
SensorComms.log	The history of the last 100 network communication attempts between the sensor and the CB Response server.
SensorComponents.log	The current state of the internal sensor components.

The following example demonstrates these conditions:

- Missing `Diagnostics` directory
- `sc control carbonblack 201` and expected `sc.exe` output
- Populated `Diagnostics` directory that contains the `SensorComms.log`

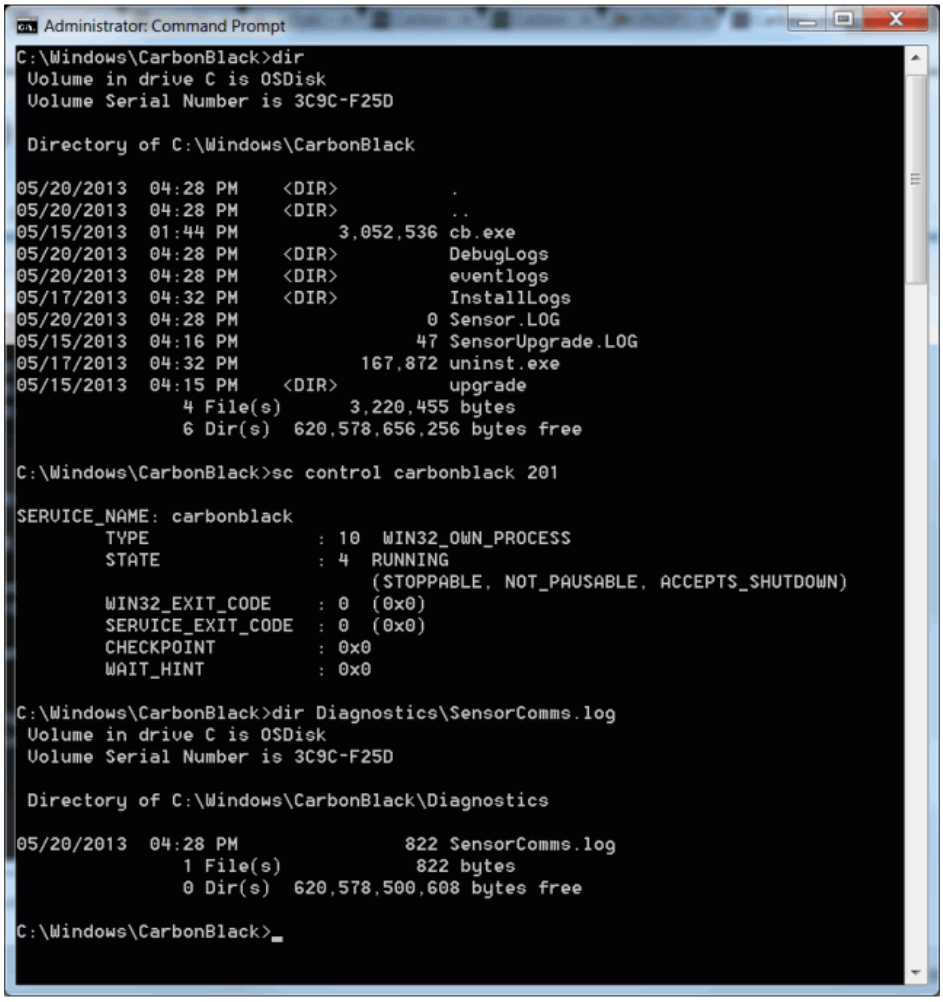

Debugging Sensor Communications

After you run the `sc control carbonblack 201` command, the `%WINDIR%\CarbonBlack\Diagnostics\` directory includes `SensorComms.log`.

This log file contains data in the following format:

```
Server URL: https://x.x.x.x:443

Time                 | URL                                      |
HRESULT
-------------------- + ---------------------------------------- + -
------

2013-05-20 21:28:38 | https://x.x.x.x:443/sensor/register      |
0x00000000

2013-05-20 21:28:38 | https://x.x.x.x:443/sensor/checkin       |
0x00000000

2013-05-20 21:28:38 | https://x.x.x.x:443/data/eventlog/submit |
0x00000000
```

continuation of log:

```
Code | DurationMs | TxBytes | RxBytes | Throttle KB/s
-----+ ---------- + ------- + ------- + -------------
0    | 577        | 300     | 10      | 100
0    | 312        | 402     | 104     | 100
0    | 249        |4328     | 0       | 0
```

The information in the `sensorComms.log` file is described in the following table.

Column	Description
Time	The time (UTC) of the connection attempt.
URL	The URL that was used during the communication.
HRESULT	The result of the operation as a raw HRESULT (0x00000000 is success).
Code	The result processed code. This can vary, based on the HRESULT source, but it can be the HTTP code (404, 500), a Win32 error (net helpmsg code), or other codes.
DurationMS	The duration of the connection attempt in milliseconds.
TxBytes	The number of bytes transmitted, not including HTTP headers.
RxBytes	The number of bytes received, not including HTTP headers.
Throttle KB/s	The rate at which the connection was throttled in kilobytes; 0 indicates that it was not throttled.

Troubleshooting Linux Sensor Installations

This section describes places to check to potential troubleshoot errors during a Linux sensor installation.

General Logging

The user mode portion of the sensor creates an execution log in the following locations:

- For a version 6.1.x Linux sensor:

  `/var/log/cbsensor/cbdaemon.INFO`

  This log file is a symbolic link that is recreated each time the daemon runs. The default log level is set to `WARNING`. This results in the generation of log files for `WARNING` and `ERROR` levels:

  `/var/log/cbsensor/cbdaemon.WARNING`

  `/var/log/cbsensor/cbdaemon.ERROR`

The kernel module logs messages to `/var/log/messages`.

Issue this command in a terminal to dump kernel messages in real time:

```
sudo tail -f /var/log/messages | grep CbSensor
```

- For a version 6.2.x or higher Linux sensor, the current log is:

  `/var/opt/carbonblack/response/log/cbdaemon.log`

  When the current log file reaches a size threshold (currently 100MB), it rolls over to cbdaemon1.log and a new cbdaemon.log is started. You might see log files named `cbdaemon[1-5].log`, with `cbdaemon5.log` being the oldest.
- The kernel module logs messages to `/var/log/messages`.

  Issue this command in a terminal to dump kernel messages in real time:

  ```
  sudo tail -f /var/log/messages | grep CbSensor
  ```

Installation Verification

The following is a list of some key files to check for to confirm installation success:

Path	Description
`/etc/init.d/cbdaemon`	Sensor daemon script
`/usr/sbin/cbdaemon`	Sensor daemon executable
`/lib/modules/$(uname -r)/kernel/lib/cbsensor.ko`	Sensor kernel module (6.1.x)
`/opt/carbonblack/response/modules/cbsensor.ko`	Sensor kernel module (6.2.x)
`/etc/sysconfig/modules/cbsensor.modules`	Kernel autostart file
`/opt/cbsensor/sensordiag.sh`	Sensor diagnostics file (6.1.x)
`/opt/carbonblack/response/bin/sensordiag.sh`	Sensor diagnostics file (6.2.x)
`/opt/cbsensor/sensoruninstall.sh`	Sensor uninstall file (6.1.x)
`/opt/carbonblack/response/bin/sensoruninstall.sh`	Sensor uninstall file (6.2.x)
`/var/lib/cb/config.ini`	Configuration file (6.1.x)
`/var/opt/carbonblack/response/config.ini`	Configuration file (6.2.x)
`/var/lib/cb/sensorsettings.ini`	Settings file (6.1.x)
`/var/opt/carbonblack/response/sensorsettings.ini`	Settings file (6.2.x)

- To verify that the sensor daemon is running, issue the following command:

  ```
  pidof cbdaemon
  ```

  There should be exactly one PID returned.
- To verify that the sensor kernel module is running, issue this command:

  ```
  lsmod | grep cbsensor
  ```

  The output should show one item, if the sensor kernel module is running.

Installation Failures

To check if the sensor is installed correctly, issue this command:

```
rpm -qa cbsensor
```

If the sensor is installed, then a single line will be displayed on your screen showing the version and build numbers. The following is an example:

```
cbsensor-v6.2.0.60603-1.x86_64
```

Note

The version number depends on the version installed.

Sensor Communication History

Running inside a terminal as root and sending the SIGUSR2 signal (via su), issue this command:

```
kill -n 12 $(pidof cbdaemon)
```

The log is located at `/var/tmp/cb/sensor_comms.log`. Each transaction has a `HRESULT`, which can be one of the following:

Facility Number	Description	Error Code Value
204	OS level errors	Maps to errno
25	HTTP errors	HTTP error code
200	Curl errors	See http://curl.haxx.se/libcurl/c/libcurl-errors.html.
201	Curl form errors	

Manual Sensor Daemon Start and Stop

- To restart the service, open a terminal and issue this command:

  ```
  sudo service cbdaemon restart
  ```
- To start the service, open a terminal and issue this command:

  ```
  sudo service cbdaemon start
  ```
- To stop the service, open a terminal and issue this command:

  ```
  sudo service cbdaemon stop
  ```

Note that if the sensor crashes, you may need to delete the following file to eliminate an error when trying to start the daemon: `/var/run/cbdaemon.pid`

Determine Server URL

To determine the server URL used by the sensor, follow the instructions in "Sensor Communication History" on page 135 to create a communication log and dump the contents of the generated log file. The server URL appears at the top.

Initiate an Immediate Checkin to the Server

Running inside a terminal as root and sending the SIGUSR1 signal (via su), issue this command:

```
kill -n 10 $(pidof cbdaemon)
```

Driver Debug Parameters

Two arguments can be passed to the driver to control the debug behavior:

- `g_traceLevel` – Controls debug trace output flags (6.1. only).
- `g_eventFilter` – Controls which event types are generated.

These arguments can be passed in the `/etc/sysconfig/modules/cbsensor.modules` file or as described in the version-specific sections below.

Sensor Version 6.1.x Driver Debug Parameters

For version 6.1.x sensors, use the following command:

`sudo insmod cbsensor.ko g_traceLevel=<value> g_eventFilter=<value>`:

```
insmod cbsensor.ko g_traceLevel=0x00200000
```

or

```
modprobe cbsensor g_traceLevel=0x00200000
```

where `0x00200000` is hook tracing

The full list of levels is:

```
#define  DL_INIT           0x00000001
#define  DL_SHUTDOWN       0x00000002
#define  DL_WARNING        0x00000004
#define  DL_ERROR          0x00000008
#define  DL_INFO           0x00000010
#define  DL_REQUEST        0x00000100
#define  DL_HOOK           0x00200000
#define  DL_VERBOSE        0x08000000
#define  DL_ENTRY          0x10000000
#define  DL_EXIT           0x20000000
```

To create the log level mask you want, `OR` the individual levels.

Sensor Version 6.2.x Driver Debug Parameters

For version 6.2.2 sensors and later, the preferred method for setting the debug logging subsystem level is by editing `/proc/debug-subsystems`. In this file, the debug subsystem value is reported in hexadecimal as an 'or' of the values shown below.

An alternative to editing debug-subsystems is to use the following command:

```
sudo insmod cbsensor.ko g_debug_subsystem=<value>
g_eventFilter=<value>
```

The subsystems are defined as:

```
#define DS_COMMS 0x0002
#define DS_FILE 0x0004
#define DS_HASH 0x0008
#define DS_BAN 0x0010
#define DS_HOOK 0x0020
#define DS_ISOLATE 0x0040
#define DS_LOG 0x0080
#define DS_LSM 0x0100
#define DS_MOD 0x0200
#define DS_NET 0x0400
#define DS_PROC 0x0800
#define DS_PROCFS 0x1000
#define DS_TEST 0x8000
```

If DS_TEST is specified, the enabled log levels and subsystems will be logged to the system log (use `dmesg` to view). This is done at module load time when specified as a parameter.

Subsystem debug messages are reported at KERN_INFO level. The minimum kernel logging level must be at least KERN_INFO for these messages to be reported. It can be set at runtime via `/proc/sys/kernel/printk` or at boot time by setting the `loglevel` kernel command line parameter.

Daemon Debug Options

Debugging Parameters for 6.1.x Sensors

For 6.1.x version sensors, there are options for the CB Response initialization file that you can use for daemon debugging. This file is located at:

```
/var/lib/cb/sensorsettings.ini
```

When you set one or both of these options, you can use SIGHUP to reread the sensorsettings.ini file and update the log settings.

```
$ sudo killall cbdaemon -SIGHUP
```

The two debugging options are:

- **DaemonLogLevel** -- You can set this to any of the following: Error, Warning, Info, Debug1, Debug2, Debug3, Debug4, Debug5. DaemonLogLevel is not set in the ini file, the default DaemonLogLevel is Warning.
- **SensorLogLevel** -- This option is a string value in the form

  ```
  <base level>[/<extra level>[/..]]
  ```

  where the possible values of <base level> are Error, Warning, Info, and Debug, and the possible values of <extra level> are Hook, Request, Entry, Exit, Comms, and Trace.

  In most cases only <base level> will need to be set. If SensorLogLevel is not set in the ini file, the default is Warning.

Notes

If you deployed sensor versions prior to 6.1, be aware of the following change:

- Due to the additional granularity now available in the daemon, many of the debugging messages have been decreased in priority.
- The Info level can provide a good overview of observed events without providing excessive output.
- Some of the log messages are more uniform. The most notable change is that "FILE EVENT" messages now display the type of file event as a string instead of an integer value. For example, "FILE_OPEN EVENT" and "FILE_CLOSE EVENT".

Debugging Parameters for 6.2.x Sensors and Later

For sensors at version 6.2.1 and later, use the preceding procedures described for daemon debugging, noting the following changes:

- The initialization file that you edit is in a different location than earlier sensors:

  `/var/opt/carbonblack/response/sensorsettings.ini`
- The **DaemonLogLevel** values were changed for the 6.2.x sensor. The new levels are: None. Error. Warning. Info. Verbose
- As of version 6.2.2, **SensorLogLeve**l is no longer a valid option and will be ignored.

Determine Sensor Version

To determine the version of cbdaemon running, from a terminal, issue this command:

```
cbdaemon -v
```

Trigger a Diagnostic Data Dump

Running the `sensordiag.sh` script dumps and collects network event logs, cbdaemon log files, and important sensor and system configurations that can help diagnose sensor issues. The script packages up these files into a single compressed file that you can deliver to Carbon Black for analysis.

Generate a `.tar.gz` file in the current directory for diagnostic purposes by issuing this command:

```
sudo /opt/carbonblack/response/bin/sensordiag.sh
```

Troubleshooting OSX Sensor Installations

This section describes places to troubleshoot potential errors during OSX sensor installation.

Installation Verification

The following is a list of key files that should be present if installation is successful:

Path	Description
`/Applications/CarbonBlack/CbOsxSensorService`	Sensor service
`/Applications/CarbonBlack/sensoruninst.sh`	Uninstall script
`/System/Library/Extensions/CbOsxSensorNetmon.kext`	Network monitor
`/System/Library/Extensions/CbOsxSensorProcmon.kext`	Process monitor
`/var/lib/cb/sensorsettings.ini`	Settings file

Installation Failures

The installation process can fail if the `sensorsettings.ini` is not located in the same directory as `Installer.pkg`.

If the installation does not complete successfully, the installer reverts all the changes made to the system but leaves the `cblog.log` file intact. For troubleshooting, collect the installer log file created at `/var/log/cblog.log` and send it to Community Resources for assistance.

Communications Logging

1. Determine the PID of the CB Response sensor by issuing this command:

   ```
   ps -ax | grep CbOsxSensorService
   ```

2. Start the communications log dump by issuing this command:

   ```
   sudo kill -s USR2 <pid of CbOsxSensorService>
   ```

You can locate the log at `/var/lib/cb/sensor_comms.log`. Each transaction has a `HRESULT` (see description at http://msdn.microsoft.com/en-us/library/cc231198.aspx) that can be one of the following:

Facility Number	Description	Error Code Value
203	OS level errors	Maps to errno
25	HTTP errors	HTTP error code
200	Curl errors	See http://curl.haxx.se/libcurl/c/libcurl-errors.html.
201	Curl form errors	

Manual Sensor Daemon Start and Stop

To manually start and stop the sensor daemon service, open a terminal and issue these commands:

```
sudo launchctl unload /Library/LaunchDaemons/
com.carbonblack.daemon.plist
```

```
sudo launchctl load /Library/LaunchDaemons/
com.carbonblack.daemon.plist
```

Determining Sensor Version

To determine the sensor's version, open a terminal and issue this command:

```
/Applications/CarbonBlack/CbOsxSensorService -v
```

Determine Server URL

To determine the server URL used by the sensor, follow the instructions in "Installation Failures" on page 135 to create a communication log and dump the contents of the generated log file. The server URL appears at the top.

Initiate an Immediate Checkin to the Server

To initiate an immediate checkin to the server, open a terminal and issue this command:

```
sudo kill -s USR1 <pid of CbOsxSensorService>
```

Initiate a Diagnostic Data Dump

You can run a command to dump and collect network event logs, cbdaemon log files, and important sensor and system configurations that can help diagnose sensor issues. The files are packaged up into a single compressed file that you can deliver to Community Resources.

To dump a communication and event tracking logs, open a terminal and issue this command:

```
sudo kill -s USR2 <pid of CbOsxSensorService>
```

Diagnostic Uploads Utility

Beginning with CB Response server version 6.2.2 and macOS (OS X) sensor version 6.2.0, a new sensor diagnostics tool can collect diagnostic data packages from the endpoint and upload them to a cloud location for analysis, using the CB Response server as an intermediary. This data can help Carbon Black representatives troubleshoot crashes, performance problems, or other situations in which you believe there is an issue with a sensor. This feature as available on both on-premise and cloud servers. It is currently available only on the latest macOS (OS X) sensor.

There are three different categories of data that can be uploaded using this feature:

- **Crash data (automatic or manual):** This option returns crash reports for Carbon Black user-mode Service and Sensor Diags. You can choose to package and upload

crash data manually or set it for automatic packaging and upload when there is a crash.

- **Diagnostics data (manual):** This option returns information about the sensor. The data includes a sample of the Carbon Black user-mode Service, Carbon Black user-mode Service statistics, cblog.log (installer log), any diag files for Carbon Black user-mode Service, system log messages containing "Carbon Black" and all daemon log files. This can be useful for situations in which, while there has not been a crash, other behavior suggests a problem in sensor operation. This option must be run manually.
- **Environment data (manual):** This option returns a list of all open files, a list of all running processes and the amount of CPU they are using, and computer information including Power-On Self Test, Memory, System Software Version, Boot Device, Computer Name, User Name, and a list of all kernel extensions. This option must be run manually.

Automatic Crash Data Upload

By default, a Carbon Black customer service representative must ask you to generate and manually upload (or provide ssh access to) diagnostic files. In a crash situation, this can lead to time-consuming back and forth when you need to get a system running and protected again as quickly as possible. Beginning with server release 6.2.2, diagnostics from sensor crashes may be collected and uploaded automatically for storage in a cloud location. This provides rapid access to the data by Carbon Black sensor experts when an issue requires troubleshooting, without requiring additional steps.

- Automatic upload of sensor crash data is disabled by default. You must opt-in to enable it.
- This feature is currently available only on OS X sensors.

When a sensor checks-in with a Response server, it receives the current setting for **Allow Upload of Sensor Diagnostics Data**. If the setting Manual or Automatic, the sensor will allow manual uploads of diagnostic files. In the case of Automatic, if there is crash data available, the sensor will initiate upload of that data. The server will reject any uploads if Disabled is chosen.

Manual Upload Option (Command Line Utility)

If you choose **Manual** for Allow Upload of Sensor Diagnostics Data on the Shared Settings page, you initiate data collection and uploads by executing a command line utility. The command line syntax for macOS/OS X is as follows:

```
sensordiag -type CDE [-startdate YYYY-MM-DD [THH:MM:SSZZZZ] ]
[-enddate YYYY-MM-DD [THH:MM:SSZZZZ] ] [-upload [<number of
seconds>] ] [-remember ]
```

The “type” options determine which type(s) of data is uploaded:

- C: Crash reports for Carbon Black user-mode Service and Sensor Diags
- D: Diagnostics reports
- E: Environment reports

The other options are:

- `startdate/enddate`: For manually collected sensor diags, you can specify the range of diagnostic files to include in the zip file. This is based on their modified date (date created or dates inside files are not considered for this parameter).

- `upload`: When you run the sensor diagnostics command manually, this option must be specified if you want the resulting zip file uploaded to the CB Response server – otherwise it just remains on the sensor. If a time argument is specified, the tool will only look for files that were modified within the start and end dates specified. If a time argument is not specified, the tool will capture logs from the beginning of the day until the current time.
- `remember`: This option uses the end date of the most recent sensor diagnostics zip file as the startdate for a new one.

Enabling Sensor Diagnostics Uploads

To use the sensor diagnostic uploads feature, you first enable it through the Shared Settings page (click ***Username*** **> Sharing Settings**). There are three options for crash data uploads, and two options for diagnostics and environment data:

- **Disabled** – The default setting for each diagnostics type. Neither automatic crash file uploads nor manual triggering of any uploads is available.
- **Manual** – You manually start the `sensordiag` utility tool via the command line. No data is automatically uploaded. Running the tool collects and uploads the data type that is specified by a command-line switch (Crash, Diagnostic, Environment, or any combination).
- **Automatic** – The sensor automatically collects and uploads crash data when a crash is detected on the sensor. Although diagnostic and environment data cannot be uploaded automatically, selecting **Automatic** also enables the `sensordiag` utility so that you can manually collect and upload these data types.

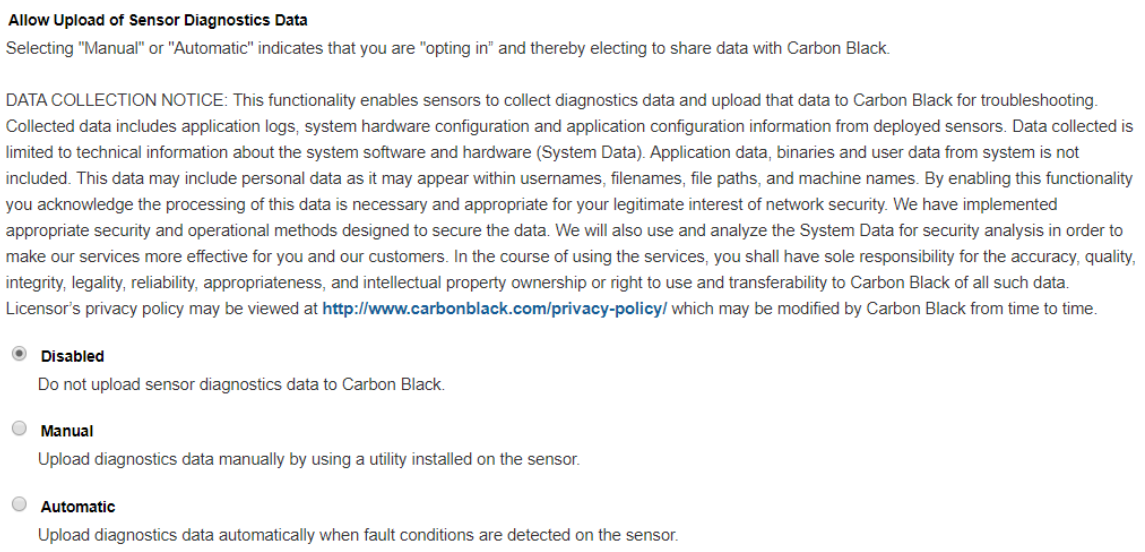

File Transfer and Security

Each collection of sensor diagnostic files is packaged as a zip file on the sensor before uploaded to the CB Response server. From the CB Response server, uploaded sensor diagnostic files are sent to the Carbon Black cloud, and are encrypted until and unless they are accessed by authorized Carbon Black representatives.

To avoid sending oversized files over an HTTP request, the uploaded file size is limited to 5MB. Files larger than that are split into multiple files during transit.

Files are uploaded to the CB Response server on first-come-first-serve basis. Only one file can be uploaded to a server at a time. If additional sensors check in for an upload while an upload is already in progress, a “Try Again” message is sent to sensor so that it can try again at a later time. In a clustered environment, multiple simultaneous uploads are possible.

After the upload succeeds, the zip file is removed from the sensor.

Data Collected by Sensor Diagnostics

Your organization might have specific standards for what kind of information may be uploaded to a third party such as Carbon Black. To help you decide whether the sensor diagnostic upload features meet those standards, the following table shows the type of information each option will upload.

	Crash Logs	Diagnostic Reports	Environment Data
File name	No	Yes (1)	Yes (5)
File path	No	Yes (1)	Yes (5)
IP address	No	Yes (1)	No
Command line	No	Yes (2)	No
Network operations (device names)	No	Yes (3)	No
File writer name	No	No	No
Audit logs (user name)	No	No	No
Username associated with process	No	No	No
Hostname	Yes (4)	Yes (4)	Yes (4)
Full binary	No	No	No
File metadata	No	No	No
Email address	No	No	No

	Crash Logs	Diagnostic Reports	Environment Data

(1)- File names and file paths (which include the file name) and IP addresses appear in `CbOsxSensorService` log files.

(2)- Command lines appear in log files in certain error situations.

(3)- Device names appear in the log together with IP addresses (1).

(4)- The hostname is part of the name of the zip file that is sent to Carbon Black.

(5)- System logs are collected which may contain path names logged from other processes.

(6)- Hostname is shown in system profiler information.

Chapter 10

Responding to Endpoint Incidents

This chapter describes how to respond to endpoint incidents by isolating endpoints by using CB Live Response and banning process hashes.

Sections

Topic	Page
Overview of Incident Response	146
Isolating an Endpoint	147
Using Live Response	150
Extending Live Response	161
Live Response Activity Logging and Downloads	161
Banning Process Hashes	162

Overview of Incident Response

When you discover a malicious file or process on your endpoint(s) using CB Response, you can address the issue in a variety of ways. CB Response provides the following methods for responding to threats directly from the console:

- **Endpoint Isolation** – You can isolate an endpoint from the rest of the network, leaving only the connections that are needed for access to its sensor by the CB Response server.
- **CB Response Live Response (Live Response)** – Live Response opens a command interface for direct access to any connected host running the CB Response sensor. Responders can perform remote live investigations, intervene in ongoing attacks, and instantly remediate endpoint threats.
- **Process Hash Banning** – You can ban a process hash so that the process cannot be run again on hosts reporting to this CB Response server, and any running version of it is terminated.

These features can be used together or separately. For example, you can isolate an endpoint immediately to prevent the spread of the problem and then use Live Response to end the process and perform any other file removal or needed repairs.

On the other hand, if the incident is not ongoing, isolation might not be necessary. In that case, you can use Live Response to remediate or further investigate the issue on affected endpoints, or simply ban the hash for the malicious process.

CB Response does not present a message on the affected endpoint when any of these features is used on an affected sensor. With endpoint isolation, a user would likely become aware quickly that they had lost network access but would not know why. With Live Response, actions you take on a computer might affect a user's access to files or programs, but there would be no indication that CB Response tools are responsible, unless you have chosen to make the user aware of that. Also, when there is an attempt to run a process that is banned by hash, the operating system might display a dialog indicating a lack of access, or the process might silently fail to run.

If you also have the CB Protection agent on your endpoints, you can use CB Protection control features to investigate incidents and modify rules to prevent future occurrences. See "Integrating CB Response with CB Protection" in the *CB Response Integration Guide* for details on the features that are available when the two platforms are connected.

Note

To use the features described in this chapter, a user must be one of the following:

- a user that has the enhanced Analyst permission for the feature and is a member of a team that has the Analyst role for the sensor group for the endpoint being acted upon (or for any sensor group to ban hashes)
- for on-premise installations, a Global Administrator
- for cloud installations, an Administrator

See "Managing User Accounts for On-Premise Servers" on page 51 or "Managing User Accounts for Cloud Servers" on page 65 for more information about user roles and privileges.

Isolating an Endpoint

You can isolate one or more endpoints from the rest of your network and the Internet through the CB Response console. When an endpoint is isolated, its connectivity is limited to the following (unless you have created network isolation exclusions as described in “Isolation Exclusions” on page 149):

- The CB Response server can communicate with an isolated computer.
- To allow the sensor to communicate with the CB Response server, ARP, DNS, and DHCP services remain operational on the sensor’s host. (For Windows operating systems prior to Vista, ICMP (for example, ping) will remain operational.)
- DNS and DHCP are allowed through on all platforms. This is required for proper communications to the CB Response server. Protocols are allowed by UDP/53, UDP/67, and UDP/68.
- ICMP is allowed on the following operating systems:
  - Windows (operating systems prior to Vista)
  - OSX
  - Linux
- UDP is blocked on all platforms.

In order to isolate an endpoint, you must be one of the following:

- a user with the enhanced Analyst permission for isolating sensors and is a member of a team that has the Analyst role for the sensor group of the endpoint being acted upon
- for on-premise installations, a Global Administrator
- for cloud installations, an Administrator

To isolate endpoints:

1. Log in as a Global Administrator (on premises) or Administrator (cloud), or as a user who has Isolate sensor permission and membership on a team with the Analyst role for the endpoint to isolate.
2. On the navigation bar, click **Sensors**.
3. Check the box next to each endpoint to isolate.
4. From the **Actions** drop-down list, select **Isolate**.
5. Click **OK** to confirm that you want to isolate these endpoints.

   The endpoint is isolated from all but the CB Response server and the network services that are required to connect the two, in addition to any addresses that are allowed due to network isolation exclusions.

When you designate an endpoint for isolation, its status on the server first moves into in “isolation configured” state waiting for its next check-in. Because of this, several minutes can pass before the endpoint is actually isolated. When it checks in, the server tells the sensor to isolate the endpoint, and when the sensor responds, its state changes to “isolated”.

After it is isolated, endpoints normally remain isolated until the isolation is ended through the console. However, if an isolated system is rebooted, it is not isolated again until it checks in with the CB Response server, which could take several minutes.

Having isolated endpoints, you can proceed with remediation steps. For example, you might use Live Response to investigate or modify an endpoint. When you are finished, restore connectivity to the endpoints that you isolated.

To end network isolation for endpoints:

1. Log in as a Global Administrator (on premises) or Administrator (cloud), or as a user with Isolate sensor permission and membership on a team with the Analyst role for the endpoint to restore.
2. On the navigation bar, click **Sensors**.
3. Check the box next the endpoints for which to restore network connectivity.
4. From the **Actions** drop-down list, select **Remove isolation**.
5. Click **OK** to confirm the restoration.

   The computers return to the network with the same access they had before they were isolated (unless you made access changes through Live Response).

Isolation Exclusions

Starting with CB Response version 6.5.0, Windows sensors version 6.2.4 and higher and OSX sensor versions 6.2.7 and higher support isolation exclusions. You can add one or more IPv4 addresses or domain URLs that isolated sensors can access in isolation mode, in addition to the CB Response server. This setting is applied on a per-sensor-group basis.

This feature is disabled by default; to enable it, you must edit the `cb.conf` file. See the *CB Response cb.conf Guide* for instructions.

To create an isolation exclusion:

1. Log in as a Global Administrator (on premises) or Administrator (cloud), or as a user who has *Isolate sensor* permission and membership on a team with the Analyst role for the sensor group to which you want to add exclusions.
2. On the navigation bar, click **Sensors**.
3. Click the gear icon next to the sensor group for which you want to add isolation exclusions.
4. Click **Isolation Exclusions** and then click **Add Exclusion**.
5. Enter a description that identifies the exclusion (50 character maximum), and the IPv4 address or domain URL that specifies the exclusion (253 character maximum).

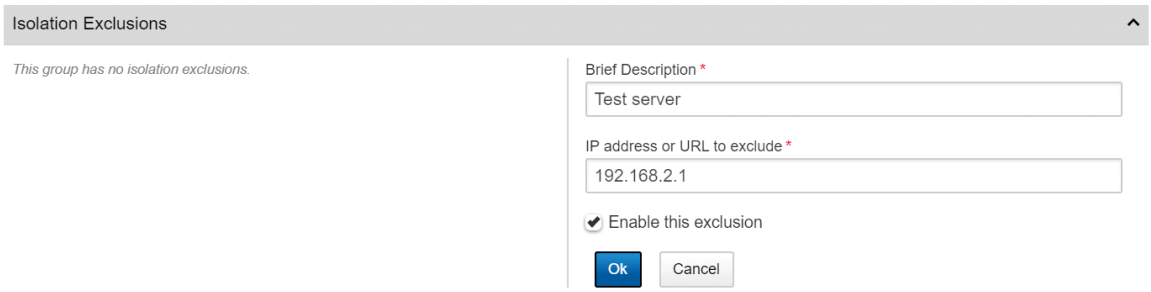

6. Select **Enable this exclusion** and click **OK**.
7. Click **Save Group**.

After you have created an exclusion, you can edit it by clicking the pencil icon, or you can remove the exclusion by clicking the trash can icon.

> **Note**
>
> Duplicate exclusions are not allowed. If you enter the same IP address or URL for more than one exclusion, the last entry that was submitted is retained, but the duplicated entry is removed.

Using Live Response

Live Response opens a command line interface for direct access to any connected host that is running the CB Response sensor. Responders can perform remote live investigations, intervene in ongoing attacks, and instantly remediate endpoint threats. For example, Live Response allows a responder to view directory contents, kill processes, modify the registry, and get files from sensor-managed computers.

Live Response is disabled by default on newly installed CB Response systems. "Enabling and Configuring Live Response" on page 150 describes ways to enable the feature and adjust its data settings.

After Live Response is enabled on a CB Response server, a user must be one of the following to use it:

- a user with the enhanced Analyst permission for Live Response and is a member of a team that has the Analyst role for the sensor group for the endpoint being accessed
- for on-premise installations, a Global Administrator
- for cloud installations, an Administrator

Important

The Live Response feature should be used in full compliance with your organization's policy on accessing other user's computers and files. Consider the capabilities described here before giving users access to the feature and choosing the Sensor Group in which you will place computers.

If you do not want console administrators for on-premises installations to be able to activate Live Response, make sure `CbLREnabled=False` is set in your `cb.conf` file and is not commented out.

There are two Live Response modes:

- **Attached Mode** – When you activate Live Response for a specific endpoint, you create and attach to a *session*. The interface for a session includes information about the endpoint and a command window for interacting with the endpoint. See "Live Response Endpoint Sessions" on page 152.
- **Detached Mode** – You can enter Live Response without being attached to a particular session through the **Go Live** command on the console menu. This interface includes commands to manage and access existing sessions as well as commands that are useful outside of a session. See "Detached Session Management Mode" on page 160.

Enabling and Configuring Live Response

There are two ways to enable and disable Live Response: through the `cb.conf` file or (if not set in `cb.conf`), through the CB Response console.

If CbLREnabled has no value (or is commented out) in the cb.conf file, an administrator can enable or disable Live Response in the console using a switch on the Advanced Settings page. *This is the on-premises default in version 6.3.0 and later.*

To enable or disable Live Response via the CB Response console:

1. Log in as a Global Administrator (on premises) or Administrator (cloud).
2. Click **Username > Settings**.
3. Click **Advanced Settings** and scroll to the **CB Live Response** section.
4. Check or uncheck **Enable CbLR** and click the **Save changes** button.

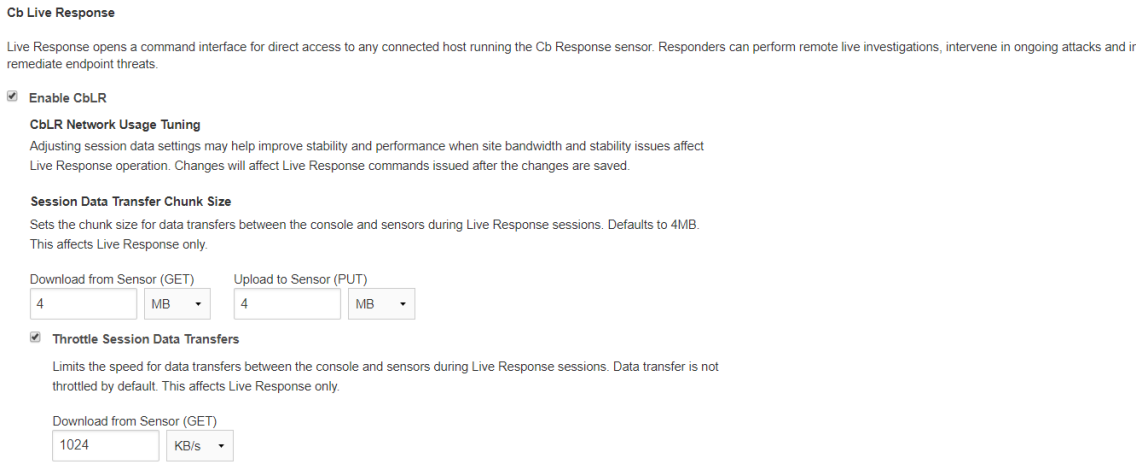

Note: If the **Enable CbLR** box is grayed out and unresponsive, the value is set in `cb.conf` and cannot be changed via the console.

For on-premise servers, you can edit the `cb.conf` file to fix the state of Live Response so that it cannot be enabled or disabled through the console.

To enable or disable Live Response via cb.conf (on-premise servers only):

1. On the CB Response server, open `/etc/cb/cb.conf` for editing.
2. Add or uncomment the following setting in the cb.conf file and set its value:

   ```
   CbLREnabled=True
   ```

   -or-

   ```
   CbLREnabled=False
   ```

3. Save the `cb.conf` file.
4. You must stop and restart the standalone server or cluster to make the new setting effective:
   - For standalone server:

     ```
     sudo service cb-enterprise restart
     ```

   - For clusters:

     ```
     sudo cbcluster stop
     ```

     (...wait for all the nodes to shut down, and then...)

     ```
     sudo cbcluster start
     ```

Tuning Live Response Network Usage

Sites that have bandwidth or stability issues might experience performance problems or failures with Live Response. To help mitigate these issues, you can adjust the data transfer chunk size and also enable throttling of data transfers between the CB Response console and sensors during a Live Response session.

Changes to these settings affect only Live Response, and are effective only on commands that you issue after the changes are saved. The settings apply to all users.

To modify Live Response network usage:

1. Log in as a Global Administrator (on premises) or Administrator (cloud).
2. Click ***Username*** **> Settings** .
3. Cick **Advanced Settings** and scroll to the **CB Live Response** section.
4. Enable Live Response if it is not already enabled (you cannot modify settings when it is disabled).
5. In the **CbLR Network Usage Tuning** section, under *Data Transfer Chunk Size*, set new **Download from Sensor (GET)** and/or **Upload to Sensor (PUT)** values. The default value for both settings is 4MB.
6. In the same section, check the box for **Throttle Session Data Transfers** if you want to activate throttling. Throttling is turned off by default, and it must be turned on before you can edit its value.
7. Set the throttling speed in **Download from Sensor (GET)** to a new value. The default setting when throttling is first activated is 512KB.
8. Click the **Save changes** button.

Live Response Endpoint Sessions

To access an endpoint using Live Response, a user must either have Global Administrator (or cloud Administrator) privileges, or be on a team with the Analyst role for that endpoint. A "session" must first be created with the sensor you want to access. A session indicates that the sensor is connected to the CB Response server to receive real-time commands.

Sessions are created and attached automatically when you click the **Go Live** button on the **Sensor Details** or **Process Analysis** pages. If you enter the Live Response console using the **Go Live** command from the console menu, access to an endpoint requires that you first create and attach a session:

```
[Live Response]# session new [sensor_id]
[Live Response]# attach [provided_session_id]
```

You can have sessions with multiple sensors active at the same time. Use the `detach` command to detach from a session but leave it active.

Use the `session close` command to end a session with the sensor. Sessions will timeout when they are not attached and active for five minutes.

Each session has a unique numeric ID. Up to 10 sessions can be running at one time, and multiple users can be attached to the same session.

Note

More than one CB Response console user can attach to the same session with an endpoint at the same time. If more than one user submits a command through the session at approximately the same time, each command must finish executing before the next one can begin. Also, one user can undo or otherwise modify what another user is doing. Consider this if more than one user has Live Response access to an endpoint.

To create and attach to a Live Response sensor session:

1. Log in as a Global Administrator (on premises) or Administrator (cloud), or as a user with Live Response permission and membership on a team with the Analyst role for the endpoint you want to access.
2. On the navigation bar, click **Sensors** and then click the name of the endpoint to access.
3. Click **Go Live**.

The **Live Response** page appears with a command window on the left and an information panel on the right. The command window prompt shows the name of the host and the current directory in which Live Response is active. The information panel includes:

- **Host Details**
- **Alerts** related to the host
- **Running Processes** on the host

A status indicator (dot) and a message appear immediately above the command window. The dot has the following color code:

- **Green** – The sensor is connected and a session has been established. The host name is shown.
- **Orange** – The CB Response server is waiting for the sensor to check in, or no host is connected because no session is attached.
- **Gray** – A session cannot be established with the sensor because the host is offline, the sensor is disabled, or the sensor is not a version that supports Live Response.

4. To view a list of the available commands, click in the command window area and enter the `help` command. You can get information about a specific command by entering:

   `help` *`commandname`*

The following table shows the complete set of Live Response commands. In the descriptions, *remote host* refers to the host that is being accessed through Live Response, and *local host* refers to the host on which the user is running the CB Response console. These commands are all run in the SYSTEM context.

Command	Description
archive	Obtain an archive (gzip tarball) of all the session data for this session, including commands run and files downloaded. The archive is downloaded to the computer on which you are running the CB Response console using the browser's download method.
argparse	Test how Live Response parses CLI arguments. This command helps determine if there are any interpretation issues. For example, it can reveal whether spaces or other special characters are properly escaped.
cd ***[dir]***	Change the current working directory. Options include absolute, relative, drive-specific, and network share paths.
clear	Clear the console screen; the `cls` command can also be used for this purpose.
delete ***[path]***	Delete the file specified in the path argument. The file is permanently deleted, not sent to the Recycle Bin.
detach	Detach from the current Live Response session. If a session has no attachments, it remains live until it times out (five minutes by default).
dir	Return a list of files in the current directory or the specified directory if it is added to the command, (for example, `dir c:\temp` or `dir /tmp`)
drives	List the drives on the remote host. This is for Windows only.
exec[processpath]	Execute a background process specified in the processpath argument on the current remote host. By default, process execution returns immediately and output is to stdout and stderr. Options may be combined: • **exec** `-o outputfile processpath` – Redirect the process output to the specified remote file, which you can download. • **exec** `-w processpath` – Wait for the process to exit before returning. You could combine the options as shown in the example below to execute and capture the output from a script: **exec** `-o c:\output.txt -w c:\scripts\some_script.cmd` You must provide the full path to the process for the processpath argument. For example: `c:\windows\system32\notepad.exe`

Command	Description
execfg ***[processpath]***	Execute a process on the remote host and return stdout/stderr. For example, this command prints the output of ipconfig to the screen: `execfg c:\windows\system32\ipconfig /all`
files ***[-s session] [action] [option]***	Perform actions over cache-stored session files. All files transferred to/from an endpoint with every CB Live Response session are cached on the server for a period of time after a session is closed. If there is any kind of interruption in the connection between a user's browser and the CB Response server, files may be retrieved directly from the cache instead of connecting to the sensor again. This command is valid in both the global and session scopes when attached to a sensor. In the global scope, the session ID must be defined with `-s`. A list of sessions is available through the `sessions` command. If attached to a sensor, the current session is assumed unless otherwise specified. There are three available actions: • `list` – List all the cached files that are available in the specified session by file ID. • `get [id]` – Get the file [id] from the cache. • `delete [id]` – Remove the file [id] from the cache.
get ***[path]***	Obtain the file specified in the path argument from remote host and download it to the host running the CB Response console for this session. Progress of the download is indicated in the Live Response window as described in "Status, Error and Progress Messages" on page 157.
help	Show the Live Response session commands with a brief description of each. If a command name is added, show the description of the specified command, with additional details (such as options) if available. For example: `help dir`
hexdump	Output the first 50 bytes of the file in a hexdump format.
kill	Terminate the specified process.

Command	Description
memdump ***[filepath]***	Take a full system memory dump and store it to the given file path, which must include a file name. When the memory dump is completed, it is automatically compressed and uploaded to the server. If you name the file with a .zip extension, it will be uploaded using the file name you provided. Otherwise, Live Response will append .zip to the name you provide. Once uploaded, the .zip file can be downloaded through the CB Response console. Memory dumps can take several minutes. Progress is indicated in the Live Response window as described in "Status, Error and Progress Messages" on page 157. The memdump command is for Windows hosts only.
mkdir	Make a directory on the remote host.
ps	Obtain a list of processes from the remote host. In the output from this command, the listing for each process includes an **Analyze** link. Clicking the link opens the **Process Analysis** page for the process. Note that analysis information for a newly discovered process might not yet be fully committed to the CB Response database and therefore not viewable. Clicking the link navigates away from the Live Response console and loses whatever context you had there.
put ***[remotepath]***	Put a file from the host on which the console is being run onto the remote host at the specified path. You specify the file in the **Open** dialog of the browser, after the command is entered in Live Response. Progress of the upload is indicated in the CBLR console as described in "Status, Error and Progress Messages" on page 157.
pwd	Print the current working directory.
reg	View or modify Windows registry settings. The syntax of this command is: **reg** *[action] [key] [options]* See "Registry Access in Live Response" on page 158 or use `help reg` in the Live Response command window for details. This is for Windows only.

As shown in the preceding table, some commands provide information and others allow you to modify an endpoint.

Note

Be sure to use the commands and options as documented here. Although some of the Live Response commands are the same as commands in the DOS command interface, the available options are specific to Live Response.

Status, Error and Progress Messages

Status and error messages should inform you of any connection or command error issues, but you can also use the `dir` or `pwd` commands to confirm your connection.

For commands that involve file transfers (get, put, and memdump), Live Response reports on the progress of the transfer. As soon as you begin a session, Live Response monitors for file transfer activity.

When one of the file transfer commands is executed, a rotating asterisk character appears in the Live Response window indicating that the transfer is happening, and a series of progress messages appear.

The messages depend upon the size of the transfer and whether you initiated the transfer in the current session or have attached to an existing session:

- For file transfers that take more than a few seconds, the Live Response window shows the number of bytes and the percentage of the total already transferred, updated every five seconds. This information is displayed to both the user who initiated the command and to any user who is attached to the same session.

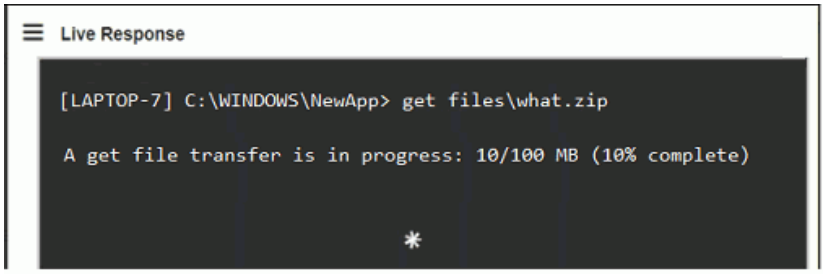

- There are other status messages, such as “Preparing file for transfer.” If a transfer is small and completed almost immediately, this might be all you see, without byte or percentage numbers.

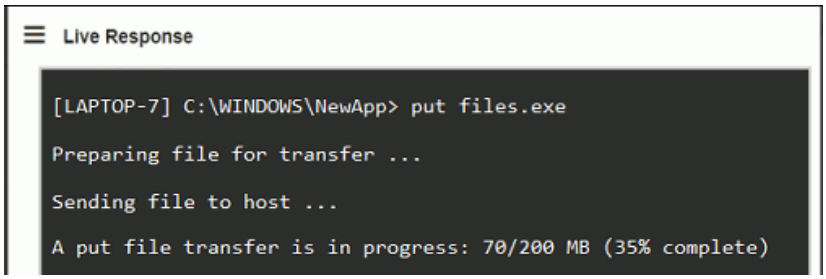

- While the transfer is in progress, users attaching to a session that has a transfer in progress will see only the numerical progress indicators.
- A successful transfer results in a “File transfer complete” message for all users who are attached to the session, and a command prompt returns.

Note

While a file transfer is in progress, no other commands can be entered in the Live Response window.

Ending Live Response Sessions

To end a Live Response session:

- In the Live Response command window, enter the `detach` command.

  The session with that computer ends and the general `[Live Response]#` prompt replaces the computer-specific prompt.

Sessions also timeout after a lack of activity. The default timeout value is five minutes. You can change this value in the `CbLRSessionTimeout` setting in the `cb.conf` file. For more information, see the *Carbon Black Response Server Configuration (cb.conf) Guide*.

Registry Access in Live Response

In a Live Response session for a Windows sensor, the `reg` command provides direct access to the remote computer's Windows Registry.

The syntax of the Live Response `reg` command is:

```
[Live Response]# reg [action] [key or value] [options]
```

The following table shows the `reg` command actions and their options. These options are intended to mirror the Windows default `reg.exe` command syntax. For all `reg` command actions, key paths can take hive references in either short or long form: `HKLM` or `HKEY_LOCAL_MACHINE`.

Action	Description
query	Format: **`reg query`** *`[key or value] [options]`* Options: *(none)* – If no option switch is specified, query for the specified key **`-v`** – Query for the specified value For example: `reg query HKLM\Software\Microsoft\Windows\CurrentVersion\Run`
add	Format: **`reg add`** *`[key] [options]`* Options: **`-v`** – Value for the key to be added **`-d`** – Data for the key to be added **`-t`** – Type of the key to be added; accepted types are: • `REG_NONE` • `REG_BINARY` • `REG_SZ` • `REG_EXPAND_SZ` • `REG_MULTI_SZ` • `REG_DWORD` • `REG_DWORD_BIG_ENDIAN` • `REG_QWORD` For example: `reg add HKLM\Software\Microsoft\Windows\CurrentVersion\Run -v calc -t REG_SZ -d c:\windows\system32\calc.exe`
delete	Format: **`reg delete`** *`[key or value] [options]`* Options: *(none)* – If no option switch is specified, delete the specified key **`-v`** – Delete the specified value For example: `reg delete HKLM\Software\Microsoft\Windows\CurrentVersion\Run -v calc`

Detached Session Management Mode

You can enter Live Response without a specific session. In this mode, you can take certain actions that do not require access to an endpoint, such as viewing the sessions that are active or examining files uploaded to the server as a result of a session. You can attach to (join) an existing session or create a new one.

Some commands in detached mode are accessible by users who do not have Global Administrator privileges, but most are not, and attempting to use them returns an error message in the command window.

To open a Live Response command window without a session:

- On the navigation bar, click **Go Live**.

  The Live Response page appears. In this mode, the prompt in the command window shows **`[Live Response]#`** without the name of an endpoint.

The following table shows the available commands in Live Response Management Mode.

Command	Description
archive ***[id]***	Obtain an archive (gzip tarball) of all the session data for the session whose ID is provided.
argparse	Test how Live Response parses CLI arguments. This command helps determine whether there are any interpretation issues.
attach ***[id]***	Attach to the session whose ID is provided. The `session` command can be used to find the ID of an existing session or create a new one. A session must be in **active** or **pending** state to be attached.
clear	Clear the console screen. You can also use the `cls` command for this purpose.
files -s ***[id]***	Perform actions over cache-stored files for the session whose ID is provided.
help	Show the commands available in this mode with a brief description of each.
help ***command***	Show the description of the specified command with additional details (such as options) if available. **For example:** `help dir`
sensor ***[options]***	List sensors managed by this CB Response server. Options: **`-i`** *`[1.2.3.4]`* – Return all sensors with specified IP address **`-n`** *`[host_str]`* – Return all sensors with matching host name **`-a`** – Return all sensors Searches are case-sensitive substring searches for both host name and IP address. You must use an option with this command. If both **`-n`** and **`-i`** are specified, only **`-i`** is used.

Command	Description
session	Manage Live Response sessions. With no argument, lists all open sessions and their ID numbers, which can be used with the `attach` command. Options: • **`session new`** *`[id]`* – Create a new session for the sensor whose ID number is provided. You must provide a *sensor* ID, not a *session* ID. • **`session list [-v]`** – List existing sessions. If the **`-v`** option is included, closed sessions are included. This option (without **`-v`**) is the default when no additional arguments are used. • **`close`** *`[id]`* – Close the session whose ID is provided.

Extending Live Response

Because the built-in commands in Live Response include `put` to put a file on the remote system and `exec` and `execfg` to execute processes on the system, responders can arbitrarily extend the capabilities of Live Response beyond the built-ins commands.

For example, an investigator could take the following series of actions:

- Upload `yara.exe` and search memory for your custom yara signatures.
- Upload `winpmem.exe` and dump a memory image.
- Upload `sbag.exe` and parse the registry for Shellbags artifacts.
- Upload a custom PowerShell script and execute it with `powershell.exe`.

Although the library of built-in commands in Live Response will grow, it will never include every command for every situation. The ability to use `put file` and `create process` together assures that you have the freedom to add utilities you need for forensics and incident response. Additional capabilities are provided by a Live Response API, described at:

https://github.com/carbonblack/cbapi/tree/5.0.0/sensor_apis#carbon-black-live-response-sensor-api

Live Response Activity Logging and Downloads

Live Response activity is logged on both the CB Response server running Live Response and the sensors it accesses.

For any sensor that is accessed by Live Response, commands executed during the session are logged in the `sensor.log` file, which is in the CB Response sensor installation folder on the endpoint.

On the CB Response server, Live Response activity can be reviewed in the following files:

- `/var/log/cb/liveresponse/debug.log` – Begin troubleshooting a Live Response issue here. This log contains debug information that is related to the functional operation of the Live Response components and communication between sensor and server.

- `/etc/cb/liveresponse-logger.conf` – You can change the level of information in the `debug.log`.
- `/var/log/cb/audit/live-response.log` – This file audits Live Response activity. It keeps a log of all commands that are executed on an endpoint, the sensor ID, IP address, and hostname of the endpoint, and the username and account of the user who executed each command.
- `/var/cb/data/liveresponse` – This directory stores “get” and “put” files. It also contains the output of all executed commands. For example, if you perform a process listing, the list goes into this directory in JSON format. If you download a file (for example, using the archive command), it appears in this directory (under `/tmp`) and on the host that is running the CB Response browser.

You can change the length of time Live Response data is retained by editing the `CbLrDefaultSessionTTLDays` parameter in the `cb.conf` file. By default this is 7 days. For more information, see the *Carbon Black Response Server Configuration (cb.conf) Guide*.

Banning Process Hashes

A CB Response investigation might reveal that known malware has been allowed to run on endpoints without being blocked. This could be because of a gap in updating your endpoint protection software or a more general gap in protection capabilities. Another possibility is that you receive notification of a threat not yet encountered on your endpoints, and you are not certain that you are fully protected against it.

While not intended to replace endpoint protection products, CB Response provides a hash banning feature to prevent malware processes from running in the future. This feature will also terminate the process for a newly banned hash if it is running when the ban is created. You can use this feature to prevent further actions from a threat until your endpoint protection is able to do so.

In order to ban hashes, you must be one of the following:

- a user that has the enhanced Analyst permission for banning hashes and is a member of a team that has the Analyst role for any sensor group
- for on-premise installations, a Global Administrator
- for cloud installations, an Administrator

Notes

The CB Response banning feature identifies and bans processes based on their MD5 hash.

- Hash banning does not ban shared libraries, such as DLLs, SYSs, CPLs, and OCXs. You can follow the steps to ban these files, but it will have no effect.
- Banning does not use SHA-256 hashes, even if they are available.
- If an endpoint is restarted, any banned process that runs on restart will terminate as soon as the CB Response sensor begins to run.

Creating Process Hash Bans

You can ban a process MD5 hash from several locations in the CB Response console:

- The **Binary Details** page has a **Ban this hash** button if the binary is an EXE.
- The **Process Analysis** page has a **Ban this hash** command on the **Actions** menu.
- The **Manage Banned Hashes** page includes the **Ban More Hashes** button. You can click this button to specify one or more MD5 hashes to be banned.
- The **Manage Banned Hashes** page has checkboxes that allow previously configured bans to be disabled and re-enabled.

> **Note**
>
> Banning a process hash without knowing the purpose of the process can have serious consequences. While CB Response sensors prevent you from banning most critical processes, it is still be possible for a user to ban a process that is required for proper operation of either your computers or your business applications. Be sure all CB Response console users understand this before they use the banning feature.

The following procedure describes how to ban an MD5 hash.

To ban a process MD5 hash from the Process Analysis page:

1. Log in as a Global Administrator (on premises) or Administrator (cloud), or as a user who has Ban hashes permission and membership on a team with the Analyst role for any sensor group.
2. On the navigation bar, click **Process Search** and search for the process. See “Process Search and Analysis” on page 172 for details on searching.
3. Click the name of the process to ban.
4. On the Process Analysis page, click **Ban this hash**.

> **Note**
>
> This button only appears if the binary is an EXE. DLLs cannot be banned.

5. The **Confirm Banned Hashes** page appears and lists this hash, indicates whether it is known, and the number of endpoints at this site on which the hash for this process has been seen.

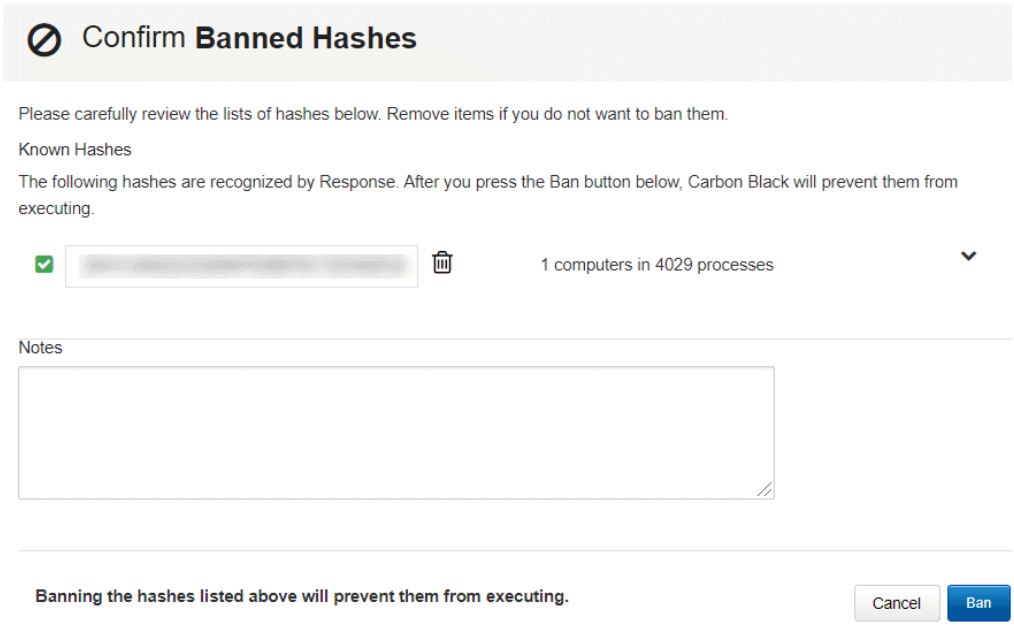

6. Clicking the **Trashcan** icon deletes the hash from the list. For single-hash-ban operations, click **Cancel** at the bottom of the page.
7. Add information in the **Notes** textbox, to explain why you banned the hash. This can include a file name, threat report identification, or anything else that is helpful to examine the ban.
8. Click **Ban** to ban the hash. The ban is added to the list on the Manage Banned Hashes page, and is enabled. By default, the list is arranged in alphanumeric order by MD5 hash.

Banning a List of Hashes

You might have a list of process hashes from CB Response or another source that you want to ban. For example, a warning from a threat intelligence source might provide a list of malware hashes. You can ban these processes in bulk on the Manage Banned Hashes page, including processes that are not yet observed by sensors reporting to your CB Response server.

To ban a list of process hashes:

1. Log in as a Global Administrator (on premises) or Administrator (cloud), or as a user with Ban hashes permission and membership on a team with the Analyst role for any Sensor Group.
2. On the navigation bar, click **Banned Hashes**.
3. Click the **Ban More Hashes** button.

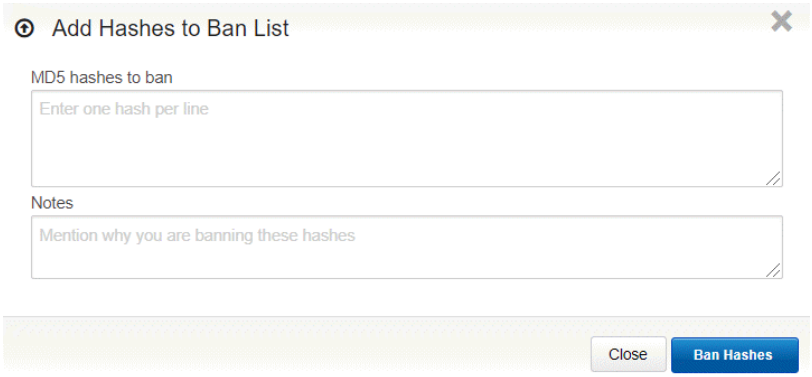

4. In the **MD5 hashes to ban** field, enter the MD5 hashes for the processes to ban. Each hash must be on its own line.
5. In the **Notes** field, provide information about why these hashes are being banned. You might also want to add names for each of the hashes, if available.
6. When you have entered the hashes and notes, click **Ban Hashes** to display the Confirm Banned Hashes page.

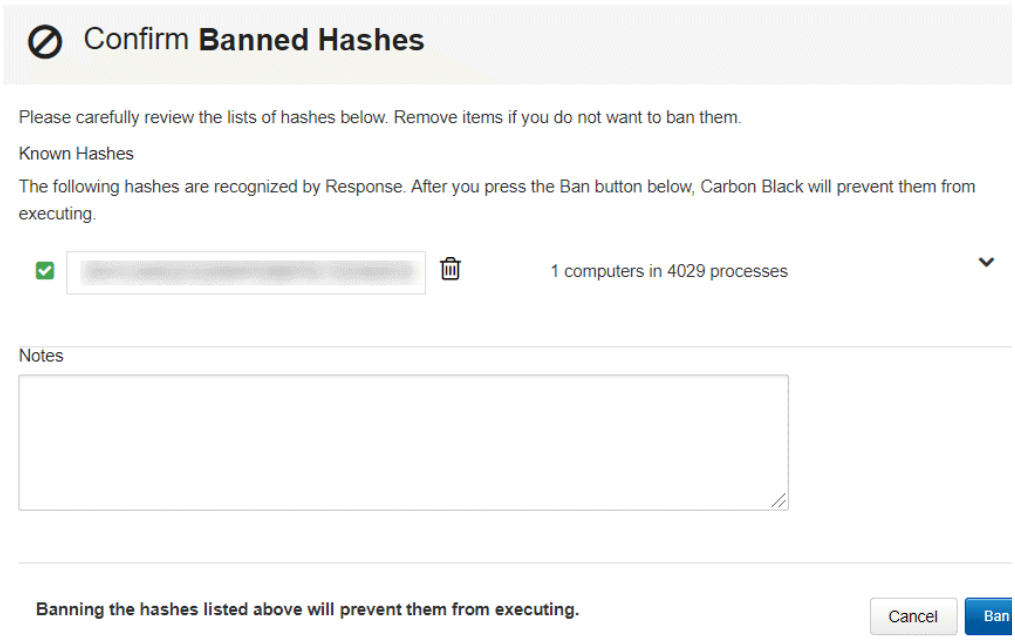

Note

The page indicates whether the hash is already known to this CB Response server, and if so, how many instances of the process have been seen and on how many computers. This page also allows you to modify the **Notes** for this ban before finalizing it.

7. For more information about a known hash, click the down-arrow to the right of it.
8. If you decide not to ban a hash, click the **Trash can** icon next to it.
9. Click **Ban** to ban all listed hashes.

The bans are added to the list on the **Manage Banned Hashes** page and are enabled.

The **Notes** you entered appear next to each hash you included in this ban. By default, the list is arranged in alphanumeric order by MD5 hash.

Managing and Monitoring Hash Bans

After you begin banning process hashes, several options are available for managing and making use of bans. You can:

- View data about bans on the Manage Banned Hashes page. See “The Manage Banned Hashes Page” on page 166.
- View block events on the Process Analysis page. See “To view all block events for a parent process:” on page 168.
- Enable alerts and syslog event recording for process blocks caused by bans, using the **Banning Events** feed on the Threat Intelligence Feeds page. See “Enabling Alerts and Syslog Output for Banning Events” on page 170.
- Monitor banning alerts on the Triage Alerts page. See “To view banned hash alerts:” on page 170.
- Enable and disable bans on the Manage Banned Hashes page. See “Disabling a Hash Ban” on page 171.

The Manage Banned Hashes Page

The Manage Banned Hashes page lets you add, manage, and get information about process hash bans created on your CB Response server:

- **Table of Bans** – Any hash bans that have been created on your CB Response server are listed in a table, including bans that are enabled and bans that are not currently enabled. There is also an indicator at the top-right corner of the page that shows the total number of bans (both enabled and disabled) that have been created.
- **Access to Additional Ban Information** – Some information about each ban is shown in the table rows, and additional information is available through drill-down features for each ban.
- **Toggling of Ban Status** – The status of each ban is displayed in the **Banned** column. You can enable or disable any ban.
- **Ban More Hashes** – This button opens the **Add Hashes to Ban List** dialog, where you can enter one or more hash values on the page to create new bans.

The table of hashes lists each hash that has been created on this server. You can also search for hash bans by the MD5 hash of the process, and you can control the display of the entire table using the following controls:

- **View** – You can click different buttons in the **View** field to display **All** bans (the default), currently **Banned** hashes, and **Previously Banned** hashes (ban disabled).
- **Sort By** – You can sort the table by **MD5** hash (the default sort), **Date Added**, or **User**. Radio buttons change the sort order from ascending to descending.

The following table describes fields on this page. Note that the table data that reports on blocks caused by bans requires that the Banning Events feed on the Threat Intelligence Feeds page is enabled. (See "Threat Intelligence Feeds" on page 244.)

Column	Description
Hash	The MD5 hash of the process that is or was banned. Clicking on the hash opens the Binary Details page for the hash.
Notes	Any user-created notes about the ban or hash.
Latest Block	The length of time since the process identified by the MD5 hash was blocked on a system reporting to the CB Response server.
Total Blocks	The total number of times this process has been blocked by the ban.
Hosts w/ Blocks	The number of systems on which this process was blocked at least once. If a host name appears, clicking on it opens the Sensor Details page for that host.
Banned	This checkbox controls the status of the ban. When the box is checked, the ban is enabled. When the checkbox is not checked, the ban is disabled.
(more details)	Clicking the blue down arrow icon expands the row for a hash ban to provide additional details.

When you expand the row for a ban using the blue down arrow, information about the ban and its process appears in the panel. You also can use navigation links to go to other pages for more information.

The table data requires that the Banning Events feed on the Threat Intelligence Feeds page is enabled. (See "Threat Intelligence Feeds" on page 244.)

The following table describes the process hash ban details:

Column	Description
Hosts / Processes	Shows how many hosts have run the process identified by this MD5 hash and how many times the process ran before it was banned.
Meta data	The name of the CB Response console user who created the ban, when the ban was added, and the date and time of the most recent block caused by the ban. Clicking the user name navigates to the table of users on the User Management page.
Hosts	The endpoints on which the process controlled by this ban has been blocked.
Notes	Any user-created notes about the ban or the hash. Notes can be edited.
View ban history	Opens a separate **Ban History** window that shows status changes (enabled, disabled) for the ban, who made them, and when they were made.

Column	Description
(process search)	Clicking the blue magnifying glass icon navigates to the Process Search page with the results of a search for this process.

Monitoring Banning Events

When a process is blocked because of a CB Response hash ban, that is an indication that some user or process attempted an unwanted activity. Even though the activity was blocked, you might want to investigate the attempt.

CB Response reports an event each time a hash ban attempts to block a process, even if the block fails (for example because of an attempt to block a critical system or CB Response process). The event appears on the Process Analysis page of the parent process. If a process was running at the time a ban was created and then terminated by the ban, a banner reports that fact on the Process Analysis page.

Blocking events can also trigger alerts and be included in the syslog output from CB Response. See “Enabling Alerts and Syslog Output for Banning Events” on page 170.

To view all block events for a parent process:

- On the Process Analysis page for the parent process, search for **blocked** in the **Type** filter. (See “Process Search and Analysis” on page 172.)

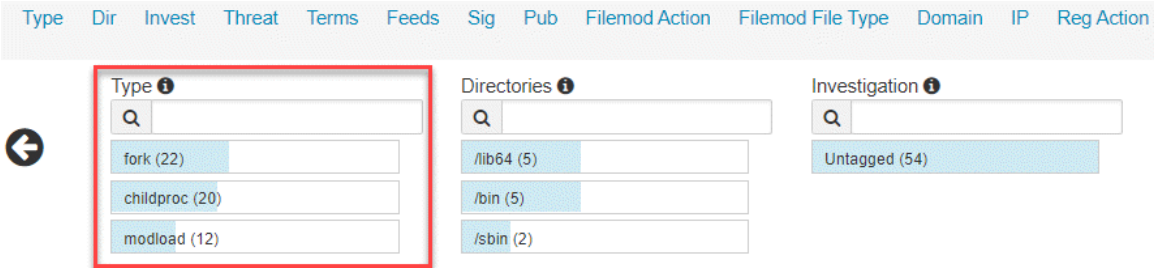

Searching for Blocked Processes

The Process Search page lets you search for processes that have been affected by a process hash ban. This includes processes that were successfully blocked as well as those that could not be blocked for various reasons.

To search for processes that have block events:

1. On the navigation bar, click **Process Search**.
2. Click **Add Search Terms** and then click **Choose Criteria**.

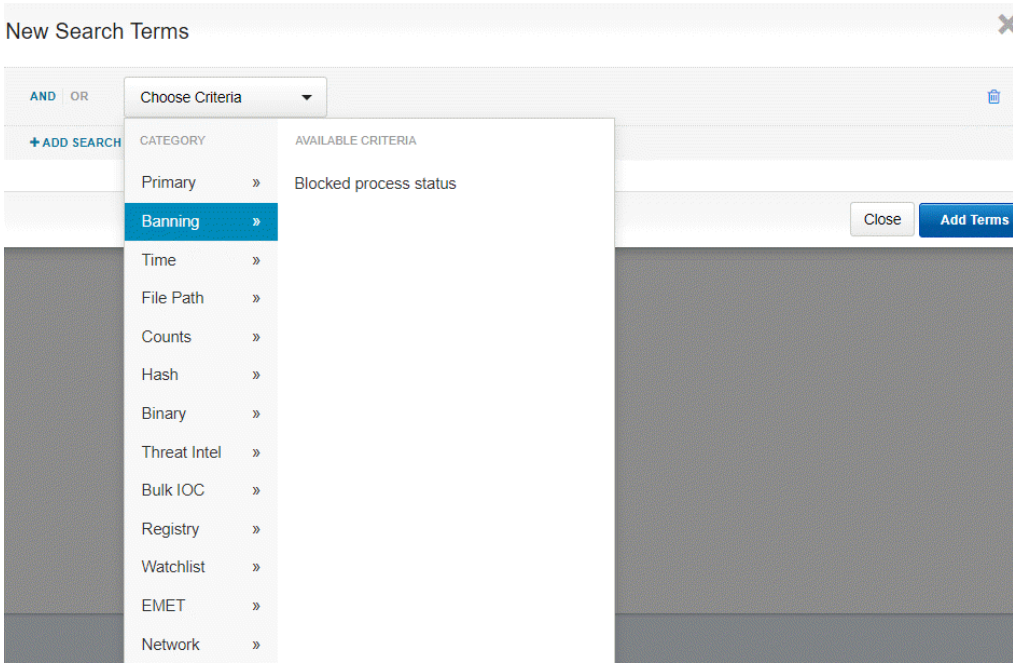

3. Click **Banning** and then click **Blocked process status**.
4. You can select block conditions to search for.

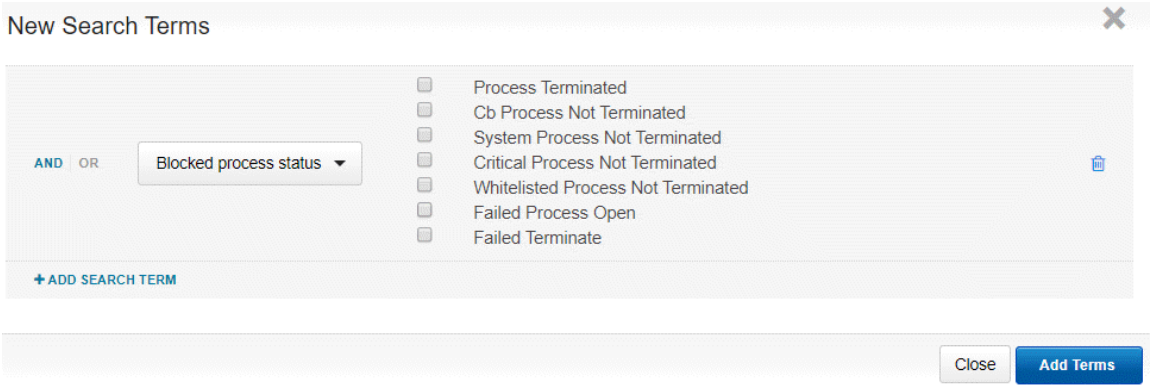

5. To search for successful blocks only, select **Process Terminated**. To search for other block events, check the boxes for those events. When you have selected all relevant boxes, click **Add Terms** and run the search. The search results are updated to match your criteria .

Note

You can manually enter one of the following queries for blocked processes:

- `block_status:processterminated`
- `processblock_count:[1 TO *]`

Enabling Alerts and Syslog Output for Banning Events

Unless banning is disabled entirely, process block events are sent to the CB Response server and viewable on the Process Analysis page. (See “Process Search and Analysis” on page 172.)

To configure alerts and syslog output for process blocks, a special **Banning Events** panel is available on the Threat Intelligence Feeds page. (See “Threat Intelligence Feeds” on page 244.) This is not a feed in the normal sense, because the events for blocks are sent to the server regardless of whether the feed is enabled. However, the feed must be enabled if you want to configure notifications for banning events.

The **Banning Events** feed is available by default and does not require enabling communication with Carbon Black Threat Intel.

To enable alerts and syslog recording of blocking events due to hash bans:

1. On the navigation bar, click **Threat Intelligence**.
2. Locate the **Banning Events** feed.

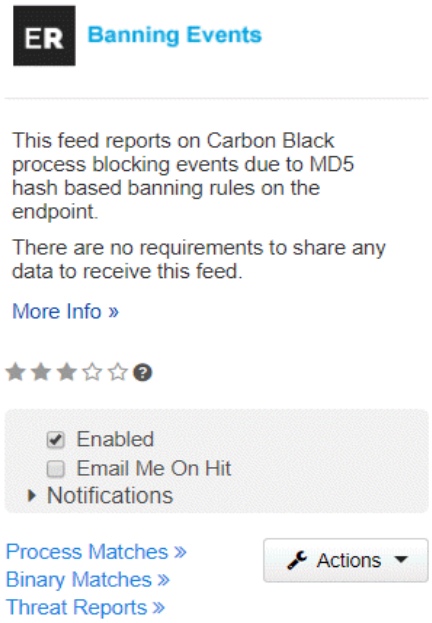

3. Click **Notifications** and select the notification types to create: **Create Alert** and/or **Log to Syslog**.
4. To receive email when a block event occurs, select **Email Me On Hit** .

   You will now receive alerts when there is an attempt to run a banned process, and the Manage Banned Hashes page will report on the number of blocks and the time of the most recent attempt and the system where it occurred.

To view banned hash alerts:

1. On the navigation bar, click **Triage Alerts**. For more information, see “Managing Alerts on the Triage Alerts Page” on page 289.
2. In the search box for the **Feed** filter, enter `cbbanning` and press **`Enter`**, or if it already appears on the list, click `cbbanning`.

   Banning Events alerts appear in the results table.

In addition to triggering alerts (if enabled), processes that are blocked due to a hash ban generate events that appear on the Process Analysis page.

For example, if you receive a Process Blocking alert for a process, the Process Analysis page for the parent process appears and includes a **blocked** event.

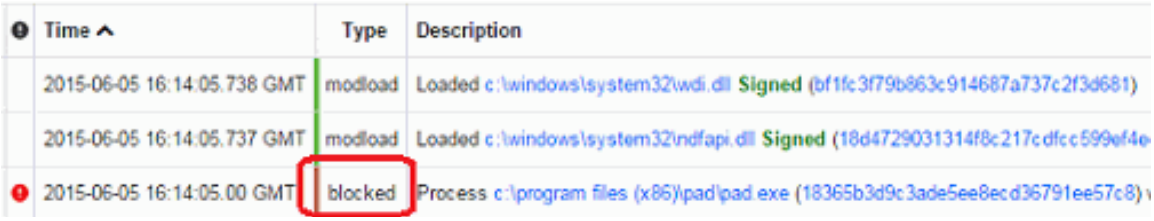

Disabling a Hash Ban

After a hash ban is created, it always appears on the Manage Banned Hashes page. You can turn bans on and off.

To disable a process hash ban:

1. Log in as a Global Administrator (on premises) or Administrator (cloud), or as a user who has Ban hashes permission and membership on a team with the Analyst role for any sensor group.
2. On the navigation bar, click **Banned Hashes**.
3. Locate the row for the hash ban to remove, and deselect the check box in the **Banned** column.

Disabling or Restricting the Hash Ban Feature

The ability to ban process hashes is enabled by default. However, you can disable or restrict it in certain ways.

Disabling Bans in a Sensor Group

You can disable banning on all hosts in a specified sensor group. In this case, any process hash bans that are configured on the server are ignored by sensors in that group, and no processes are blocked by CB Response on those sensors.

To disable process hash bans in a sensor group:

1. On the navigation bar, click **Sensors**.
2. In the **Groups** panel, click the sensor group to exempt from hash bans.
3. At the top of the **<*name*> Group** panel, click **Edit**, and then click **Advanced** to expand the advanced sensor group settings.
4. Deselect **Process Banning**.
5. Click **Save Group**.

Chapter 11

Process Search and Analysis

This chapter describes how to perform detailed process searches and in-depth analyses of processes that are found in search results.

Sections

Topic	Page
Overview of Process Search	173
Results Table	182
Process Analysis Page	184
Analysis Preview Page	199

Overview of Process Search

When you become aware of an incident that could be a threat, you can search all your systems and endpoints for processes that have Indicators of Compromise (IOCs). For example, you might receive a call reporting unusual software behavior or an alert from a threat intelligence report or watchlist. CB Response sensors collect data automatically so that you can immediately start analyzing issues and finding solutions.

Use the Process Search page to begin investigating potential threats. This section describes how to perform basic process searches using search strings and predefined search criteria.

> **Note**
>
> Some process data might not be displayed depending upon certain `cb.conf` settings. The **SensorLookupInactiveFilterDays** setting determines whether sensors that have not checked in for a specified number of days are filtered out of the console Sensors page. If the value of this setting is less than the value of **MaxEventStoreDays** (30 days by default), process data from inactive sensors is not included in search results.

To Access the Process Search Page:

- On the navigation bar, click **Process Search**.

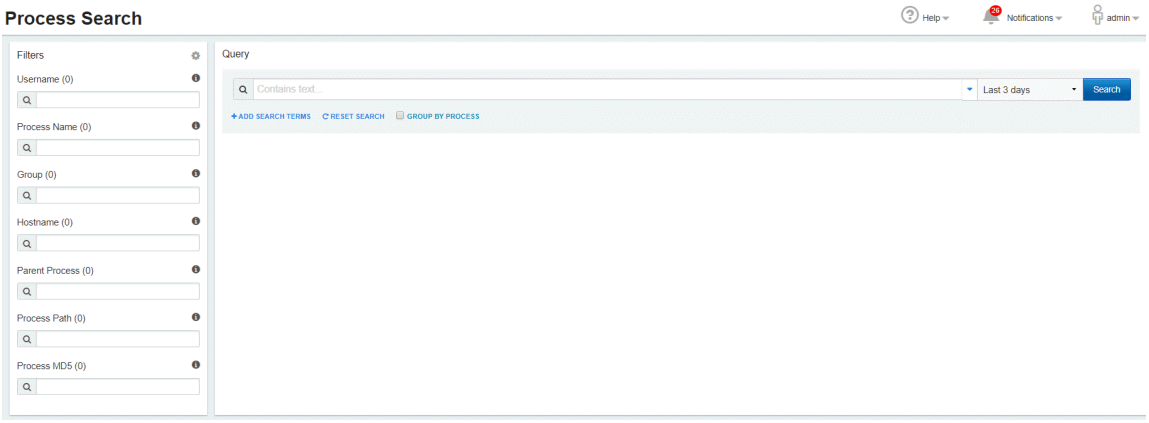

Time Filter

In the **Process Search** page, you can specify that the results show only processes that appear in events that occurred within a specified time period.

To use the time filter:

1. Select a time filter from the dropdown menu.

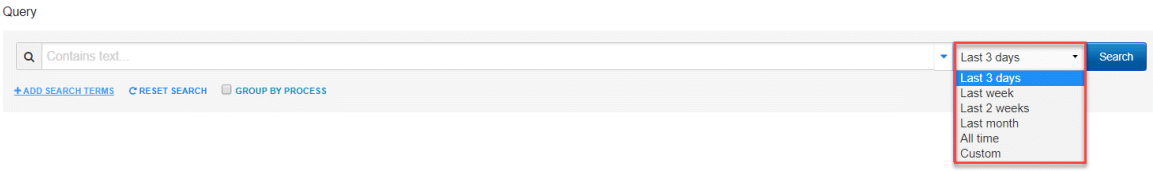

2. Click **Search**.

Search Filters

Search filters provide ways to specify and narrow a search. Each filter represents terms that have actually been seen in various fields, such as **Process Name** or **Hostname**. The percentage next to each term shows the relative frequency with which the term appears in the field.

No content appears in the search filters until after you have initiated a search. Then, the search filters populate according to their match to the search results.

Available filters include:

- **Process Name** – Unique names of processes that match your search criteria.
- **Group** – The activity distributed among the configured sensor groups whose processes match the search criteria.
- **Hostname** – The hostnames of the currently installed sensors that have processes that match the search criteria.
- **Parent Process** – The parent processes that create child processes and match the search criteria.
- **Process Path**– The full physical path of the executables from which a process was executed.
- **Process MD5** – The MD5 hash value of the executable for each matching process.

Enable/Disable Filters

To display only certain search filters on the Process Search page:

1. Click the **Gear** icon to the right of **Filters**.
2. Select checkboxes to enable or disable the filters to display.

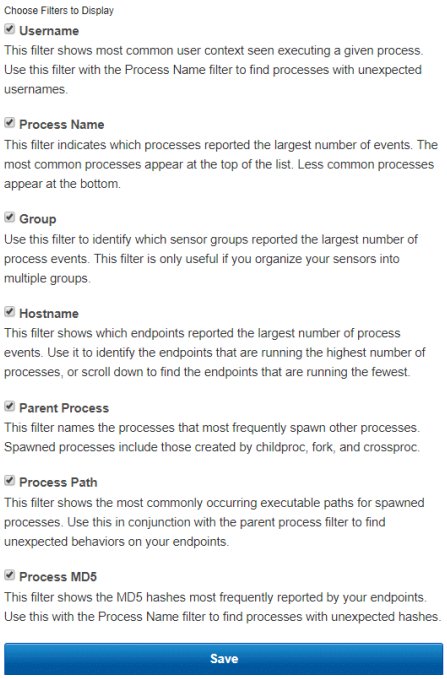

Disabling a filter removes it from view, and if it is part of the search query, those pieces of the query are removed. Enabling a filter places it back into view.

3. Click **Save**.

Select Multiple Filter Rows

You can select specific filter rows within a filter table by using your cursor. The search results are updated based on these selections.

- Selecting multiple rows within a single filter updates the query with a logical OR between those filters. For example, choosing “bash” and “nginx” in the **Process Name** filter shows events related to either bash or nginx.
- Selecting multiple rows across multiple filters updates the query with a logical AND between those filters. For example, choosing “bash” in the **Process Name** filter and “python” in the **Parent Process** filter shows instances of bash that were spawned by Python.

Selected filter rows are highlighted in yellow. Click a filter row to deselect it.

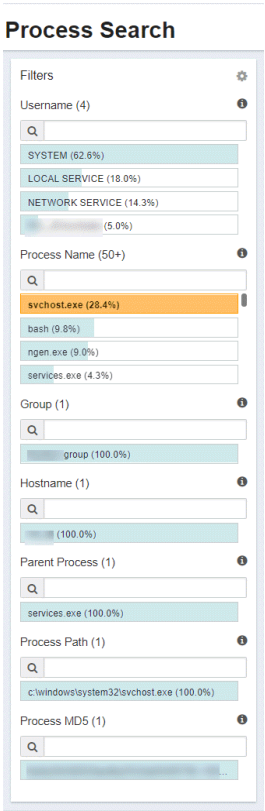

Filter Row Percentages

The top row within a filter has occurred more than any other process within that filter. Filter row percentages indicate the percentage of processes that have occurred in a particular filter. This is always equivalent to 100% when you add up all filter rows in a filter.

Filter Search Fields

Each filter contains a **Search** field in which you can enter search parameters to further refine search results.

Search Field

You can manually enter keyword searches or predefined search criteria in the **Search** field at the top of the Process Search page. While you enter search criteria, the correct syntax appears. If you do not enter any search criteria, the system runs a search using ***.*** This displays every process that has executed. The processes are ranked according to the process start time — the most frequently executed processes are at the top. You can sort the results according to the count of events or last update time.

The **Search** field and criteria fields can be used independently or in combination. When used in combination, the system combines them using an `AND` operator.

Clicking the blue **Search** button executes a search using selected parameters. By default, searches are constrained to the past three days. You can select a different time range for your search. When viewing events on the Process Search page, the time displayed next to the process is the time that the sensor recorded the event, not the time that the server received the event.

Save Searches

You can save frequently executed searches by selecting the **Favorite** (star) icon to the right of the **Search** field. A confirmation appears in the top-right corner of the console indicating that the search has been saved.

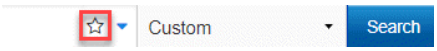

To execute a saved search:

- Click the down arrow to the right of the **Favorite** (star) icon and select the saved search from the dropdown list.

  The selected saved search is loaded and executed:

Clear Saved Searches

You can clear saved searches by accessing the **My Profile** page and clicking **Clear Preferences**.

To clear saved searches:

1. Click ***Username*** **> My Profile**.
2. Click **Clear Preferences**.

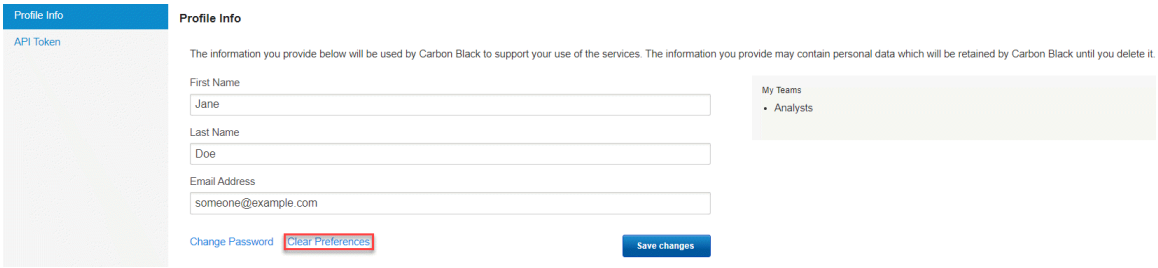

Add Search Terms

Process searches explicitly support `AND/OR` operators. You can select from an array of filters to form your search using these `AND/OR` operators.

To add search terms:

1. Add search terms (in the form of `AND/OR` operators) by clicking **Add Search Terms** on the **Process Search** page beneath the **Search** field.

2. Select a search term type from the **Choose Criteria** dropdown menu.

3. Click **Add Search Term** to add terms. When you are finished, click **Add Terms**.

   For example, the following search terms will display processes that have started **i**n the last 30 minutes in the example.com domain.

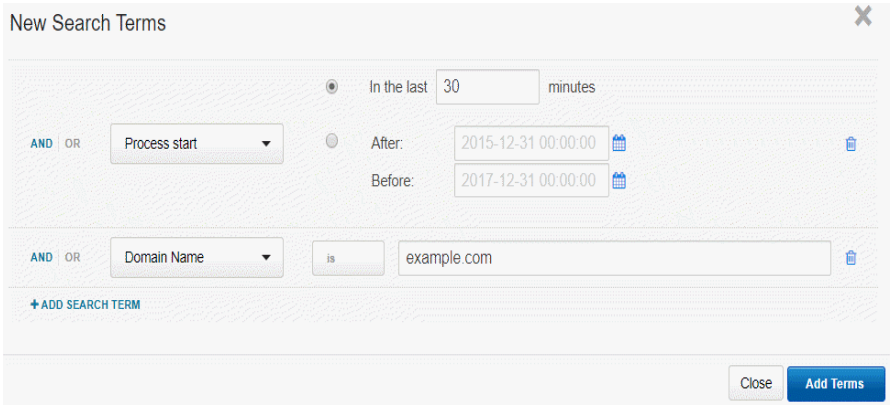

   To remove a search term, click the **Delete** (trashcan) icon to the right of the search term.

Reset Search

To reset and remove all search terms, click **Reset Search** on the **Process Search** page beneath the **Search** field.

Group By Process

To combine results with the same process into groups, select **Group By Process** on the Process Search page:

Search Result Messages

If no search results match your search criteria, the Process Search page displays a message that recommends that you widen your search by deselecting filters.

If search results are too large, you can use filters to narrow your search or you can view all results.

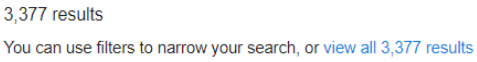

Get Comprehensive Results

A **Get Comprehensive Results** button appears on the Process Search page if a search query spans both current data and older data that was collected prior to version 6.1, and if the query has complex search terms that require special processing on the server.

- If you do not request comprehensive results, the server returns correct search results for old data, but results might be incomplete for data that was collected prior to 6.1.
- If you request comprehensive results, the server returns full search results for current data, but excludes data that was collected prior to 6.1.

Example Process Search

This section explains at a high level how to perform a process search.

To perform a process search:

1. On the navigation bar, click **Process Search**.
2. Enter search criteria by performing one (or combining both) of these tasks:
   a. Enter keyword searches or predefined search criteria in the **Search** field.
   b. Click **Choose Criteria** to display a list of searchable criteria. Select the search criteria, such as **Banning > Blocked Process Status > Process Terminated**. Then, click **Add Search Term** to add an additional set of search criteria, such as **Time > Process Start > In the last 90 minutes**. Repeat this process to add more search criteria.
3. When you are finished entering your search criteria, click **Search**.

The search results appears in the **Results** table. For more information, see “Results Table” on page 182

Note

For information on performing advanced queries, see “Advanced Search Queries” on page 213.

Managing High-Impact Queries

Certain process searches can cause significant performance problems in CB Response. Two types of searches that can have a negative impact are:

- Searches with leading wildcards
- Searches with binary terms (which require a join between the process and module databases) if you have very large modules cores - see “Searching with Binary Joins” on page 233.

Beginning with CB Response 6.2.3, these searches are blocked by default when executed through the console. However, there are options in both the console interface and the server configuration file (`cb.conf`) for blocking and unblocking these types of process searches.

The blocking features, both from `cb.conf` and through the console, applies only to interactive searches in the console. Searches executed via the API, existing watchlists or feeds are not impacted by these settings.

Responding to Blocked Searches

Users attempting blocked search types will see a message describing why the search was blocked. If you determine that a setting is preventing searches that you expect to succeed, you can either modify the settings or modify your search. If you unblock one of the search types, monitor the performance impact to determine whether you can operate successfully in that mode.

You can also reconfigure the search to avoid the blocked condition. See Chapter 13, “Advanced Search Queries,” for more information about creating complex process searches.

Process Search Settings in the Console

The settings for blocking searches with leading wildcards and searches with joins of large module cores are on the **Advanced Settings** tab on the Settings page. By default, both process search settings are checked, thereby blocking these searches.

You must have Global administrator privileges to change these settings. In addition, if the settings are controlled by a `cb.conf` file configuration, they will be grayed out and unavailable for editing. See “Process Search Settings in cb.conf” on page 181.

To block or allow high-impact process searches:

1. Log in to CB Response as a Global Administrator (for on-premises installations) or an Administrator (for the cloud).
2. Click ***Username*** **> Settings** and click **Advanced Settings**.
3. Under **Process Search Settings**, check (or uncheck) the box for the search type to block or unblock.
4. Click the **Save changes** button.

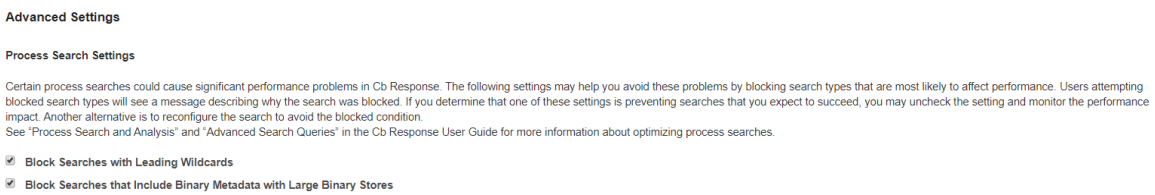

Process Search Settings in cb.conf

Two settings in the `cb.conf` file affect whether process searches with possibly significant performance impact are blocked, allowed, or configurable through the console interface:

- `ForceBlockLeadingWildcardsInSearchTo` (interacts with **Block Searches with Leading Wildcards** in the console).
- `ForceBlockCoreJoinsInSearchTo` (interacts with **Block Searches that included Binary Metadata with Large Binary Stores** in the console).

By default, neither of these settings are present in the `cb.conf` file.

If a setting is `True`, process searches in the relevant category are blocked, and no user, including a Global Admin, can change this setting through the console.

If a setting is `False`, process searches in the relevant category are allowed, and no user, including a Global Admin, can change this setting through the console.

A third setting in `cb.conf`, `ModuleCoreDocumentCountWarningThreshold`, sets the number of module core documents that is considered large enough to be blocked when **Block Searches that included Binary Metadata with Large Binary Stores** is activated. By default, this setting has a value of ten million.

Results Table

At the bottom of the Process Search page, the **Results** table appears. Each row contains details about an executed process that matches the search criteria.

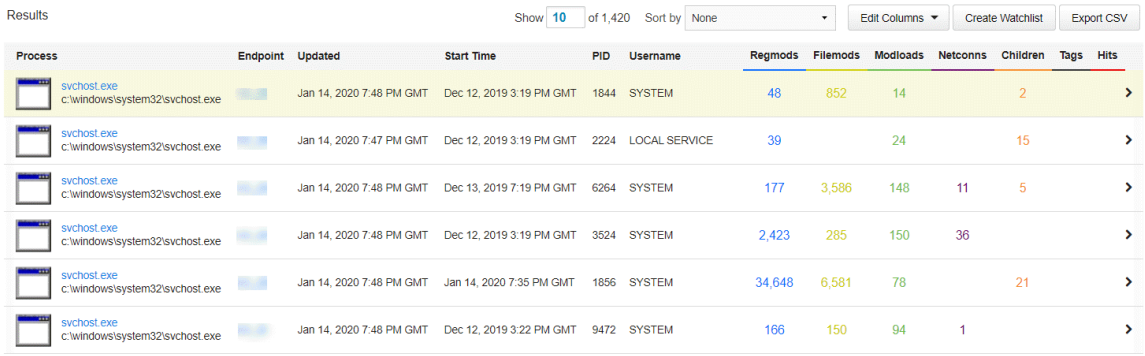

Results Table Options

A few options appear above the **Results** table:

- **Showing** – Use this option to adjust the maximum number of search results that display on a page. The default setting is 10 results per page.
- **Sort by** – Use this menu to sort search criteria by these options:
  - **None**
  - **Process last update time**
  - **Process start time**
  - **Process name**
  - **Network connections**
  - **Registry modifications**
  - **File modifications**
  - **Binary loads**
- **Edit Columns** – Use this option to select which columns are visible in the search results. Users can also toggle between showing event counts or summary information for each column:

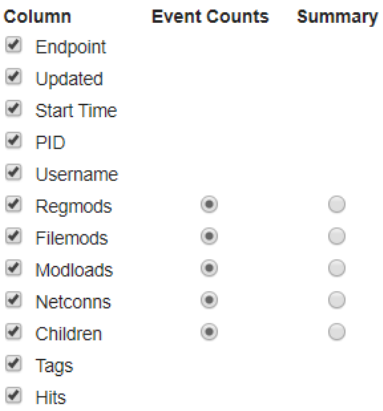

- **Create Watchlist** – Use this option to create a watchlist that is based on the current query string. A watchlist is a saved search that you can use to track specific IOCs. See Chapter 18, “Watchlists”.
- **Export CSV** – Use this option to export the first 1000 process search results to a comma-separated CSV file for reporting, retention, or compliance. Each row contains a URL to access the details of each result on the table.

> **Note**
>
> To export more than 1000 rows, you must configure API functionality to capture and save the data. See the Carbon Black Developer Network at https://developer.carbonblack.com/reference/enterprise-response/.

Results Table Row Details

On each row within the **Results** table, the following information appears:

Process Analysis Page

Title	Description
Process	The icon of the process or program that was executed and the name of the executable file that was run; for example, `notepad.exe`. The file system path from which the process was executed appears.
Endpoint	The elapsed time since the most recent execution of the process, and the endpoint that is associated with the result.
Updated	The timestamp when the process was last updated.
Start Time	The timestamp when the process started.
PID	The Process ID.
Username	The username that is associated with this process.
Regmods	The number of Windows registry modifications that were made by the execution of this process. Regmods are color-coded in blue.
Filemods	Contains a color-coded dot if the execution of the process resulted in file modifications. Filemods are color-coded in yellow.
Modloads	Contains a color-coded dot if the execution of the process resulted in loaded modules. Modloads are color-coded in green.
Netconns	Contains a color-coded dot if the execution of the process resulted in attempted or established network connections. Netconns are color-coded in purple.
Children	Contains a color-coded dot if the execution of the process resulted in generated child processes. Children are color-coded in orange.
Tags	Contains a color-coded dot if the execution of the process resulted in events that were tagged in a CB Response investigation. Tags are color-coded in black.
Hits	Contains a color-coded dot if the execution of the process resulted in watchlist or feed hits. Hits are color-coded in red.
>	The Process Analysis page with details about the process executable file. See “Process Analysis Page” on page 184.

After you have detected a threat and searched process executables, when you find a process that merits investigation, you can open the **Process Analysis** page.

To open the Process Analysis page:

1. Execute a query as discussed in “Overview of Process Search” on page 173.
2. In the **Results** table, locate the process to analyze and click the arrow (**>**).

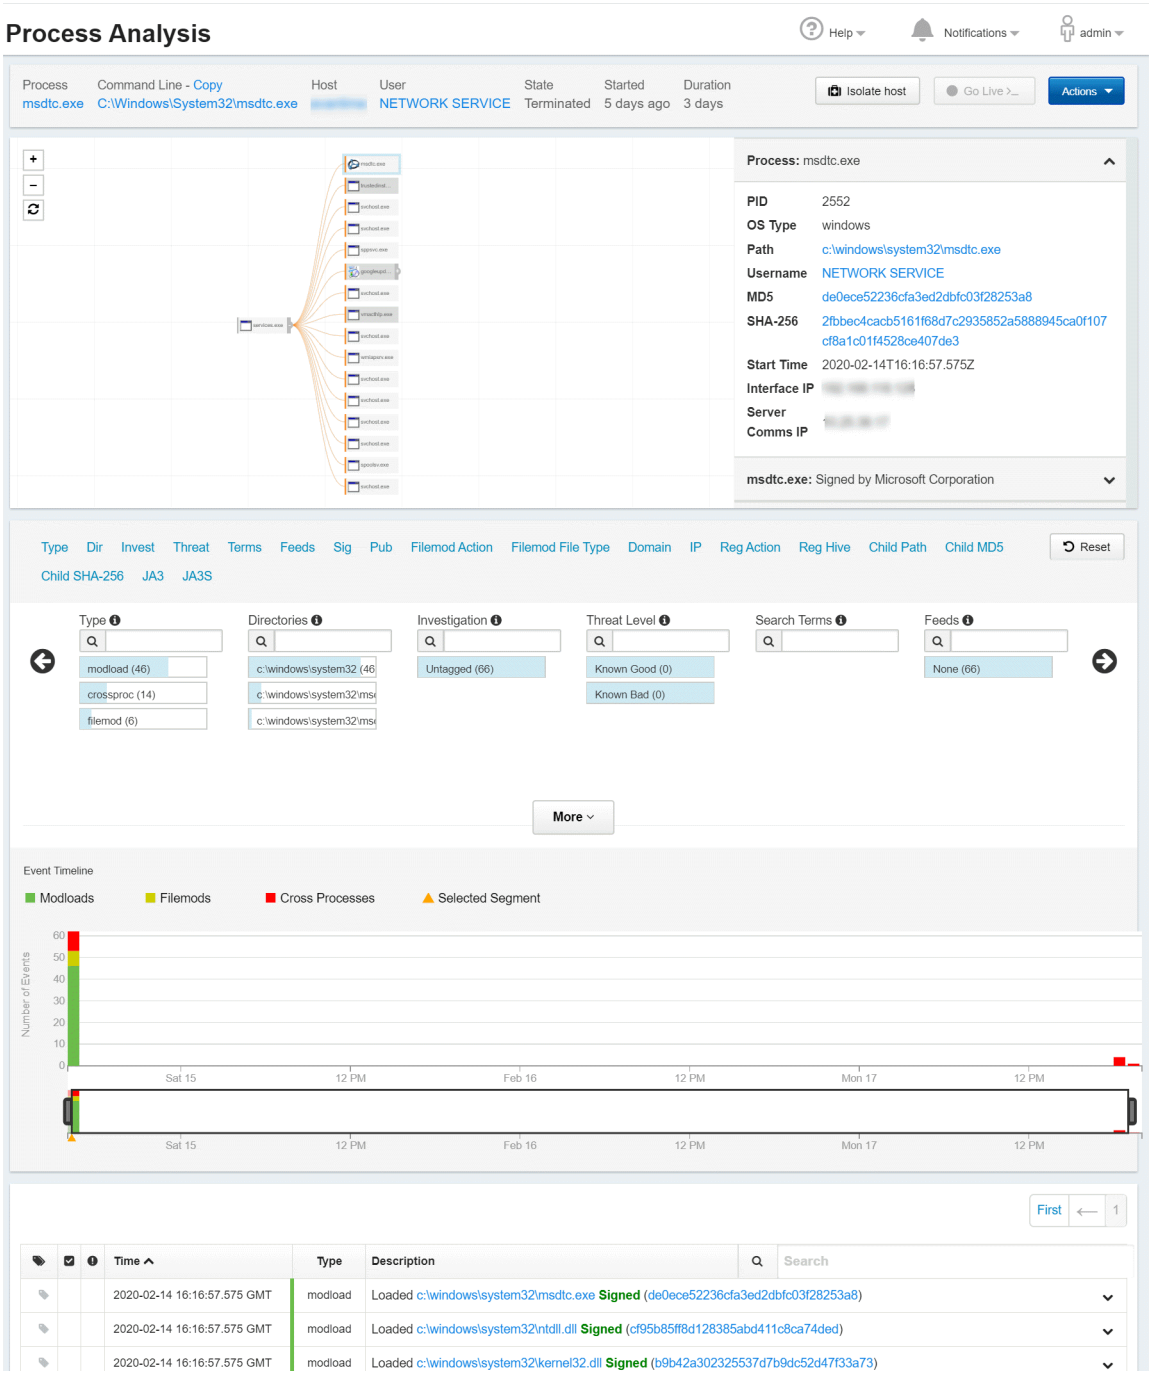
Process Analysis
Help
Notifications
admin
Process
msdtc.exe
Command Line - Copy
C:\Windows\System32\msdtc.exe
Host
User
NETWORK SERVICE
State
Terminated
Started
5 days ago
Duration
3 days
Isolate host
Go Live >_
Actions
Process: msdtc.exe
PID 2552
OS Type windows
Path c:\windows\system32\msdtc.exe
Username NETWORK SERVICE
MD5 de0ece52236cfa3ed2dbfc03f28253a8
SHA-256 2fbbec4cacb5161f68d7c2935852a5888945ca0f107cf8a1c01f4528ce407de3
Start Time 2020-02-14T16:16:57.575Z
Interface IP
Server Comms IP
msdtc.exe: Signed by Microsoft Corporation
Type Dir Invest Threat Terms Feeds Sig Pub Filemod Action Filemod File Type Domain IP Reg Action Reg Hive Child Path Child MD5 Child SHA-256 JA3 JA3S
Reset
Type
modload (46)
crossproc (14)
filemod (6)
Directories
c:\windows\system32 (46
Investigation
Untagged (66)
Threat Level
Known Good (0)
Known Bad (0)
Search Terms
Feeds
None (66)
More
Event Timeline
Modloads
Filemods
Cross Processes
Selected Segment
Number of Events
Sat 15
12 PM
Feb 16
Mon 17
First
Time
Type
Description
Search
2020-02-14 16:16:57.575 GMT modload Loaded c:\windows\system32\msdtc.exe Signed (de0ece52236cfa3ed2dbfc03f28253a8)
2020-02-14 16:16:57.575 GMT modload Loaded c:\windows\system32\ntdll.dll Signed (cf95b85ff8d128385abd411c8ca74ded)
2020-02-14 16:16:57.575 GMT modload Loaded c:\windows\system32\kernel32.dll Signed (b9b42a302325537d7b9dc52d47f33a73)

Process Analysis Features

The **Process Analysis** page contains the following features to help you investigate process details, such as:

- "Process Summary" on page 186
- "Interactive Process Tree" on page 188
- "Process Execution Details" on page 189
- "Binary Metadata" on page 190
- "On Demand Feeds shows if the process event details had any hits from on-demand feeds." on page 190
- "EMET Protections Enabled (Windows Only)" on page 191
- "Process Event Filters" on page 191
- "Event Timeline" on page 193
- "Process Event Details" on page 193

Process Summary

The process summary information is located in the top panel of the **Process Analysis** page and displays general process execution details. For example:

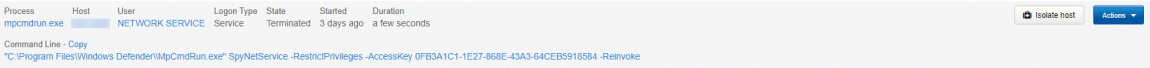

The summary information consists of the following data:

- **Process**: Identifies the main process for which the analysis is displayed.
- **Command Line**: Describes the command that initiated the process.
- **Host**: Identifies the host upon which the command was initiated.
- **User**: Identifies the user who was logged in at the endpoint when the command was initiated.
- **Logon Type**: Describes the logon type for the process; valid logon types are:
  - System (0)
  - Interactive (2)
  - Network (3)
  - Batch (4)
  - Service (5)
  - Proxy (6)
  - Unlock (7)
  - Network Cleartext (8)
  - New Credentials (9)
  - Remote Interactive (10)
  - Cached Interactive (11)
  - Cached Remote Interactive (12)
  - Cached Unlock (13)
- **State**: The state can be **Running** or **Offline**.
- **Last Activity**: The last time the process was run.

- **Duration**: The number of hours that the process has been running.
- **Isolate Host**: Click the **Isolate host** button to isolate an endpoint. For example, you might discover that suspicious files are executing from a particular endpoint and you want to prevent them from spreading to other endpoints in your network.

  When an endpoint is isolated, connections to the CB Response server (such as DHCP and DNS) are maintained, but all other connections are blocked or terminated. The user is not notified by CB Response, but the endpoint will not work as expected.

  > **Note**
  >
  > To isolate an endpoint, you must be a Global Administrator for on-premises installations, an Administrator for cloud installations, or a user on a team who has Analyst privileges for the endpoint to isolate.

  The computer remains isolated until this option is disabled or the computer reboots. See “Isolating an Endpoint” on page 147.
- **Go Live**: The **Go Live** button is useful when you are investigating an IOC. After you have identified an endpoint that has suspicious activity, you can directly access the content on that endpoint. You can open an interactive live session to the endpoint host and execute commands in real time to help isolate or eradicate the threat. See “Using Live Response” on page 150.
- **Actions**: The **Actions** dropdown menu includes the following options:
  - **Ban this hash** – Creates a ban of the process. If process hash banning is enabled for a sensor group, hosts attempting to run this process will find it blocked, and any running instances of the process are terminated. See “Banning Process Hashes” on page 162.
  - **Export events to CSV** – Downloads a `Report.zip` archive to your local computer. The files contain the information in the **Description** fields for each **Type** filter that appears in the results table at the bottom of the **Process Analysis** window. See “Process Event Filters” on page 191.
  - **Share** – Opens the CB Response user’s default email client, creates an email, and includes the details from the `summary.txt` file (path, MD5, start timestamp,

last updated timestamp, hostname, and full command line), and a URL that accesses the same page in which **Share** was clicked.

Interactive Process Tree

By default, the interactive process tree view displays the parent process of the selected process executable file in a search result together with the relevant child processes:

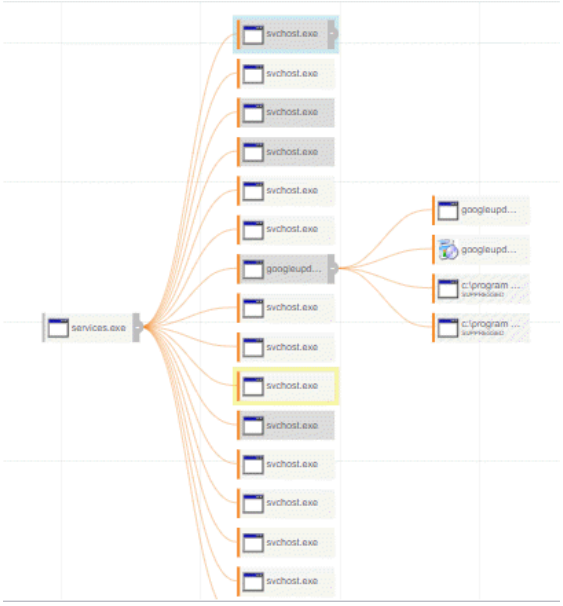

You can interact with the process tree by clicking child and parent processes to identify issues. This view shows the selected process event and includes its parent process and child processes. Siblings to the selected process also appear.

To expand or collapse nodes in the process tree, click a parent or child node.

To view additional nodes, left-click and hold any part of the tree while moving your cursor

Clicking other child or parent processes updates the Process Analysis page in context to show the newly selected process details, including the summary tables and graphs.

Note

The process tree can display up to 15 child processes; either 15 unsuppressed, 15 suppressed, or 15 of both types.

For processes that have more than 15 unsuppressed and 15 suppressed child processes, the tree shows unsuppressed processes first, and then suppressed processes, until a total of 15 child processes appear.

Process Execution Details

Process execution details appear in the panel to the right of the Process Tree:

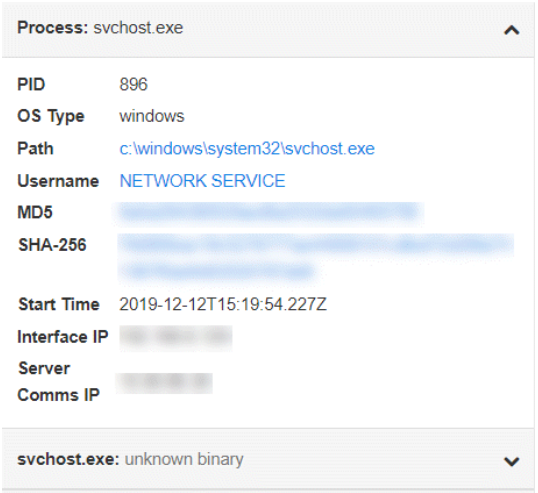

If the process is an executable, the following information is displayed:

Field	Description
Process	The name of the process executable file.
PID	The Process Identification (PID) number of the process.
OS Type	The operating system on which the process was executed.
Path	The physical path from which the process was executed.
Username	The name of the user who executed the process.
MD5	The MD5 hash value of the process.
SHA-256	The SHA-256 hash value of the process. **Note:** Availability of SHA-256 hash data is dependent upon sensor capabilities. The macOS (OS X) sensor version 6.2.4, which is packaged with CB Response Server version 6.3, sends SHA-256 hashes to the server. Check the Carbon Black User Exchange or Carbon Black Support for information about other sensors that can generate SHA-256 hashes. For files that were originally discovered by a sensor that did not provide SHA-256 hashes, process information for new executions show SHA-256 hashes, but binary entries show SHA-256 as “(unknown)” until they appear as new files on a sensor that supports SHA-256.
Start Time	The date and time of the process execution.
Interface IP	The IP address of the network adapter on the sensor. **Note:** Pre-5.1 sensors do not report an Interface IP.

Field	Description
Server Comms IP	The IP address from which the server recognizes the sensor that is reporting data. If the sensor is communicating through a Proxy or NAT, the address will be for the Proxy or NAT, not the sensor itself.

Binary Metadata

To view digital signature information and metadata about the process executable file, click the down-arrow next to the process executable file name at the bottom of the **Process Execution Details** panel:

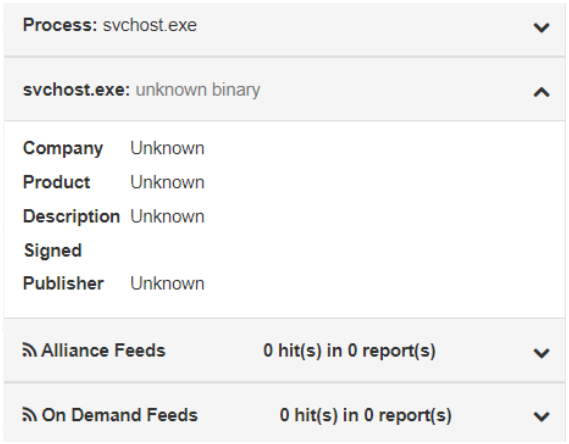

Alliance Feeds shows if the process event details had any hits from CB Threat Intel partner feeds.

If there are any hits, the results appear below the CB Threat Intel feeds in rows that are expanded by default. Each row shows:

- The source of the feed.
- A link to information about the threat that was detected.
- The date and score of the hit.
- The IOC (Indicator of Compromise) value of the process event that caused the hit.

Click the IOC hash value to go directly to the process event row for that event.

On Demand Feeds shows if the process event details had any hits from on-demand feeds.

On-demand feeds provide information from the CB Threat Intel “on demand” when a process that is part of the CB Threat Intel database is viewed on the Process Analysis page. This information includes domain classification and threat geolocation. There might not be any on-demand data available for a process that you are analyzing.

Click the **Sharing Settings** link to access the Sharing page to set this up. For more information, see “On-Demand Feeds from CB Threat Intel” on page 253.

EMET Protections Enabled (Windows Only)

The **EMET Protections Enabled** panel appears if Enhanced Mitigation Experience Toolkit (EMET) is installed on the host that reported the process and EMET Protection is enabled for the process on that host.

Process Event Filters

Several filters provide high-level details on events that occurred in the process executable file. You can access a filter by clicking the blue filter name.

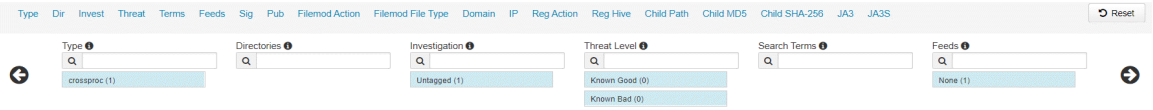

Process event filters provide a further refinement of the data that is already on the Process Analysis page – they do not do a search for new results.

Filter details are provided in the filter rows beneath the blue filter name you select. The left/right arrows let you scroll through the available filters.

You can also hover over filter rows to see the number of events that were affected.

The following table provides a description of each filter. The **Menu Bar Name** column contains the abbreviated name of the filter, as shown in the filter menu bar.

Filter	Menu Bar Name	Description
Type	**Type**	Shows process event types. See "Process Event Details" on page 193. • **filemod** – file modifications • **modload** – number of modules loaded • **regmod** – (Windows only) registry modifications • **netconn** – number of network connections enabled • **childproc** – child processes • **fork** – (OS X and Linux only) fork processes • **posix_exec** – (OS X and Linux only) posix_exec processes • **crossproc** – (Windows only - not supported on Windows XP/2003) cross processes • **blocked** – process blocked due to ban • **emet** – (Windows only) EMET mitigation
Directories	**Dir**	The directories used in this process.
Investigation	**Invest**	The tagged status for events in this process for any investigations.

Filter	Menu Bar Name	Description
Threat Level	**Threat**	Shows report scores for events associated with CB Threat Intel hits in this process. See "Threat Intelligence Feed Scores" on page 246
Search Terms	**Terms**	Shows matching query terms used in searching for processes.
Feeds	**Feeds**	Shows CB Threat Intel feed hits found in this process.
Signature	**Sig**	The signature status types of all modules that were loaded by this process (for example, signed, unsigned, or expired).
Publisher	**Pub**	The publishers of all the modules that were loaded by this process.
FileMod Action	**FileMod Action**	The types of file modifications that occurred during the execution of this process (create, delete, first write, last write) and the number of times those actions occurred.
FileMod File Type	**FileMod File Type**	The types of the files that were modified.
Domain	**Domain**	The domain (DNS) names that are associated with network connections that were made by this process.
IP Address	**IP**	The IP addresses that are associated with network connections that were made by this process.
RegMod Action	**Reg Action**	(Windows only) The type of registry modification (created, deleted key, deleted value, first write, last write).
RegMod Hive	**Reg Hive**	The location of the registry that is associated with registry modification events.
Childproc Filepaths	**Child Path**	Paths to child processes that were created by this process.
Childproc md5s	**Child MD5**	MD5 files of child processes that were created by this process.
Childproc sha-256	**Child SHA-256**	SHA-256 of child processes that were created by this process.
JA3	**JA3**	JA3 fingerprint of the client TLS hello packet.
JA3S	**JA3S**	JA3S fingerprint of the server TLS hello packet.

On the far right of the filter menu bar, you can click the **Reset** button to reset all of the filters to their original state.

Event Timeline

The **Event Timeline** is useful for investigating IOCs for events that occurred at a specific time.

A legend of color-coded event types appears at the top of the timeline. These colors are carried over to the bottom two timeline graphs to represent particular event types.

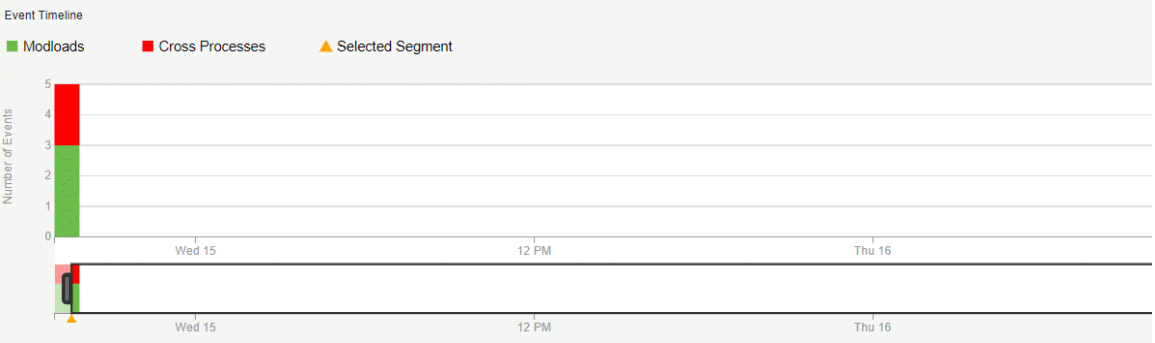

The bottom graph contains an interactive range selector widget that users can expand or collapse to zoom in on and out of the timeline. You can do this by placing your cursor on the left or right side and pressing your left mouse button; then, slide the range selector widget back and forth across the timeline. As you move the range selector widget, the Process Event Details are updated. See “Process Event Details” on page 193.

The top graph in the timeline displays event counts, which are broken down into event type segments. The top graph expands/collapses and slides back and forth in conjunction with the range selector widget. Users can zoom in on event segments to view event counts for specific time segments.

Process Event Details

The **Process Events Details** view for a selected process appears as a table with several rows at the bottom of the Process Analysis page:

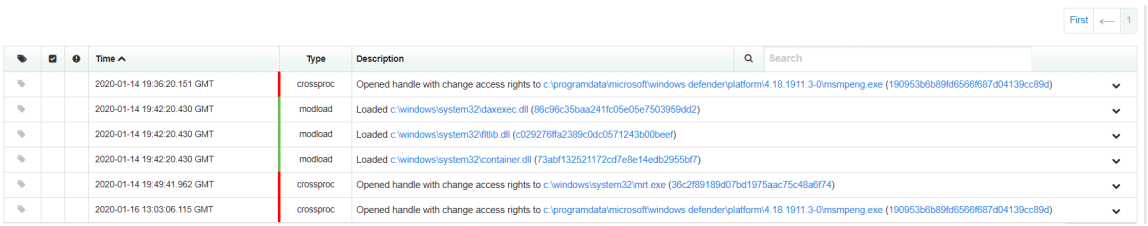

The process events rows show the following details:

Heading	Description
Tag	Shows if an event is tagged for an investigation. You can click the tag icon to select this event for a future investigation. After you select the tag icon, it turns blue to indicate that it is now included in an investigation.
Trusted Events	Shows if the event is trusted. When you click on the row, the trust information appears with a link to the source.

Heading	Description
Threat Intelligence Feed Hits	Shows if this event has matched a threat intelligence feed.
Time	The time that the event occurred in Greenwich Mean Time (GMT).
Type	The process event type. See "Process Event Types" on page 196. • **crossproc** (cross process) – appears with a red bar (Windows only - not supported on Windows XP/2003). • **child process** (child process) – appears with an orange bar. • **fork** (fork process) – appears with a yellow-orange bar (OS X and Linux only). • **filemod** (file modification) – appears with a yellow bar. • **modload** (number of modules loaded – appears with a green bar. • **posix_exec** (posix_exec process) – appears with a blue green bar (OS X and Linux only). • **regmod** (registry modification) – (Windows only) appears with a blue bar. • **netconn** (number of network connections enabled) – appears with a purple bar. • **blocked** (process blocked by hash ban) – appears with a brown bar. • **emet** (EMET mitigation) – appears with a gray bar (Windows only).

Heading	Description
Description	The operation that the **Type** event performed. See "Process Event Types" on page 196. The **Description** column might contain: • **filemod** – "Deleted" or "Created" and then provide the path to the file that was modified. • **modload** – The module that was loaded by the process. Modload descriptions can also include the path of the module that was loaded, if the module was signed or unsigned by the publisher, and the unique MD5 hash. • **regmod** – The Windows registry key that was created or modified. • **netconn** – The connection made, including the IP address (including hostname unless DNS resolution is excluded for the host), port, and protocol. • **childproc** – The child process start time, end time, and PID of the selected parent process. • **fork** – (OS X and Linux only) The instance's parent process, forked with a different Process ID (PID). • **posix_exec** – (OS X and Linux only) The instance's process that is loaded and the new binary image. • **crossproc** – (Windows only - not supported on Windows XP/2003) The action it performed; for example, opening a handle or starting or ending processes. • **blocked** – Blocked events. These are associated with the banning functionality. • **emet** – (Windows only) The EMET mitigation type reported when this process was invoked and the filename that was used in the attempt to run the process.
Search	Lets you reduce the number the events that display and focus the results based on terms entered into the **Search** box. For example, entering "Microsoft" into the **Search** box would display only Microsoft events.

You can expand an event in the results table by clicking the down arrow on the right:

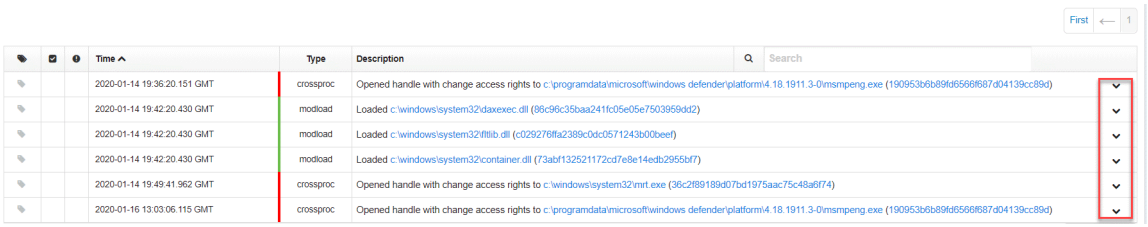

Details about the event appear. This example shows details for an event with the type **modload**:

Process Event Types

Different types of details display for each type of event, as shown in the following table:

Event Type	Details
filemod	The number of endpoints that have seen this file modification and the number of processes in which the file modification occurred on those endpoints.
modload	• The number of endpoints that have seen the MD5 hash for the module that was loaded and the number of processes in which the MD5 appears on those endpoints. • Binary information – SHA-256 hash (if available), company name, product name, a description of the binary, signature status, and publisher • CB Threat Intel information – the source of the threat intelligence feed, a link to the report for the MD5 hash, the MD5 score, and the MD5 trust status.
regmod	Windows sensors only. The number of endpoints that have seen a modification of a registry key, and the number of processes in which the registry modification occurred on those endpoints.
netconn	The number of network connections that the execution of this process either attempted or established.

Event Type	Details
childproc	• The number of endpoints that have seen the MD5 in the description and the number of processes in which the MD5 was observed. Lists the names of the processes. • Process metadata – The length of time for which the process was active, and when the process execution occurred, username of the user who is executing the process, MD5 hash, SHA-256 hash (if available), and the command line of the process executable file. • Binary information – SHA-256 hash (if available), company name, product name, product description, signature status, and publisher. • If the child process is suppressed due to Retention Maximization, then it also shows the username and command line. You choose maximization levels in the Edit Group Settings and Create Group pages. See "Advanced Settings" on page 107. This image shows suppressed vs. unsuppressed child processes. Suppressed child processes are labeled **Suppressed** in the process tree. The process tree shows a maximum of 15 child processes: either 15 unsuppressed, 15 suppressed, or 15 of both types. For processes that have more than 15 unsuppressed and 15 suppressed child processes, the tree shows unsuppressed processes first, and then suppressed processes, until a total of 15 child processes appear in the tree.
fork	(OSX and Linux only) Indicates that this is a fork process and shows the instance's parent process, forked with a different Process ID (PID). When a process performs a fork() system call, all activity for that process continues to be associated with the parent. A new **fork** event type is displayed on the Process Analysis page of the parent, indicating that the parent process performed a fork. The PID of the forked process and the timestamp of when the fork occurred is recorded

Event Type	Details
posix_exec	(OS X and Linux only) Indicates this is a posix_exec process and shows the instance's process that is loaded and the new binary image. If a process performs an exec() system call, a new process document will not be created. This activity will be reported as a new **posix_exec** event type within the process, and the process metadata will be updated to reflect the new image and command line associated with the exec() system call.
crossproc	Windows only (not supported on Windows XP/2003): Shows occurrences of processes that cross the security boundary of other processes: • Description of the OpenProcess API call for the cross process. CB Response records all OpenProcess API calls that request `PROCESS_CREATE_PROCESS`, `PROCESS_CREATE_THREAD`, `PROCESS_DUP_HANDLE`, `PROCESS_SUSPEND_RESUME`, `PROCESS_VM_OPERATION`, or `PROCESS_VM_WRITE` access rights. These access rights allow this process to change the behavior of the target process. Windows sensors only. • Process metadata – the length of time the cross process was active, username of the user who executed the process, MD5 hash, SHA-256 hash (if available), and the command line of the process executable file. • Binary metadata – SHA-256 hash (if available), the company name, product name, product description, signature status, and publisher.
blocked	The path and hash of a process that is blocked by a CB Response process hash ban. When expanded, metadata for the process and its binary appear: • Process metadata – when the process was terminated, username of the user attempting to run the process, process MD5, command line path for the process. • Binary metadata – SHA-256 hash (if available), company name, product name, product description, signature status, publisher.
emet	(Windows only) The EMET mitigation type reported when this process was invoked and the filename used in the attempt to run the process. Additional details include number of endpoints and processes that have seen the event, the time of the EMET mitigation, the EMET ID of the event, and any warnings. Output from EMET might provide additional details.

Analysis Preview Page

On the Process Search page (discussed in “Overview of Process Search” on page 173), scroll to the **Results** table (discussed in “Results Table” on page 182). Click anywhere in a query result row (except for a hyperlinked item or the **>** icon).

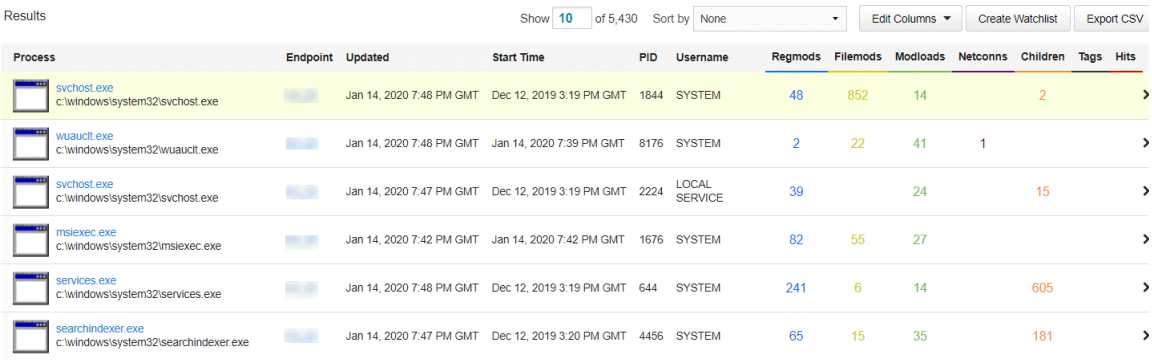

The **Analysis Preview** page appears and provides a brief overview of the process that you selected:

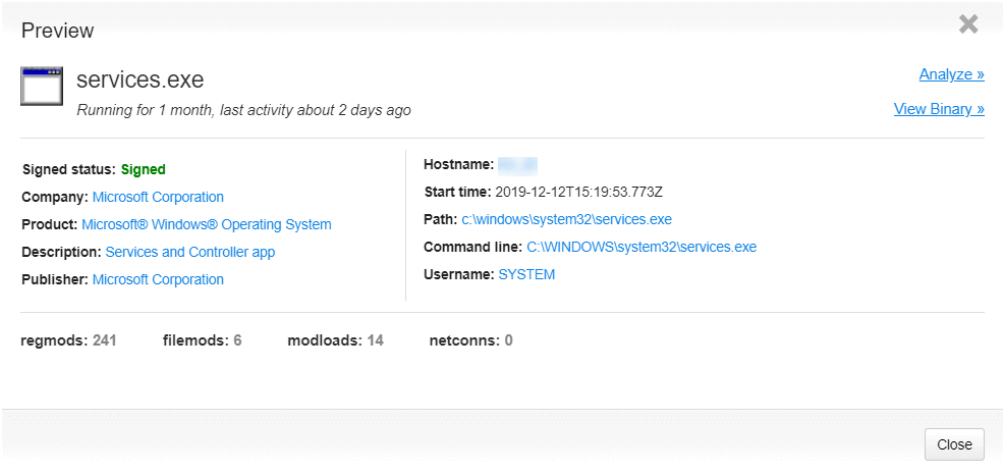

Title	Description
Signed status	Shows if the process executable file is signed by the publisher.
Company	The company name of the process executable file.
Product	The product for which the process executable file was created.
Description	A text description of the process executable file.
Publisher	The official publisher of the process executable file.
Hostname	The name of the host (endpoint) on which the process was run.
Start time	The full timestamp for the time when the process was run.
Path	The physical path from which the process was run.
Command line	The full command line specific to the execution of this process.

Title	Description
Username	The user on the given host who executed the process. The format is `<domain>\<username>`.
Regmods	The number of Windows registry modifications that were made by the process execution.
Filemods	The number of files that were modified by the execution of this process.
Modloads	The status of modules that were loaded by this process execution.
Netconns	The number of network connections that this process execution either attempted or established.
Analyze	Click to open the **Process Analysis** page for a more granular analysis of the process executable file. See “Process Analysis Page” on page 184.
View Binary	Click to view the detailed binary analysis page for the process executable file. See “Binary Search and Analysis” on page 201.

Chapter 12

Binary Search and Analysis

This chapter explains how to search for and analyze binary metadata.

Sections

Topic	Page
Overview of Binary Search	202
Entering Search Criteria	202
High-level Result Summaries	204
Related Metadata	205
Binary Search Results Table	206
Binary Preview	207
Binary Analysis	207

Overview of Binary Search

CB Response sensors begin tracking binaries when they are executed by a process. You can perform a binary search to explore the metadata of a binary.

On the navigation bar, click **Binary Search**.

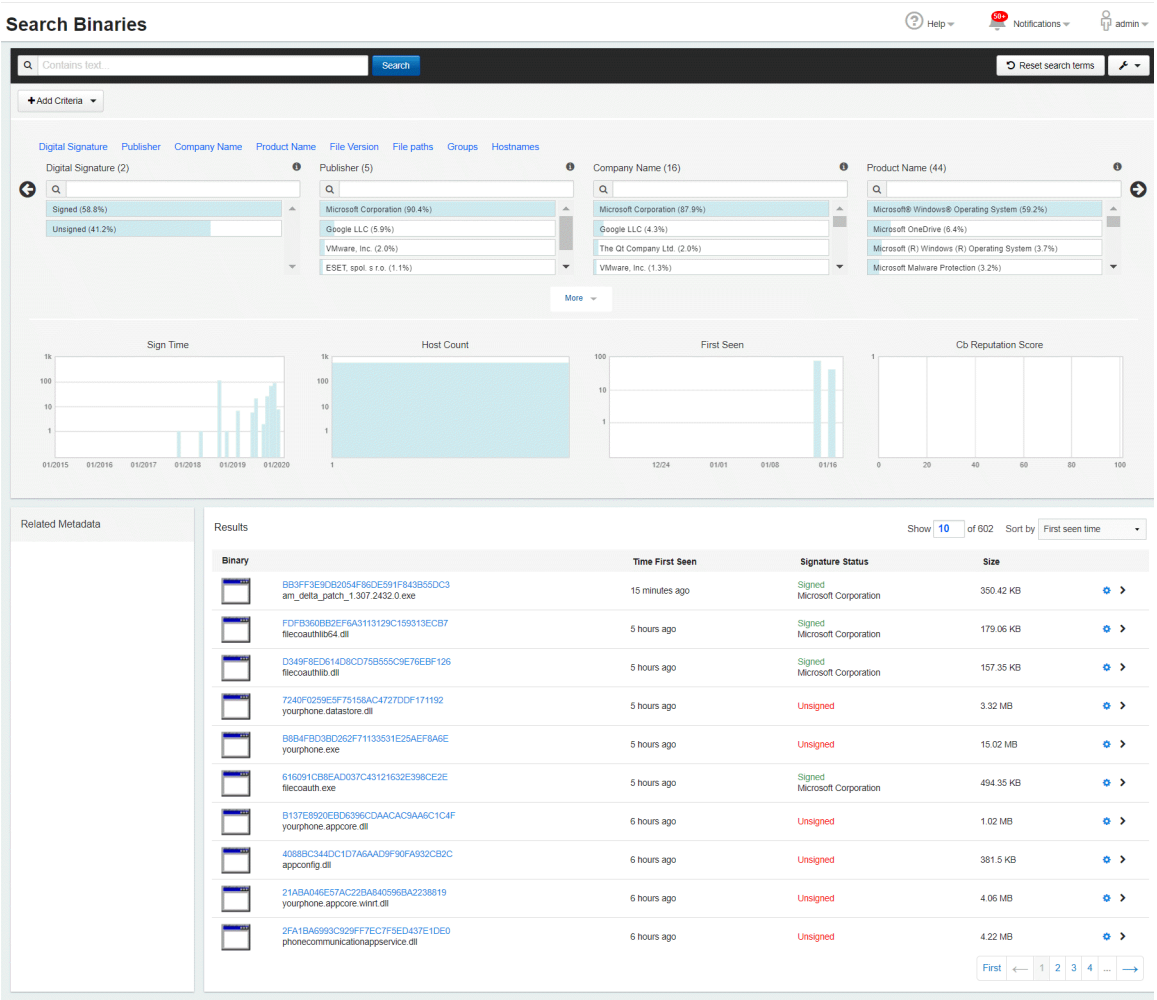

Entering Search Criteria

You can enter keyword searches or pre-defined search criteria in the **Search** box at the top of the page. While you enter search criteria, the correct syntax is displayed. However, the search not only auto-completes your criteria but estimates results as well.

If you do not enter any search criteria, the system runs a search with ***.***, which includes every binary that has executed in your environment. The results appear with a single instance of each binary and its metadata. Each binary is identified by its MD5 hash value.

To perform a binary search:

1. On the navigation bar, click **Binary Search**.
2. In the **Search** box, enter a search string (formatted with the correct syntax) or click **Add Criteria** to display predefined search criteria options:

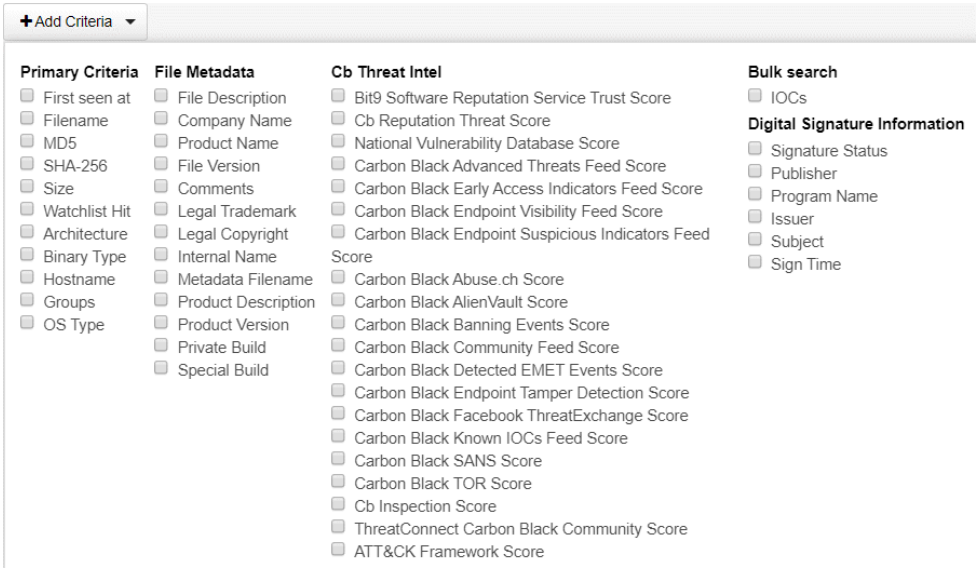

3. If you select a search criteria option, you must specify details for that search criteria option. For example, if you select the **OS Type** search criteria option, you must select one or more OS types for this search and then click **Update.**

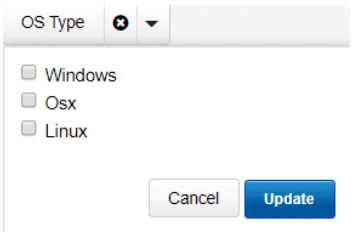

   If you add multiple search criteria fields, they are combined using an `AND` operator.

4. When you finish entering search criteria, click **Search**.

   Search results appear in a series of facets and graphs, along with a results table.

> **Note**
>
> For detailed information about using queries in CB Response, see "Advanced Search Queries" on page 213.

Additional Binary Search Page Features

In the top-right corner of the page, an **Actions** menu provides several options:

- **Share** – Share query strings. You can email the URL of the CB Response server with a query string to another CB Response user. That user can then use the string to view the same results in their own CB Response console.
- **Add Watchlist** – Create a watchlist that is based on the current query string. A watchlist is a saved search that you can use to track specific IOCs. See Chapter 18, “Watchlists”.
- **Export CSV** – Export the first 1000 process search results to a CSV file in a comma-separated value format for reporting, retention, or compliance. Each row contains a URL to access the result details.

> **Note**
>
> To export more than 1000 rows of data, you must configure API functionality to capture and save the data. See the Carbon Black Developers Network at https://developer.carbonblack.com/reference/enterprise-response/.

The **Reset search terms** button at the top right of the Search Binaries page removes all search criteria and restores the default view using ***.*** as the search criteria.

High-level Result Summaries

When you click **Search**, the Binary Search page updates the results data with information that is specific to your search criteria. The results are displayed in a variety of formats that allow you to quickly find suspicious binaries.

A summary of the results appears in facets (small tables and graphs that provide high-level result data). Each process that matches your search criteria appears in a row below the facets.

The following figure shows two rows of facets:

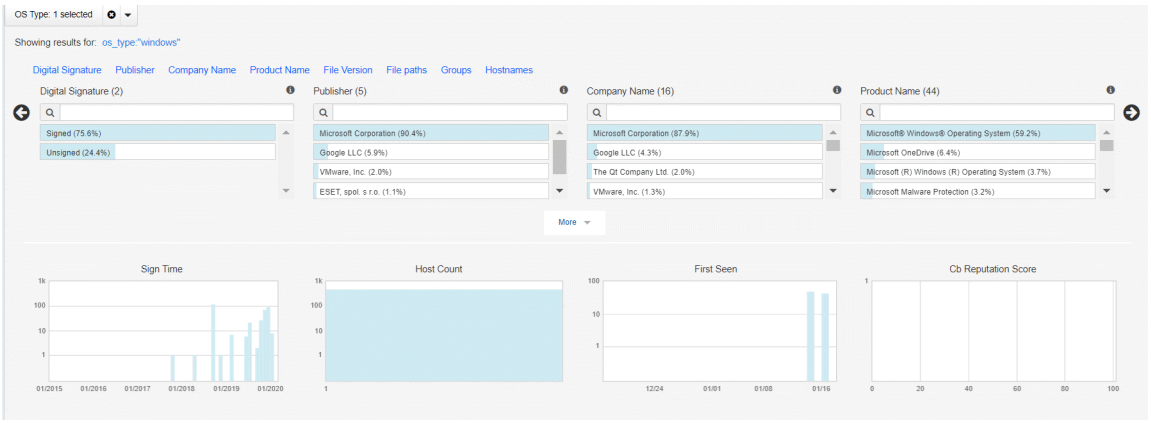

Facets provide a high-level summary of your current search results. Click the information icons to learn more about each facet.

The top row of facets contains information about the binary search results. Click the right-arrow to see all facets in this row.

- **Digital Signature** – The percentages of signed, unsigned, explicit distrust, and expired binaries.
- **Publisher** – A list of binary publishers and the percentage of binaries that have those publishers.
- **Company Name** – A list of binary publisher companies and the percentage of binaries with those company names.
- **Product Name** – The product name of the binary.
- **File Version** – The file version of the binary.
- **File Paths** – A list of file paths where files matching the current binary search have been seen.
- **Groups** – A list of the sensor groups that have identified binaries.
- **Hostnames** – A list of host names for computers on which binaries have been identified.

The second facet row contains graphs. Clicking on a facet within a graph filters the results to show the items that match that value. By default, these facets are sorted by the highest-to-lowest percentage.

Hovering over a facet within a graph displays binary counts.

The second facet row displays the following information about binaries in the results:

- **Sign Time** – The number of binaries that were signed on a particular date.
- **Host Count** – The number of binaries that were seen by CB Response on a host or a number of hosts.
- **First Seen** – The number of binaries that were first detected on a particular date.
- **CB Reputation Score** – The number of binaries that match the current search listed by CB Reputation Score.

Related Metadata

Below the facets and to the left of the binary search results, the **Related Metadata** panel appears. If you hover over an item in **Related Metadata**, rows that correspond with the selected common elements are highlighted to the right.

Binary Search Results Table

At the bottom of the page (to the right of **Related Metadata**) the binary search results table appears. Each row provides details about binary metadata that matches the search criteria.

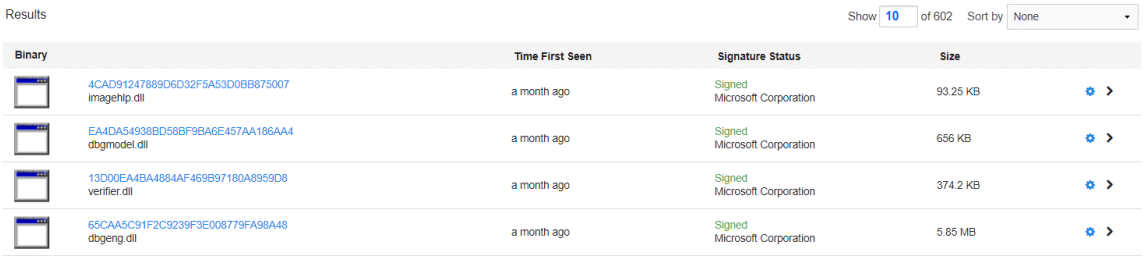

Above the search results, you can see how many binaries match the search criteria and selected filters. You can select sorting options for the list of binaries.

Search results provide the following information about the binaries in the list:

Title	Description
Icon	The icon of the file in which the binary was detected. For example: Click to display the **Binary Preview** page. See "Binary Preview" on page 207.
Binary MD5 Hash	The MD5 hash value of the binary.
Time First Seen	The first time that the binary was seen.
Signature	Shows whether the binary file is signed or unsigned.
Size	The size of the file that contains the binary.
	Indicates whether an existing watchlist identified the binary. Click the icon to open the watchlist. See Chapter 18, "Watchlists."
>	Click to display the **Binary Analysis** page. See "Binary Analysis" on page 207.

Binary Preview

Click the icon at the left of a row to view the **Binary Preview** page:

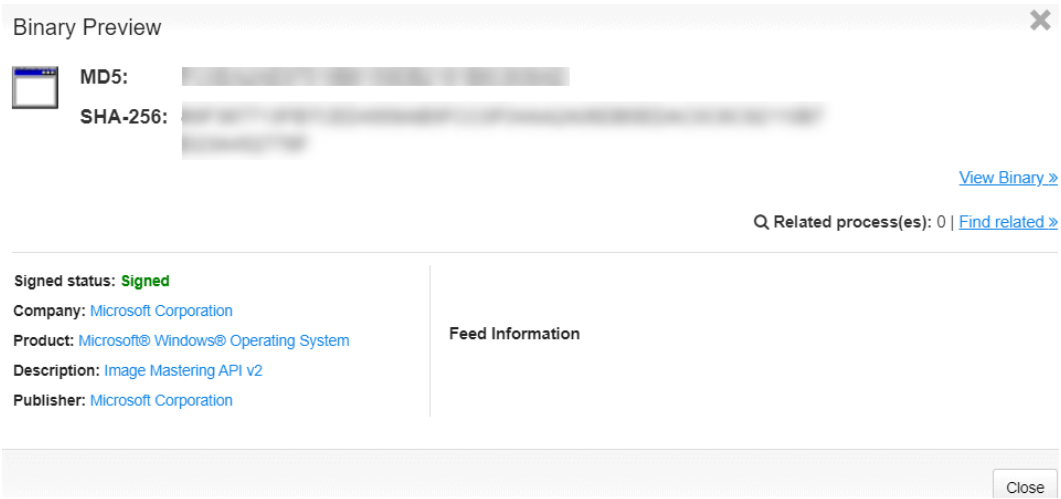

At the top of the page, the hashes of the binary (MD5 and, if available, SHA-256) appear. The file name(s) that the binary has used are listed beneath the hash value (if available).

The **Binary Preview** page provides a quick overview of the following details:

- **Metadata**:
  - **Signed status** – The status of whether the binary file is signed by the publisher.
  - **Company** – The company name identified in the metadata of the binary file.
  - **Product** – The product name identified in the metadata of the binary file.
  - **Description** – A text description of the binary file.
  - **Publisher** – The official publisher of the binary file.
- **Feed Information** – A list of CB Threat Intel feed scan results. You can click on the blue links to go to the source of the results.

At the top right of the page, the following options appear:

- **View Binary** – Click to view the detailed Binary Analysis page. See “Binary Analysis” on page 207.
- **Find related** – Click to open the Process Search page, with a predefined query for the MD5 hash value of this binary. The number of related processes displays to the left of the **Find related** link. See Chapter 11, “Process Search and Analysis.”

Binary Analysis

Use the **Binary Analysis** page to thoroughly investigate a binary. You can access the page in one of two ways:

- Click the **View Binary** link on the **Binary Preview** page.
- Click the **>** icon on the right end of a binary search results table row from the **Binary Search** page.

The **Binary Analysis** page appears:

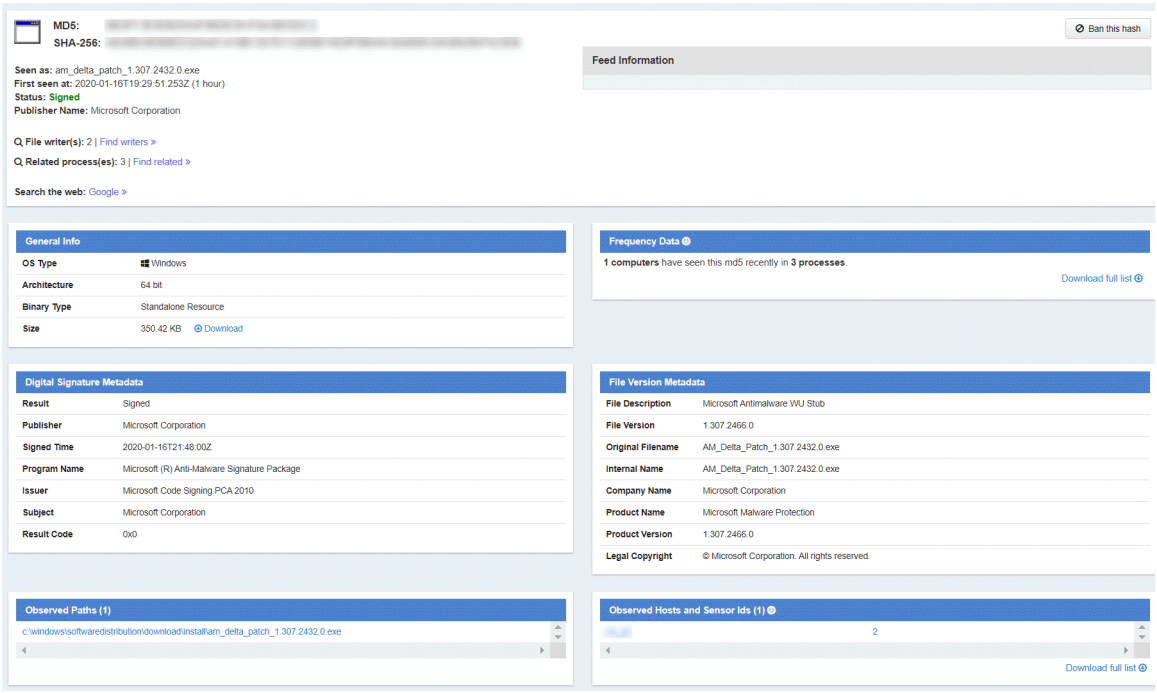

The **Binary Analysis** page contains data for investigating the binary. See the following sections for details:

- "Binary Overview" on page 208
- "General Info" on page 209
- "Frequency Data" on page 210
- "Observed Paths" on page 212
- "Observed Hosts and Sensor IDs" on page 212

Binary Overview

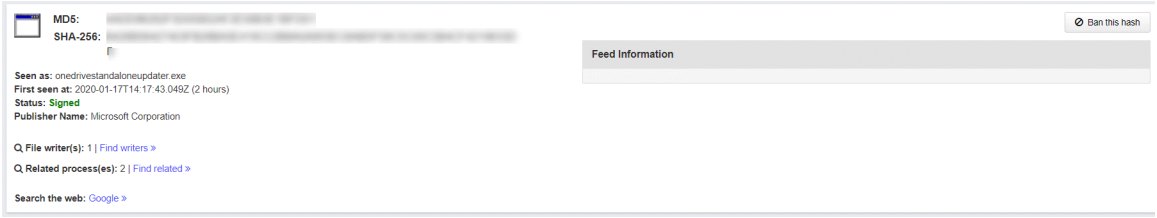

The **Binary Overview** section includes the following information:

Heading	Description
MD5 Hash Value	MD5 hash value for the binary.
SHA-256 Hash Value	The SHA-256 hash value for the binary. **Note:** Availability of SHA-256 hash data is dependent upon sensor capabilities. The macOS (OS X) sensor version 6.2.4, which is packaged with CB Response Server version 6.3, sends SHA-256 hashes to the server. Check the User Exchange or Carbon Black Support for information about other sensors that can generate SHA-256 hashes. For files that were originally discovered by a sensor that did not provide SHA-256 hashes, process information for new executions show SHA-256 hashes, but binary entries show SHA-256 as “(unknown)” until they appear as new files on a sensor that supports SHA-256.
Ban this hash	Click this button to ban this hash. Banning a hash terminates a process, if running, and prevents it from running in the future. See “Banning Process Hashes” on page 162.
Seen as	Filenames that were seen for binaries that match this MD5 hash value.
First seen at	Full time stamp of the time that this binary was last observed by currently installed sensors.
Status	Signature status — either **Signed** or **Unsigned**.
Publisher Name	Name of the binary publisher.
File writer(s)	Number and names of files the binary has written to. Click the **Find Writers** link to view the files on the Process Search page.
Related Process(es)	Number of processes that have used this binary. Click the **Find related** link to find related process on the Process Search page.
Search the web	Performs a Google search for the MD5 hash value of the binary.
Feed Information	Shows scan results for this binary from CB Threat Intel feeds. Click the links to see the results.

General Info

General Info shows the following details about the binary file:

Heading	Description
OS Type	Binary operating system.
Architecture	Binary architecture — **32-bit** or **64-bit**.
Binary Type	Binary resource type — **Standalone** or **Shared**.
Size	Size of the binary file. Also provides a link to download the physical binary.

Download	Click the **Download** button to download a copy of this binary in a zip file with a name derived from the MD5 hash of the file (for example, `A96E734A0E63B7F9B95317125DDEA2BC.zip`).
	The zip file contains two files: metadata and filedata.
	The metadata file is a text file that contains a timestamp and original filename.
	For example:
	`Timestamp: 01/17/2020 09:50:56`
	`OrigFilename: :\Device\HarddiskVolume1\ProgramData\Microsoft\Windows Defender\Definition Updates\{0B1E4D3A-9612-462F-8067-B0EDCE49CBF2}\mpengine.dll`

Frequency Data

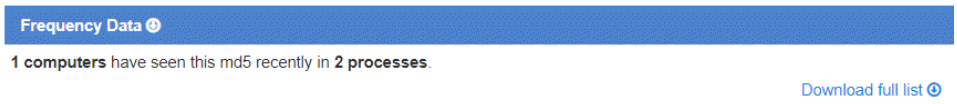

Frequency Data shows how many hosts have observed the binary that has this MD5 hash value. Click the down-arrow or click **Download full list** to download the list of hosts into a CSV file.

Digital Signature Metadata

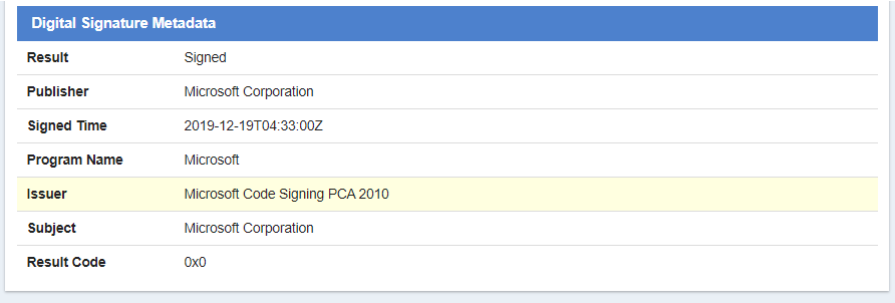

Digital signature metadata about the binary is as follows:

Heading	Description
Result	The status of the binary signature — either **Signed** or **Unsigned**.
Publisher	The publisher of the binary.
Signed Time	The time that the binary was signed.
Program Name	The binary program name.
Issuer	The binary issuer.
Subject	The binary subject.
Result Code	The result or exit code that followed the execution of the binary.

File Version Metadata

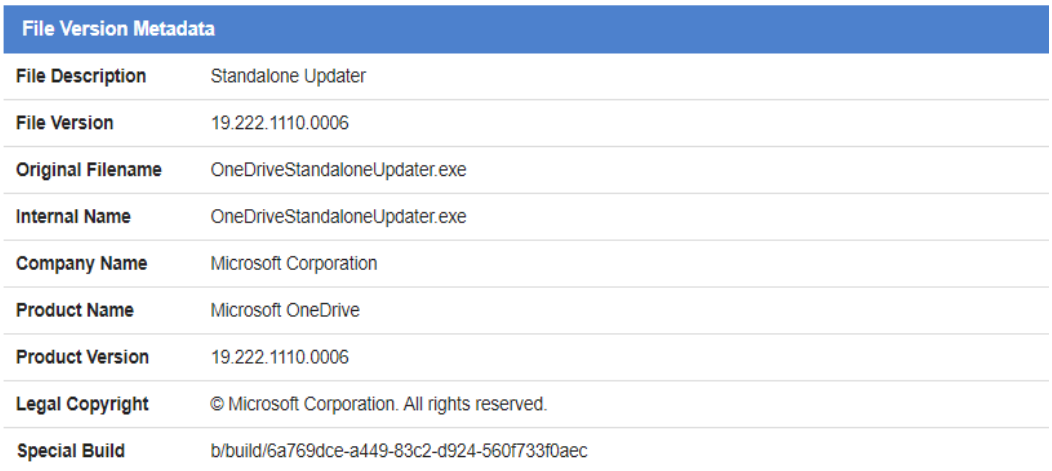

File Version Metadata	
File Description	Standalone Updater
File Version	19.222.1110.0006
Original Filename	OneDriveStandaloneUpdater.exe
Internal Name	OneDriveStandaloneUpdater.exe
Company Name	Microsoft Corporation
Product Name	Microsoft OneDrive
Product Version	19.222.1110.0006
Legal Copyright	© Microsoft Corporation. All rights reserved.
Special Build	b/build/6a769dce-a449-83c2-d924-560f733f0aec

File version metadata about the binary is as follows:

Heading	Description
File Description	Binary name (from the publisher).
File Version	Binary version
Original Filename	Binary filename.
Internal Name	Internal name of the binary
Company Name	Company name of the binary.
Product Name	Product name of the binary.
Product Version	Product version of the binary file.

Heading	Description
Legal Copyright	Copyright details for the file, including its publisher.
Special Build	Any special build data for the binary.

Observed Paths

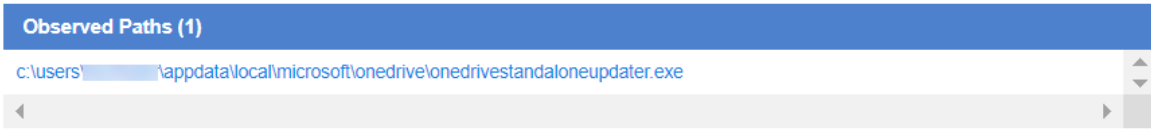

Observed paths are the full physical paths from which the binary was loaded.

Observed Hosts and Sensor IDs

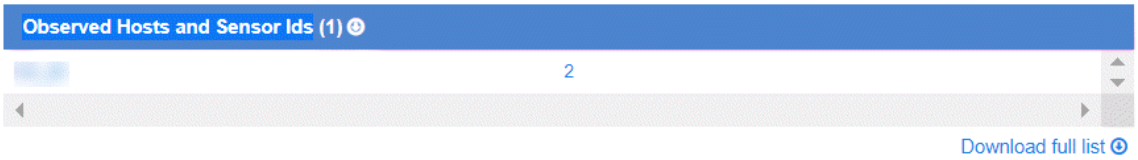

Observed Hosts and Sensor IDs shows the names of hosts on which this binary was observed together with the ID number of the sensor. Click the down-arrow or click **Download full list** to download the list of hosts into a CSV file.

Chapter 13

Advanced Search Queries

The CB Response console provides a check box interface to choose criteria for searches of processes, binaries, alerts, and threat reports. This chapter describes how to construct complex queries. The fields, field types, and examples in this chapter focus on queries to search for processes and binaries, but most of the syntax descriptions also apply to alerts and threat reports.

Sections

Topic	Page
Query Syntax Details	214
Fields in Process and Binary Searches	216
Fields in Alert and Threat Report Searches	222
Field Types	224
Searching with Multiple (Bulk) Criteria	231
Searching with Binary Joins	231
Example Searches	234

Query Syntax Details

CB Response supports multiple types of operators and syntax that can form complex queries in the **Search** boxes on the Process Search, Binary Search, Threat Report Search, and Triage Alerts pages.

Searches are generally case-insensitive.

Terms, Phrases, and Operators

A term is a single keyword (without whitespace) that is searched in the CB Response process or binary data store, or in the alerts or threat reports on your server. For example, a keyword could be: `svchost.exe.`

Terms can be combined by logical operators and nested to form complex queries; for example:

- and, AND, or whitespace — Boolean `AND` operator: `svchost.exe cmd.exe`, `svchost.exe and cmd.exe`
- or, OR — Boolean `OR` operator: `svchost.exe or cmd.exe`
- - — Boolean `NOT` operator: `-svchost.exe`
- nesting using parenthesis: `(svchost.exe or cmd.exe) powershell.exe”`
- Wildcard searches with *****; for example, `process_name:win*.exe`

Terms can be limited to a single field with *<field>:<term>* syntax; for example:

`process_name:svchost.exe`

Multiple terms are connected with `AND` if not otherwise specified.

Terms that are not preceded by fields are expanded to search all default fields.

Because terms are whitespace-delimited, use double quotes, or escape whitespaces with a single backslash, when required.

For example:

`path:"microsoft office\office15\powerpnt.exe"`

or

`path:microsoft\ office\office15\powerpnt.exe`

Terms can be combined to form phrases. A phrase is a set of terms that are separated by whitespace and enclosed in quotes. Whitespace between the terms of a quoted phrase is not treated as a logical `AND` operator. Instead, a phrase is searched as a single term.

For example: `“svchost.exe cmd.exe”`

Phrases can be combined and nested with other phrases and terms using logical operators.

For example: `"svchost.exe cmd.exe" or powershell.exe`

Restrictions on Terms

Whitespace

Whitespace is the default delimiter. A query with whitespace is “tokenized” and parsed as multiple terms.

For example:

This input: `microsoft office\office15\powerpnt.exe`

is interpreted as two terms: `microsoft AND office\office15\powerpnt.exe`

Use quotation marks to avoid automatic parsing into individual terms.

For example:

This input: `"microsoft office\office15\powerpnt.exe"`

Is interpreted as: `microsoft office\office15\powerpnt.exe`

Alternatively, you can escape whitespaces by using the backslash (\).

For example:

This input: `microsoft\ office\office15\powerpnt.exe`

Is interpreted as: `microsoft office\office15\powerpnt.exe`

See "path" on page 227 for more information about how whitespaces and slashes affect path tokenization.

Parenthesis

Parentheses are used as a delimiter for nested queries. A query with parentheses is parsed as a nested query, and if a proper nesting cannot be found, a syntax error is returned.

For example:

This input: `c:\program files (x86)\windows`

is interpreted as: `c:\program AND files AND x86 AND \windows`

Use quotation marks around the whole phrase to avoid automatic nesting. Otherwise, escape the parentheses (and whitespaces) using the backslash (\).

For example:

This input: `c:\program\ files\ \(x86\)\windows`

is interpreted as: `c:\program files (x86)\windows`

Negative Sign

The negative sign is used as logical `NOT` operator. Queries that begin with a negative sign are negated in the submitted query.

For example:

This input: `-system.exe`

is interpreted as: `not system.exe`

This input: `-alliance_score_srstrust:*`

is interpreted as: `Return all results that are not trusted by the alliance.`

You can use a phrase query to avoid automatic negation.

Double Quotes

Double quotes are used as a delimiter for phrase queries. A query in which double quotes should be taken literally must be escaped using backslash (\).

For example, the following query input:

```
cmdline:"\"c:\program files
\(x86\)\google\update\googleupdate.exe\" /svc"
```

is interpreted to match the following command line (with the command line including the quotes as shown):

```
"c:\program files (x86)\google\update\googleupdate.exe\" /svc
```

Leading Wildcards

The use of leading wildcards in a query is not recommended unless absolutely necessary, and is blocked by default. Leading wildcards carry a significant performance penalty for the search.

For example, the following query is not recommended:

```
filemod:*/system32/ntdll.dll
```

The same results would be returned by the following query, and the search would be much more efficient:

```
filemod:system32/ntdll.dll
```

Note

While process searches with leading wildcards are blocked by default beginning in CB Response 6.2.3, you can change this either through the Advanced Settings page or the `cb.conf` file. See “Managing High-Impact Queries” on page 180.

Fields in Process and Binary Searches

This section contains a complete list of fields that are searchable in CB Response process and binary searches. Some fields are valid in only one of the two, and some in both. Any binary-related field that the process search uses actually searches the executable file backing the process.

If a query specifies a term without specifying a field, the search is executed on all default fields. Default fields are indicated by `(def)`.

Note

Availability of SHA-256 hash data is dependent upon sensor capabilities. The macOS (OS X) sensor version 6.2.4, which is packaged with CB Response Server version 6.3, sends SHA-256 hashes to the server. Check the User Exchange or Carbon Black Support for information about other sensors that can generate SHA-256 hashes.

For files that were originally discovered by a sensor that did not provide SHA-256 hashes, process information for new executions show SHA-256 hashes, but binary entries show SHA-256 as “(unknown)” until they appear as new files on a sensor that supports SHA-256. This applies to all SHA-256 related fields.

Field	Process Search	Binary Search	Field Type	Description
blocked_md5	x (def)	-	md5	MD5 of a process blocked due to a banning rule.
blocked_status	x	-	status	Status of a block attempt on a running process due to a banning rule, one of the following: a-ProcessTerminated b-NotTerminatedCBProcess c-NotTerminatedSystemProcess d-NotTerminatedCriticialSystemProcess e-NotTerminatedWhiltestedPath f-NotTerminatedOpenProcessError g-NotTerminatedTerminateError
childproc_count	x	-	count	Total count of child processes created by this process.
childproc_md5	x (def)	-	md5	MD5 of the executable backing the created child processes.
childproc_sha256	x (def)	-	sha256	SHA-256 of the executable backing the created child processes (if available).
childproc_name	x (def)	-	keyword	Filename of the child process executables.
cmdline	x (def)	-	cmdline	Full command line for this process.
comments	-	x (def)	text	Comment string from the class FileVersionInfo.
company_name	x	x (def)	text	Company name string from the class FileVersionInfo.
copied_mod_len	x	x	count	Number of bytes collected.
crossproc_count	x		count	Total count of cross process actions by an actor process.
crossproc_md5	x		md5	MD5 of an actor process that performed a cross process action on a target process.
crossproc_sha256	x		sha256	SHA-256 of an actor process that performed a cross process action on a target process (if available).
crossproc_name	x		keyword	Name of an actor process that performed a cross process action on a target process.

Field	Process Search	Binary Search	Field Type	Description
crossproc_type	x (def)		processopen remotethread processopentarget remotethreadtarget	• **processopen** (or process_open) finds processes which opened a handle into another process with a set of access rights. Sample results: OpenThread() API call requested THREAD_GET_CONTEXT, THREAD_SET_CONTEXT, THREAD_SUSPEND_RESUME access rights. • **remotethread** (or remote_thread) finds processes which injected a thread into another process. Sample results: CreateRemoteThread API used to inject code into target process. • **processopentarget** is similar to processopen above, but instead of finding the actor process returns the targeted process, i.e., the process which the handle is opened into. • **remotethreadtarget** is similar to remotethread above, but instead of finding the actor process returns the targeted process, i.e., the process which the thread was injected into.
digsig_issuer	x	x (def)	text	If digitally signed, the issuer.
digsig_prog_name	x	x (def)	text	If digitally signed, the program name.
digsig_publisher	x	x (def)	text	If digitally signed, the publisher.
digsig_result	x	x (def)	sign	If digitally signed, the result. Values are: • “Bad Signature” • “Invalid Signature” • “Expired” • “Invalid Chain” • “Untrusted Root” • “Signed” • “Unsigned” • “Explicit Distrust”

Field	Process Search	Binary Search	Field Type	Description
digsig_sign_time	x	x	datetime	If digitally signed, the time of signing.
digsig_subject	x	x (def)	text	If digitally signed, the subject.
domain	x (def)	-	domain	Network connection to this domain.
file_desc	x	x (def)	text	File description string from the class FileVersionInfo.
file_version	x	x (def)	text	File version string from the class FileVersionInfo.
filemod	x (def)	-	path	Path of a file modified by this process.
filemod_count	x	-	count	Total count of file modifications by this process.
filewrite_md5	x (def)	-	md5	MD5 of file written by this process.
filewrite_sha256	x (def)	-	md5	SHA-256 of file written by this process (if available).
group	x (def)	x (def)	keyword	Sensor group this sensor was assigned to at the time of process execution.
has_emet_config	x	-	bool	True or False - Indicates whether process has EMET mitigations configured/enabled.
has_emet_event	x	-	bool	True or False - Indicates whether process has EMET mitigation events.
host_count	-	x	integer	Count of hosts that have seen a binary.
host_type	x (def)	-	keyword	Type of the computer: workstation, server, or domain controller.
hostname	x (def)	x (def)	keyword	Hostname of the computer on which the process was executed.
internal_name	x	x (def)	text	Internal name string from the class FileVersionInfo.
ipaddr	x	-	ipaddr	Network connection to or from this IP address. Only a remote (destination) IP address is searchable regardless of incoming or outgoing.

Field	Process Search	Binary Search	Field Type	Description
ipv6addr	x	-	ipv6addr	Network connection to or from this IPv6 address. Only a remote (destination) IP address is searchable regardless of incoming or outgoing.
ipport	x	-	integer	Network connection to this destination port.
is_64bit	x	x	bool	True if architecture is x64.
is_executable_image	x	x	bool	True if the binary is an EXE (versus DLL or SYS).
ja3	x	-	keyword	JA3 fingerprint of the client TLS hello packet.
ja3s	x	-	keyword	JA3S fingerprint of the server TLS hello packet.
last_server_update	x	-	datetime	Last activity in this process in the server's local time.
last_update	x	-	datetime	Last activity in this process in the computer's local time.
legal_copyright	x	x (def)	text	Legal copyright string from the class FileVersionInfo.
legal_trademark	x	x (def)	text	Legal trademark string from the class FileVersionInfo.
md5	x (def)	x (def)	md5	MD5 of the process, parent, child process, loaded module, or a written file.
sha256	x (def)	x (def)	sha256	SHA-256 of the process, parent, child process, loaded module, or a written file (if available).
modload	x (def)	-	path	Path of module loaded into this process.
modload_count	x	-	count	Total count of module loads by this process.
netconn_count	x	-	count	Total count of network connections by this process.
observed_filename	x	x (def)	path	Full path of the binary at the time of collection.
orig_mod_len	x	x	count	Size in bytes of the binary at time of collection.
original_filename	x	x (def)	text	Original name string from the class FileVersionInfo.

Field	Process Search	Binary Search	Field Type	Description
os_type	x	x	keyword	Type of the operating system: Windows, OSX or Linux.
parent_id	x	-	long	The internal CB Response process guid for the parent process.
parent_md5	x (def)	-	md5	MD5 of the executable backing the parent process.
parent_sha256	x (def)	-	sha256	SHA-256 of the executable backing the parent process (if available).
parent_name	x (def)	-	keyword	Filename of the parent process executable.
path	x (def)	-	path	Full path to the executable backing this process.
private_build	x	x (def)	text	Private build string from the class FileVersionInfo.
process_id	x	-	long	The internal CB Response process guid for the process.
process_md5	x (def)	-	md5	MD5 of the executable backing this process.
process_sha256	x (def)	-	sha256	SHA-256 of the executable backing this process (if available).
process_name	x (def)	-	keyword	Filename of the executable backing this process.
product_desc	x	x (def)	text	Product description string from the class FileVersionInfo.
product_name	x	x (def)	text	Product name string from the class FileVersionInfo.
product_version	x	x (def)	text	Product version string from the class FileVersionInfo.
regmod	x (def)	-	path	Path of a registry key modified by this process.
regmod_count	x	-	count	Total count of registry modifications by this process.
sensor_id	x	-	long	The internal CB Response sensor guid of the computer on which this process was executed.
server_added_ timestamp	-	x	datetime	Time this binary was first seen by the server.
special_build	x	x (def)	text	Special build string from the class FileVersionInfo.

Field	Process Search	Binary Search	Field Type	Description
start	x	-	datetime	Start time of this process in the computer's local time.
tampered	x	x	bool	True if attempts were made to modify the sensor's binaries, disk artifacts, or configuration
username	x (def)	-	keyword	User context with which the process was executed.
watchlist_<id>	x	x	datetime	Time that this process or binary matched the watchlist query with <id>.

Fields in Alert and Threat Report Searches

Different sets of fields are searchable on the **Triage Alerts** and **Threat Report Search** pages. As with process and binary searches, if no field is specified for a term, the search is executed on all default fields. In the tables below, default fields are indicated by (def).

Field	Field Type	Description
alert_severity	float	Overall score of the alert (combines report score, feed rating, sensor criticality). For more information, see "Threat Intelligence Feed Scores" on page 246.
alert_type	keyword	Type of the alert: one of "watchlist.hit.ingress.binary", "wathclist.hit.ingress.process", "watchlist.hit.query.process", "watchlist.hit.query.binary", "watchlist.hit.ingress.host"
assigned_to	keyword (def)	Name of the CB Response administrator who triaged (changed the status) of the alert.
create_time	datetime	Date and time this feed report was created.
created_time	datetime	Creation time of the alert.
description	text (def)	Description of the feed report, whitespace tokenized so each term is individually searchable.
domain	domain (def)	A domain IOC value in the feed report.
feed_category	text (def)	Category of this report/feed, whitespace tokenized.
feed_id	int	Numeric value of the feed id (-1 for watchlists).

Field	Field Type	Description
feed_name	keyword (def)	Name of the feed that triggered the alert. All user-created watchlists have the feed name "My Watchlists" as a special case.
group	keyword	Sensor group name of the endpoint on which the process/binary that triggered the alert was observed.
hostname	keyword (def)	Hostname of endpoint that the process/binary that triggered the alert was observed on.
ioc_value	keyword (def)	Value (IP address, MD5, or SHA-256) of the IOC that caused the alert to be triggered.
ipaddr	ipaddr	An IP address IOC value in the feed report.
ipv6addr	ipv6addr	An IPv6 address IOC value in the feed report.
is_ignored	bool	Indicates whether the report has been marked to be ignored on this server.
md5	md5 (def)	MD5 of the process that triggered the alert, or an MD5 IOC value in the feed report.
observed_filename	keyword (def)	Full path name of the process triggered the alert (not tokenized).
process_name	keyword (def)	Filename of the process that triggered the alert.
process_path	path (def)	Full path to the executable backing the process.
report_id	keyword	Name or unique identifier of the threat report that is part of the field.
report_score	float	Report score of the feed that triggered the alert. For more information, see "Threat Intelligence Feed Scores" on page 246.
resolved_time	datetime	Time this alert was triaged by a resolution action.
sha256	sha256 (def)	SHA-256 of the process that triggered the alert (if available), or a SHA-256 IOC value in the feed report.
status	keyword	Status of the alert: one of "resolved", "unresolved", "in progress", "false positive".
tags	text (def)	Tags related to this report/feed, whitespace tokenized.
title	text	Text title of the feed report, whitespace tokenized.
update_time	datetime	Date and time this feed report was last updated.
username	keyword (def)	Username in whose context the process that triggered the alert event was executed.

Field	Field Type	Description
watchlist_id	int (def)	Numeric value of the watchlist id (not applicable to feeds).
watchlist_name	keyword (def)	Name of the watchlist or the report (for feeds).

Field Types

domain

Domains are split into labels for query purposes. For example, "`example.com`" is split into "`example`" and "`com`".

If provided in a query, "dot" separator characters (.) between labels are maintained to enable position-dependent domain searches.

This has the following results:

- *Leading dot after the label, no trailing dot* – Returns results for matching labels that are at the *end* of the domain name.
- *Trailing dot after the label, no leading dot* – Returns results for matching labels that are at the *beginning* of the domain name.
- *Leading and trailing dots surrounding the label* – Returns results for matching labels that are in the middle of the domain name (i.e., not the first or last label).
- *Two labels with a dot between them* – Treated as a search for the entire phrase, and so returns results for domains that include the entire string.
- *No dot separators* – Returns results for any domain that includes the query string anywhere in the domain name.

The following table provides examples of these different domain searches:

Search	If domain is foo.com	If domain is foo.com.au
domain:com	**match**	**match**
domain:.com	**match**	no match
domain:.com.	no match	**match**
domain:com.	no match	no match
domain:example.	**match**	**match**
domain:example.com	**match**	no match

ipaddr

IP addresses are searched with a CIDR notation:

`(ip)/(netmask)`

If the netmask is omitted, it is presumed to be 32.

For example:

`ipaddr:192.168.0.0/16` or `ipaddr:10.0.1.1`

ipv6addr

IPv6 addresses are searched with a CIDR notation:

`(ip)/(netmask)`

If the netmask is omitted, it is assumed to be 32.

For example:

`ipv6addr:fe00:b9:266:2011:28dc:43d4:3298:12e2` or `ipv6addr:fe00:b9:266:2011::0/50`

text

Text fields are tokenized on whitespace and punctuation. Searches are case-insensitive.

For example, the string from the product_name field:

`Microsoft Visual Studio 2010`

is interpreted as `microsoft AND visual AND studio AND 2010`.

Searches for any of these strings will match on the binary. Phrase queries for any two consecutive terms also match on the binary.

For example:

`product_name: "visual studio"`

count

An integer value. If it exists, the values are from `0` to `MAXINT`. It supports two types of search syntaxes:

- `X:` Matches all fields with precisely `X`. For example, `modload_count:34` for processes with exactly 34 modloads.
- `[X TO Y]:` Matches all fields with counts `>= X` and `<= Y`. For example, `modload_count:[1 TO 10]` for processes with 1 to 10 modloads.

In both cases, either `X` or `Y` can be replaced by the wildcard *. For example:

`netconn_count:*` for any process where the `netconn_count` field exists.
`netconn_count:[10 TO *]` for any process with more than 10 network connections.

datetime

Datetime fields have five types of search syntaxes:

- `YYYY-MM-DD` matches all entries on this day, for example, `start:2016-12-01` for all processes started on Dec 1, 2016.
- `YYYY-MM-DDThh:mm:dd` matches all entries within the next 24 hours from this date and time, for example, `start:2016-12-01T22:15:00` for all processes started between Dec 1, 2016 at 22:15:00 to Dec 2, 2016 at 22:14:59.
- `[YYYY-MM-DD TO YYYY-MM-DD]` matches all entries between, for example, `start:[2016-12-01 TO 2016-12-31]` for all processes started in Dec 2016.

- `[YYYY-MM-DDThh:mm:ss TO YYYY-MM-DDThh:mm:ss]` matches all entries between, for example, `start:[2016-12-01T22:15:00 TO 2016-12-01:23:14:59]` for all processes started in Dec 1, 2016 within the given time frame.
- `-Xh` relative time calculations matches all entries with a time between `NOW-10h` and `NOW`. Support units supported are h: hours, m: minutes, s: seconds as observed on the host, for example, `start:-24h` for all processes started in the last 24 hours.

As with counts, `YYYYMMDD` can be replaced the wildcard *, for example, `start:[2016-01-01 TO *]` for any process started after 1 Jan 2016.

keyword

Keywords are `text` fields with no tokenization. The term that is searched for must exactly match the value in the field, for example, `process_name:svchost.exe`.

Queries containing wildcards can be submitted with keyword queries.

For example:

`process_name:ms*.exe`.

md5

md5 fields are keyword fields with an md5 hash value.

The term searched for must exactly match the value in the field.

For example:

`process_md5:6d7c8a951af6ad6835c029b3cb88d333`.

sha256

sha256 fields are keyword fields with a SHA-256 hash value.

The term searched for must exactly match the value in the field.

For example:

`process_sha256:BCB8F25FE404CDBFCB0927048F668D7958E590357930CF620F74B59839AF2A9C`.

ja3

ja3 fields are keyword fields with a ja3 hash value. You can search for the hash value. The term searched for must exactly match the value in the field.

For example:

`ja3:669181128F1B9B03303D77C6F2EEFD128`

ja3s

ja3s fields are keyword fields with a ja3s hash value. You can search for the hash value. The term searched for must exactly match the value in the field.

For example:

`ja3s:679183361F1C6F13201C62F6F2CFED111`

path

Path fields are special text fields. They are tokenized by path hierarchy as follows:

`path:c:\windows.`

For a given path, all subpaths are tokenized. For example:

`c:\windows\system32\boot\winload.exe`

is tokenized as:

`c:\windows\system32\boot\winload.exe`

`windows\system32\boot\winload.exe`

`system32\boot\winload.exe`

`boot\winload.exe`

`winload.exe`

Wildcard Searches

For queries involving path segments that are not tokenized, wildcard searches can be submitted.

For example, you can enter:

`path:system*`

for any path that has `system` as sub-path in it.

Modload Path Searches

When performing a loadable module filename (modload) search (as shown in “path” on page 227), leading forward and back slashes are tokenized. You do not have to remove the leading slash for modload path searches, although it is recommended.

For example:

`\boot\winload.exe`

should be entered as:

`boot\winload.exe`

Regmod Path Searches

When performing a Windows registry (regmod) search, a few important search caveats exist:

- If a regmod search term contains `controlset001` or `controlset002`, the search term is normalized and tokenized as `currentcontrolset`. As a result, you should search by replacing `controlsetXXX` with `currentcontrolset`.

  For example:

  `registry\machine\system\controlset001\services\xkzc`

  should be entered as:

  `regmod:registry\machine\system\currentcontrolset\services\xkzc`
- The leading backslash on regmod search terms are not tokenized. For regmod searches, be sure to omit this character when submitting search terms.

  For example:

  `\registry\machine\system\controlset001\services\xkzc`

should become:

`regmod:registry\machine\system\currentcontrolset\services\xkzc`

bool

Boolean fields have only two possible values: the string `true` or `false`. Searches are case-insensitive.

sign

Signature fields can be one of the eight possible values:

- `Signed`
- `Unsigned`
- `Bad Signature`
- `Invalid Signature`
- `Expired`
- `Invalid Chain`
- `Untrusted Root`
- `Explicit Distrust`

Values with whitespace must be enclosed in quotes.

For example:

`digsig_result:Signed` or `digsig_result:"Invalid Chain"`

cmdline

When a process launches on an endpoint, the command line for that process is sent to the CB Response server. If the server stored the whole command line as one item and allowed open ended queries of it, query performance would be extremely poor to the point of making search unusable. Instead, the server breaks each command line up into smaller component “tokens” to be stored for use when you enter a command line query.

Tokenization requires that decisions be made about which components of a command become their own token and which components are treated as delimiters between tokens. These decisions involve trade-offs since the same character may be used in different ways in a command. The following section describes how tokenization is done for CB Response cloud instances and on-premise CB Response 6.3.0 servers (and later). If you are upgrading, see also “Tokenization Changes on Server Upgrade” on page 229.

Tokenization Rules

Characters Removed Before Tokenization

With enhanced tokenization, the following characters are converted to white spaces and so removed before the command-line is tokenized:

`\ " ' ( ) [ ] { } , = < > & | ;`

Several frequently used characters are intentionally **not removed** before tokenization. These include:

- Percent (`%`) and dollar (`$`), often used for variables
- Dash (`-`), period (`.`), and underscore (`_`), often found as parts of file names
- These additional characters: `^ @ # ! ?`

Parsing Forward Slashes

The forward slash (`/`) character is handled differently depending upon its position. If it is the start of the entire command line, it is assumed to be part of the path. If it is at the start of any other token in the command line, it is assumed to be a command line switch.

There is one situation in which this parsing rule may not produce the results you want. It is not efficient for the command line parser to distinguish between a command line switch and a Unix-style absolute path. Therefore, Linux and Mac absolute paths passed on the command line are tokenized as if the beginning of the path were a command line switch. So a command line of `/bin/ls /tmp/somefile` will produce the tokens `bin`, `ls`, `/tmp` and `somefile`, incorrectly considering `/tmp` a command line switch.

Parsing Colons

The colon (:) character is handled differently depending upon its position and whether it is repeated. If it is the end of a token, it is assumed to be something the user would want to search for like a drive letter, so it is included. If there are multiple colons at the end of a token or if the colons are not at the end of a token, they are converted to white space for tokenization purposes.

File Extension Tokens

File extension tokens allow searching for either just the file extension or the entire command or file name. In other words, “word.exe” in a command line becomes two tokens: “.exe” and “word.exe”.

Wildcards

There is support for the '?' and '*' characters as wildcards when used as a non-leading character in a query, allowing you to search for any single character or multiple variable characters within a token, respectively.

> **Note**
>
> Wildcards **should not** be used as leading characters in a search.

Tokenization Changes on Server Upgrade

This section is relevant to on-premise users upgrading from a pre-6.3.0 version of CB Response. If 6.3.0 is your first version of CB Response or if you are using a cloud instance, you do not need to review this section.

Beginning with version 6.1.0, CB Response included tokenization option that improved command-line searches. This is standard for cloud instances, and beginning with version 6.3.0, it is also standard for on-premises installations. It adds the following specific improvements, which are described in more detail below:

- More special characters are removed before tokenization.
- Forward slash “/” is interpreted as a command line switch or a path character depending upon position.
- Colon “:” is interpreted as part of a drive letter token or converted to white space depending upon position and repetition.

- File extensions are stored as a separate token as well as part of a file or path name.
- Wildcards are supported in non-leading positions within a query.

These changes result in simpler queries, better and faster search results, and reduced storage requirements for tokenized command lines.

Note

If you upgraded from a pre-6.3.0 release and configured Watchlists that use command line queries, these might require a re-write to take advantage of the new tokenization. Review your Watchlist entries to make sure they return the intended results.

Example: Enhanced vs. Legacy Tokenization

The following example shows how the enhanced tokenization in version 6.3.0 differs from the previous version. It can help you convert some older queries to the new standard:

```
"C:\Windows\system32\rundll32.exe" /d
srrstr.dll,ExecuteScheduledSPPC
```

Using **legacy** tokenization, the command was broken into the following tokens:

```
“c:
windows
system32
rundll32.exe”
d
srrstr.dll,executescheduledsppc
```

The **enhanced** tokenization in version 6.3.0 breaks the same command into the following tokens:

```
c:
windows
system32
rundll32.exe
.exe
/d
srrstr.dll
.dll
executescheduledsppc
```

Examples of new search capabilities due to this tokenization include:

- You can search for .exe or .dll as part of the command line query.
- Because of more complex parsing of the forward slash, you can explicitly search for a ‘/d’ command line argument and not have to worry about false positives from just searching for the letter ‘d’.
- You can use a wildcard and search for ‘"execute*’ if you want to find a specific term passed to the command line.

- You do not have to include extraneous single or double quote marks to find a drive letter or command path.

Retention Maximization and cmdline Searches

On the Edit Group page for a sensor group, you can specify **Retention Maximization** options that help control the information that is recorded on the server to manage bandwidth and processing costs (see “Advanced Settings” on page 107). As part of this feature, the process cmdline field for parent processes store also store the cmdlines of their child processes (childprocs) that are affected by a retention setting. This is done because these childprocs do not have process documents of their own to store this information and so the expanded parent cmdline provides a way to search cmdlines for processes no longer recorded separately.

A side-effect of including the cmdlines of these childprocs in the parent’s cmdline info is that a cmdline search intended to match only the parent process’s cmdline will also match against the children. This can result in the parent process getting falsely tagged as a feed hit based on matching a childproc that was not judged to be interesting enough to justify the creation of a complete process doc. Keep this in mind when choosing **Retention Maximization** settings.

Searching with Multiple (Bulk) Criteria

You can search for multiple IOCs by using bulk search criteria in both the Process Search and Binary Search pages. While you could just enter a chain of “ORed” terms, CB Response provides special interfaces for bulk searches that do this for you when given a list of terms. You can type or paste multiple terms into a bulk search text box, following these syntax requirements:

- Each term must be on its own line.
- No punctuation is required or allowed (for example, no comma-separated lists or parentheses).
- You must use the “ipaddr:” prefix to successfully use a list of IP addresses in a bulk search.
- For most other types of data, such as md5s, prefixes are optional but more efficient. See “Fields in Process and Binary Searches” on page 216 for a table of search criteria types and their prefixes.

If a bulk search is initiated using terms without prefixes, the search is treated as a generic text search and will match the terms listed to any field. In the case of IP addresses without the “ipaddr” prefix, the search will fail because the terms are dealt with as individual numbers rather than four-part addresses.

Bulk IOC searches can be added to other search criteria or used as the only criteria for a search.

To do a bulk IOC search on the Process Search page:

1. On the Process Search page, unless you have already entered some terms to include in your search, click the **Reset Search** button under the search box to start with a fresh search.
2. Click **Add Search Terms**. Click the **Choose Criteria** dropdown menu and click **Bulk IOC > IOCs**.

3. In the text box, type or paste the list of IOCs to search for, making sure they meet the syntax requirements described in this section. For example:

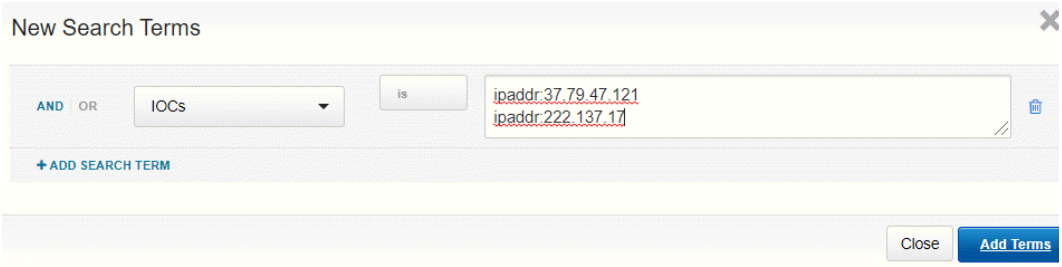

4. Although for most search criteria, you are likely to be interested in records that match one of the items on your list, you also can choose to get results that do not match your terms. Use the **is / is not** toggle in the dialog to make this choice.
5. To include additional search criteria, click the **Add Search Term** link.
6. When you have finished defining your search, click the **Add Terms** button.

   Your search is initiated and the results (if any) are shown in the table on the Process Search page. If necessary, you can continue to refine your search by using the search facet tables or you can manually enter terms.

To do a bulk IOC search on the Binary Search page:

1. On the Binary Search page, unless you have already entered some terms to include in your search, click the **Reset Search Terms** button to start with a fresh search.
2. Click the **Add Criteria** dropdown menu and, under **Bulk search**, select **IOCs**.
3. In the text box, type or paste the list of IOCs to search for, making sure they meet the syntax requirements described in this section.

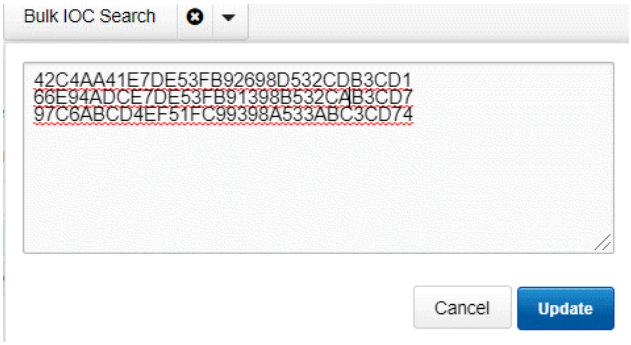

4. Click **Update** to apply the search terms.

   Your search is initiated and any results are shown in the table on the Binary Search page. If necessary, you can continue to refine your search using the search facet tables or by manually entering terms.

Searching with Binary Joins

Some binary search fields can be used as part of a process search query. (For more information, see “Fields in Process and Binary Searches” on page 216.)

In this case, the results returned are process instances that are backed by binaries that match the binary search criteria. This is called a *joined search*. For example, consider submitting the following query on the **Process Search** page:

```
digsig_result:Unsigned
```

This query returns all process instances that are backed by an unsigned MD5. By default, join searches are performed against the MD5 of the standalone process executable (`process_md5`). However, joined searches can also be performed against the MD5 of the following related events:

- `filewrites = <binary field>_filewrite`
- `parent processes = <binary_field>_parent`
- `child processes = <binary_field>_child`
- `modloads = <binary_field>_modload`

Specify the search by adding the following suffixes to the end of the binary search field:

- `filewrite`
- `parent`
- `child`
- `modload`

For example:

```
digsig_result_modload:Unsigned
```

This query returns all process instances that have loaded an unsigned module.

Note

Process searches involving large binary joins are blocked by default beginning in CB Response 6.2.3. See “Managing High-Impact Queries” on page 180 to modify this behavior.

Example Searches

Process Search Examples

Example Query Strings	Result
domain:www.carbonblack.com	Returns all processes with network connections to or from domains matching the given FQDN.
domain:.com	Returns all processes with network connections to or from domains matching `*.com`
domain:.com.	Returns all processes with network connections to or from domains matching the form `*.com.*`
domain:www.	Returns all processes with network connections to or from domains matching the form `www.*`
domain:microsoft	Returns all processes with network connections to or from domains matching `*.microsoft OR *.microsoft.* OR microsoft.*`
ipaddr:127.0.0.1	Returns all processes with network connections to or from IP address `127.0.0.1`
ipaddr:192.168.1.0/24	Returns all processes with network connections to or from IP addresses in the network subnet `192.168.1.0/24`
ipv6addr:fe00:b9:266:2011:28dc:43d4:3298:12e2	Returns all processes with network connections to or from IPv6 address `fe00:b9:266:2011:28dc:43d4:3298:12e2`
ipv6addr:fe00:b9:266:2011::0/50	Returns all processes with network connections to or from IPv6 addresses in the range of network subnet `fe00:b9:266:2011::0/50`
modload:kernel32.dll	Returns all processes that loaded a module `kernel32.dll` (accepts path hierarchies).
modload:c:\windows\system32\sxs.dll	Returns all processes that loaded a module matching path and file `sxs.dll` (accepts path hierarchies).
path:c:\windows\system32\notepad.exe	Also returns all processes with the matching path (accepts path hierarchies).

Example Query Strings	Result
regmod:\registry\machine\system\currentcontrolset\control\deviceclasses* **Notes:** Substitute “controlset001” or “controlset002” with “currentcontrolset”, as shown in this example query string. The regmod event in the process document still uses the original string, but searches must always use "currentcontrolset". regmod searches must include the complete path string or use wildcards. Searches for partial regmod paths without wildcards never yield results.	Returns all processes that modified a registry entry with the matching path (accepts path hierarchies).
path:excel.exe	Returns all processes with the matching path (accepts path hierarchies).
cmdline:backup	Returns all processes with matching command line arguments.
hostname:win-5ikqdnf9go1	Returns all processes executed on the host with matching hostname.
group:"default group"	Returns all processes executed on hosts with matching group name (use of quotes are required when submitting two-word group names).
host_type:workstation	Returns all processes executed on hosts with matching type (use of quotes are required when submitting two-word host types).
username:system	Returns all processes executed with the matching user context.
process_name:java.exe	Returns all processes with matching names.
parent_name:explorer.exe	Returns all processes executed by a parent process with matching names.
childproc_name:cmd.exe	Returns all processes that executed a child process with matching names.
md5:5a18f00ab9330ac7539675f3f326cf11	Returns all processes, modified files, or loaded modules with matching MD5 hash values.
process_md5:5a18f00ab9330ac7539675f3f326cf11	Returns all processes with matching MD5 hash values.
parent_md5:5a18f00ab9330ac7539675f3f326cf11	Returns all processes that have a parent process with the given MD5 hash value.

Example Query Strings	Result
filewrite_md5:5a18f00ab9330ac7539675f3f326cf11	Returns all processes that modified a file or module with matching MD5 hash values.
childproc_md5:5a18f00ab9330ac7539675f3f326cf11	Returns all processes that executed a child process with matching MD5 hash values.
<type>_count:*	Returns all processes that have xxx_count field > 0, where type is one of modload, filemod, regmod, netconn, or childproc.
<type>_count:10	Returns all processes that have xxx_count field = 10, where type is one of modload, filemod, regmod, netconn, or childproc.
<type>_count:[10 TO 20]	Returns all processes that have xxx_count field >= 10 and <= 20, where type is one of modload, filemod, regmod, netconn, or childproc.
<type>_count:[10 TO *]	Returns all processes that have xxx_count field >= 10, where type is one of modload, filemod, regmod, netconn, or childproc.
<type>_count:[* TO 10]	Returns all processes that have xxx_count field < 10, where type is one of modload, filemod, regmod, netconn, or childproc.
start:2011-12-31	Returns all processes with a start date of 2011-12-31 (as observed on the host).
start:[* TO 2011-12-31]	Returns all processes with a start date earlier than or equal to 2011-12-31 (as observed on the host).
start:[* TO 2011-12-31T22:15:00]	Returns all processes with a start date earlier than or equal to 2011-12-31 at 22:15:00 (as observed on the host).
start:[2011-12-31 TO *]	Returns all processes with a start date later than or equal to 2011-12-31 (as observed on the host).
start:[2011-12-31T09:45:00 TO *]	Returns all processes with a start date later than or equal to 2011-12-31 at 09:45:00 (as observed on the host).
start:*	Returns processes with any start date (as observed on the host).
start:[* TO *]	Returns processes with any start date (as observed on the host).
start:-10h	Returns all processes with a start time between NOW-10h and NOW. Units supported are, h: hours, m: minutes, s: seconds (as observed on the host).
last_update:2011-12-31	Returns all processes last updated on date 2011-12-31 (as observed on the host).

Example Query Strings	Result
last_update:[* TO 2011-12-31]	Returns all processes last updated on a date earlier than or equal to 2011-12-31 (as observed on the host).
last_update:[* TO 2011-12-31T22:15:00]	Returns all processes last updated on a date earlier than or equal to 2011-12-31 at 22:15:00 (as observed on the host).
last_update:[2011-12-31 TO *]	Returns all processes last updated on a date later than or equal to 2011-12-31 (as observed on the host).
last_server_update:[2011-12-31T09:45:00 TO *]	Returns all processes last updated on a date later than or equal to 2011-12-31 at 09:45:00 (as observed at the server).
last_server_update:*	Returns processes with any update date (as observed on the server).
last_server_update:[* TO *]	Returns processes with any update date (as observed on the server) within the range provided.
last_server_update:-10h	Returns all processes last updated between NOW-10h and NOW. Units supported are h: hours, m: minutes, s: seconds (as observed on the server).
process_id:<guid>	Returns the process with the given process id, where `<guid>` is a signed 64-bit integer.
parent_id:<guid>	Returns the process with the given parent process id, where `<guid>` is a signed 64-bit integer.
sensor_id:<guid>	Returns processes executed on host with given sensor id, where `<guid>` is an unsigned 64-bit integer.

Binary Search Examples

Example Query Strings	Result
`md5:5a18f00ab9330ac7539675f326cf11`	Returns all binaries with matching MD5 hash values.
`digsig_publisher:Oracle`	Returns all binaries with a digital signature publisher field with a matching name.
`digsig_issues:VeriSign`	Returns all binaries with a digital signature issuer field with a matching name.
`digsig_subject:Oracle`	Returns all binaries with a digital signature subject field with a matching name.
`digsig_prog_name:Java`	Returns all binaries with a digital signature program name field with a matching name.
`digsig_result:Expired`	Returns all binaries with a digital signature status of `<status>`.
`digsig_sign_time:2011-12-31`	Returns all binaries with a digital signature date of 2011-12-31.
`digsig_sign_time:[* TO 2011-12-31]`	Returns all binaries with a digital signature date earlier than or equal to 2011-12-31.
`digsig_sign_time:[2011-12-31 TO *]`	Returns all binaries with a digital signature date later than or equal to 2011-12-31.
`digsig_sign_time:*`	Returns binaries with any digital signature date.
`digsig_sign_time:[* TO *]`	Returns binaries with any digital signature date within the range provided.
`digsig_sign_time:-10h`	Returns all binaries with a start time between NOW-10h and NOW. Units supported are h: hours, m: minutes, s: seconds.
`<type>_version:7.0.170.2`	Returns all binaries with matching version, where `<type>` is product or file.
`product_name:Java`	Returns all binaries with matching product name.
`company_name:Oracle`	Returns all binaries with matching company name.
`internal_name:java`	Returns all binaries with matching internal name.
`original_filename:mtxoci.dll`	Returns all binaries with matching filename.
`observed_filename:c:\windows\system32\mtxoci.dll`	Returns all binaries that have been observed to run on or were loaded with the given path.
`<type>_mod_len:[* TO 10]`	Returns all binaries that have `<type>_mod_len` (module length in bytes) field < 4096, where type is original or copied.

Example Query Strings	Result
`<type>_desc:"database support"`	Returns all binaries that have `<type>_desc` field with matching text, where type is file or product.
`legal_<type>:Microsoft`	Returns all binaries with matching `legal_<type>` field text, where type is trademark or copyright.
`<type>_build:"Public version"`	Returns all binaries with matching `<type>_build` field text, where type is special or private.
`is_executable_image:True or False`	Boolean search (case insensitive) returning all binaries that are executable or not executable.
`is_64bit_:True or False`	Boolean search (case insensitive) returning all binaries that are 64-bit or not 64-bit.
`watchlist_4:[2014-04-01 TO 2014-09-31]`	Returns all binaries that matched watchlist 4 during the time period shown.

Threat Intelligence Search Examples

Any document matching a threat intelligence feed is tagged with an `alliance_score_<feed>` field, where the value is a score from -100 to 100.

For more information, see "Threat Intelligence Feeds" on page 244.

<feed> is the "short name" of the threat intelligence feed, such as **`nvd`** or **`isight`**.

For any threat intelligence feed, you can click the **View Hits** button to discover the feed's short name. For more information, see Chapter 15, "Threat Intelligence Feeds".

Example Query Strings	Result
`alliance_score_<feed>:*`	Returns all binaries that have *<feed>* score > 0.
`alliance_score__score_<feed>:10`	Returns all binaries that have *<feed>* score = 10.
`alliance_score__score_<feed>:[10 TO 20]`	Returns all binaries that have *<feed>* score >= 10 and <= 20.
`alliance_score__score_<feed>:[10 TO *]`	Returns all binaries that have *<feed>* score >= 10.
`alliance_score__score_<feed>:[* TO 10]`	Returns all binaries that have *<feed>* score < 10.

Chapter 14

Ingress Filtering

Ingress filtering is a technique to manage data retention. For general information about ingress filters and their use, read CB Response Managing Retention and System Capacity before you set up your ingress filters. This chapter describes how to perform common tasks related to ingress filters.

Sections

Topic	Page
Overview of Ingress Filtering	241
Viewing and Configuring Ingress Filters	241
Regex Filters	243

Overview of Ingress Filtering

In previous versions of CB Response, ingress filters could only be set up by using an API. With the release of CB Response 7.1.0, global administrators can view, create, modify, and delete ingress filters in the CB Response console.

Note

Not all features or fields are currently exposed in the console. Advanced usage might require using the API. See Ingress Filter Details for information about using the API.

Viewing and Configuring Ingress Filters

To access the Ingress Filters page:

- Click ***Username*** **> Settings** and then click **Ingress Filters**.

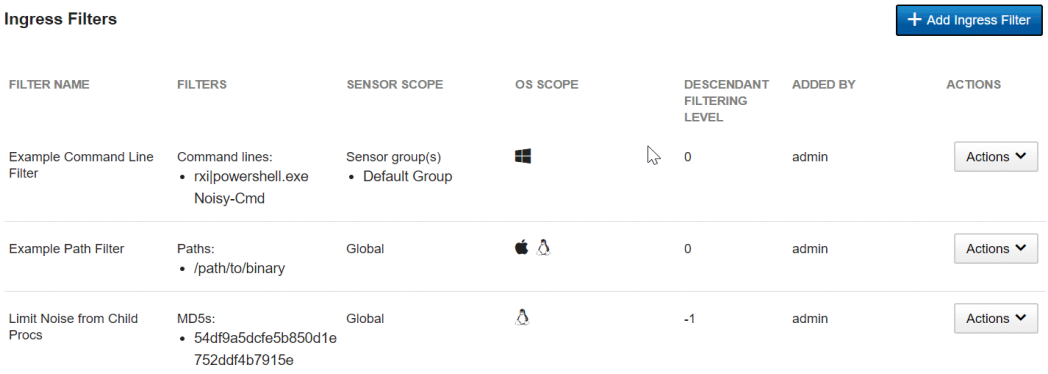

The Ingress Filters page shows the following fields:

- **Filter Name** – A unique name for this ingress filter.
- **Filters** – The ingress filter type and definition.
- **Sensor Scope** – Defines whether the ingress filter applies to specific sensor groups, individual sensors, or to all sensors.
- **OS Scope** – Defines the operating systems to which the ingress filter applies.
- **Descendant Level** – When a process is filtered, you can also filter its children, their children, etc., up to the set number of levels. For example, you can filter by:
  - -1: All descendants
  - 0: Matched process only
  - 1: Matched process and children
  - 2: Matched process, children, and next layer of descendants

    ... and so on
- **Added by** – The user who added the ingress filter

- **Actions** – Provides a dropdown menu that lets you modify and delete ingress filters.

Adding an Ingress Filter

You can add an ingress filter on the Ingress Filters page or on the Process Search page.

> **Note**
>
> Be aware that you cannot modify the name or filter type after you have added the ingress filter. Although you cannot change the filter type (for example, MD5), you can change the value of the filter (the MD5 hash).

To add an ingress filter on the Ingress Filters page:

1. Click the **Add Ingress Filter** button.
2. Fill out the following form and then click **Add**.

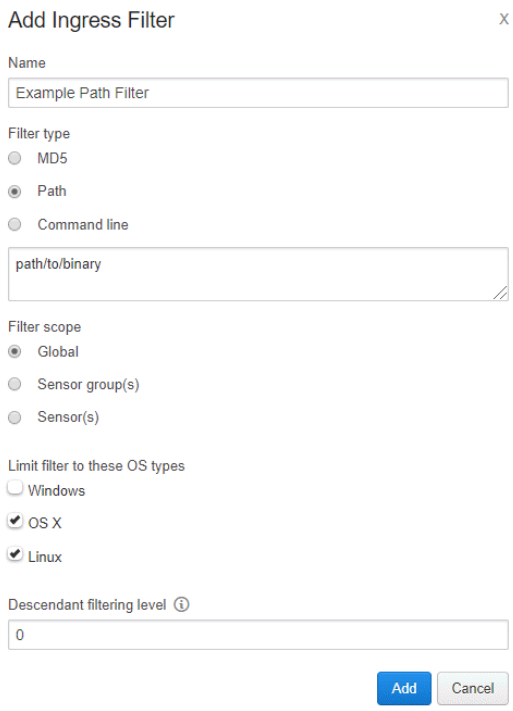

To add an ingress filter on the Process Search page:

1. On the navigation bar, click **Process Search**.
2. Perform your search and select the process from which to create an ingress filter.

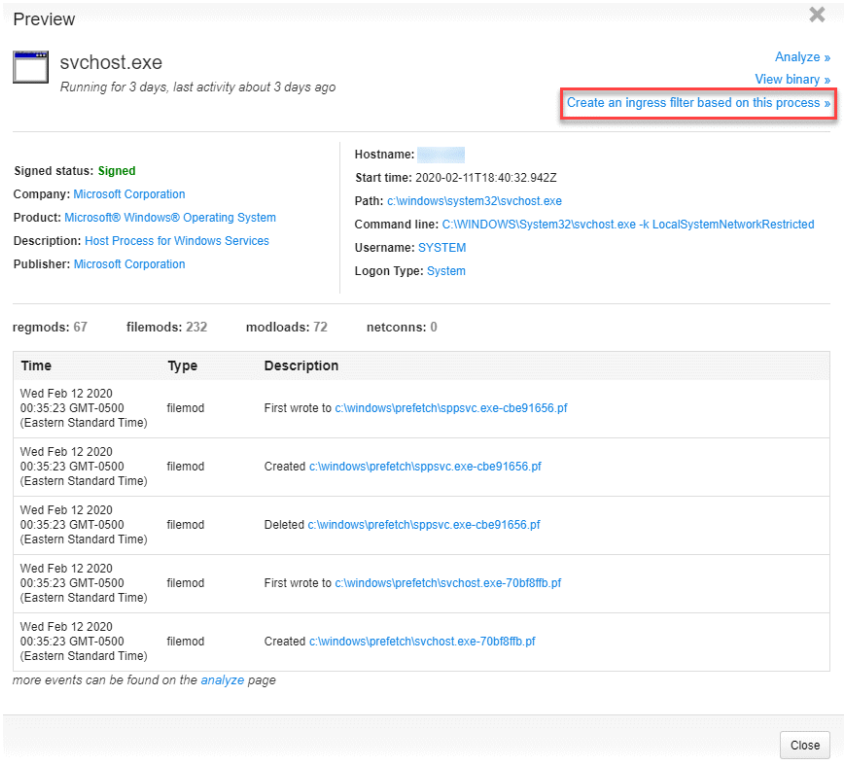

3. Select which filter type to use. It will be pre-filled from the selected process. You can edit the values. You must provide a unique filter name.

After you have added ingress filters, you can modify or delete them by using the **Actions** dropdown menu options on the Ingress Filters page.

Regex Filters

A filter can match a portion of a field. This is useful when you filter programs that are frequently executed with different command lines, such as powershell.exe or bash. To specify a regex pattern, prefix the path, command line, or MD5 value with `rx|` (or `rxi|` to specify a case-insensitive match). Regex patterns must be compatible with the Java 8 Pattern class. See https://docs.oracle.com/javase/8/docs/api/java/util/regex/Pattern.html.

For example, the following ingress filter blocks processes where the command line contains "powershell.exe Noisy-Cmd". Other invocations of powershell.exe are not blocked.

Command lines:

```
rxi|powershell.exe Noisy-Cmd
```

Chapter 15

Threat Intelligence Feeds

This chapter describes threat intelligence feeds that can be enabled on a CB Response server to enhance the verification, detection, visibility, and analysis of threats on your endpoints.

Sections

Topic	Page
Overview of Threat Intelligence Feeds	245
Managing Threat Intelligence Feeds	246
Enabling, Disabling, and Configuring a Feed	251
On-Demand Feeds from CB Threat Intel	253
Creating and Adding New Feeds	254
Searching for Threat Reports	255

Overview of Threat Intelligence Feeds

Threat intelligence feeds are streams of reports about IOCs and patterns of behaviors found in the wild by a variety of services and products. One or more feeds can be integrated into the CB Response server and console to enhance the verification, detection, visibility, and analysis of threats on your endpoints.

The source of a feed may be from:

- CB Threat Intel and the Carbon Black Threat Research Team
- A third-party Carbon Black partner
- The information and analysis collected by:
  - CB Threat Intel Reputation
  - CB Protection threat detection tools
- Shared data collected from CB Response customer enterprises

You can also create new feeds if needed. Some feeds do not require data collection from your server, while others require that you share information from your enterprise back to the feed provider to improve community intelligence data.

Available feeds appear on the Threat Intelligence Feeds page. You can enable or disable any feed on that page. The CB Response server supports the following types of IOCs:

- Binary MD5s
- Binary SHA-256s
- IPv4 addresses
- IPv6 addresses
- JA3 fingerprints
- JA3S fingerprints
- DNS names
- Query-based feeds using the CB Response process/binary search syntax to define an IOC

When a feed is enabled and IOCs from it are received, the following information and capabilities are added in CB Response:

- **Feed results added to process and binary records** – If an IOC from a feed report matches processes or binaries reported by sensors on your endpoints, the feed results are added to the records for those processes/binaries in CB Response. You can search and filter for processes or binaries using a feed report or score. For example, you can create a table of all processes whose National Vulnerability Database score is greater than 4.
- **Feed-based watchlists** – You can create a CB Response watchlist that tags a process or binary found on one of your endpoints when the score of a feed matches a specified score or falls within a specified score range.
- **Feed-based alerts** – You can configure console and email alerts when a process or binary, which is the subject of a specified feed report, is identified on an endpoint.
- **Links to feed sources** – You can link back to the source of a feed for more information, which can range from a general feed description to specific details about an IOC reported by that feed.

- **Threat Report Search** – You can search for individual threat reports from any feed that is or has been enabled.

Threat Intelligence Feed Scores

The threat intelligence feed score spectrum is as follows:

- A negative 100 (-100) score means a feed is extremely trustworthy (not in any way malicious). These scores are rare.
- A positive 100 (100) score means that a feed is extremely malicious.

Most scores will be within the 0-100 range.

Firewall Configuration for Feeds

To receive all of the threat intelligence available from CB Threat Intel, you must allow SSL access (port 443) through your firewall to the following domains:

- api.alliance.carbonblack.com:443
- threatintel.bit9.com:443

Blocking either of these will prevent your CB Response server from receiving intelligence from specific feeds as well as data, such as IP location and icon matching for files.

Managing Threat Intelligence Feeds

On the **Threat Intelligence Feeds** page, you can:

- View the available feeds and get more information about them
- Enable or disable feeds
- Configure alerts and logging for feeds
- Change the rating used to calculate the severity assigned to IOCs from a feed
- Sync one or all feeds
- Check for new feeds
- Add a new feed
- Delete user-defined feeds
- Search for threat reports

CB Threat Intel feeds are feeds that CB Response makes available from CB Response sources and third-party partners. These feeds can be enabled and (in some cases) disabled, but they cannot be deleted from the page.

Certain reports come from CB Threat Intel as on-demand feeds, and these do not provide their data until a process on the Process Analysis page matches their information. See “On-Demand Feeds from CB Threat Intel” on page 253 for more details.

The EMET Protection and Banning Events feeds send their respective events to the CB Response server regardless of whether they are enabled, but they must be enabled if you want to configure alerts and logging.

To view the Threat Intelligence Feeds page:

- On the navigation bar, click **Threat Intelligence**.

  The **Threat Intelligence Feeds** page appears, as shown in the following example.

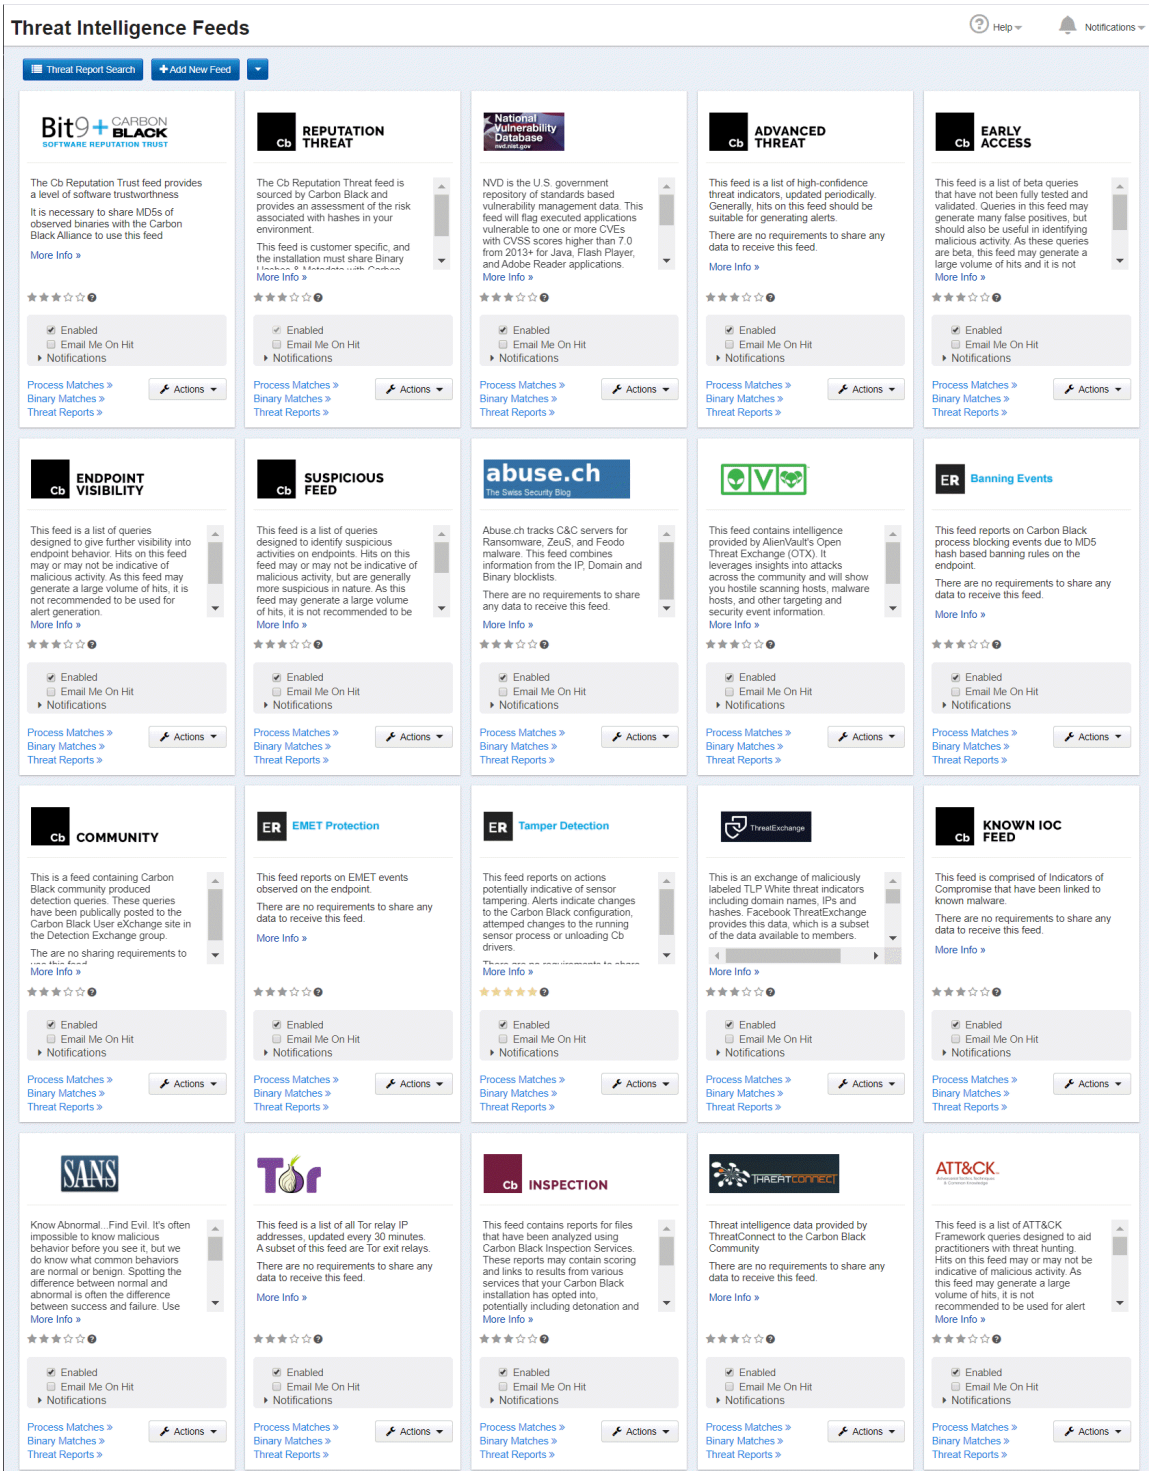

The **CB Response Tamper Detection** feed is enabled by default. It alerts on endpoint activity that indicates tampering with sensor activity:

You must enable other feeds. See “Enabling, Disabling, and Configuring a Feed” on page 251. See “Creating and Adding New Feeds” on page 254 for information about adding user-defined feeds.

Checking for New Threat Intelligence Feeds

CB Response works with a variety of partners to provide threat intelligence feeds for the CB Response server. You can add new partners and feeds to your server.

The **Check for new feeds** option on the **Threat Intelligence Feeds** page lets you check for the most recent feeds. It also removes feeds if they are no longer available. (However, existing reports and tagged processes/binaries will still identify these feeds.)

To check for new Threat Intelligence feeds:

1. On the Threat Intelligence Feeds page, click the **Action** menu (down-arrow at the top of the page) and click **Check for new feeds**.
2. If new feeds are available, they are added to the **Threat Intelligence Feeds** page.

Syncing Threat Intelligence Feeds

Threat intelligence feeds are updated periodically by the feed sources. To make certain that all feeds are up-to-date, use the **Sync All** command.

To sync all threat intelligence feeds on the page:

1. On the Threat Intelligence Feeds page, click the **Action** menu (down-arrow at the top of the page) and click **Sync All**.
2. All feeds are synced with the latest data on the **Threat Intelligence Feeds** page.

> **Note**
>
> You can sync feeds individually. Sync commands for individual feeds are available on the **Action** menu in the feed panel. See “Syncing Threat Intelligence Feeds” on page 248.

Data Sharing Settings

Most of Carbon Black’s threat intelligence feed partners provide a list of all of the IOCs they track, and almost all feeds require that you enable communication on the Sharing Settings page. Also, some feeds require that you enable data sharing.

> **Important**
>
> Management of Sharing Settings is available only to Global Administrators in on-premise installations and Administrators in cloud installations.

To enable sharing communications:

1. Click ***Username*** **> Sharing Settings**.
2. Under **General Sharing Settings**, specify sharing settings for Alliance:
   - **Enable Alliance Communication** – Enables communication with Carbon Black. It also allows download of binaries from the Alliance and the ability to retrieve Alliance feeds.
   - **Support the Alliance Threat Intelligence Community** – Enables your server to send threat intelligence statistics to Carbon Black, including alert resolutions, ignored reports, and feed ratings. These statistics improve the efficacy of Carbon Black-provided threat intelligence and give you a community consensus on the ratings of feeds and threat indicators.
3. Specify sharing settings for statistics and diagnostics data:
   - **Enable Performance Statistics** – Enables sharing of usage, resource, and sensor statistics with Carbon Black.
   - **Allow Unattended Background Upload of Diagnostics Data** – Enables the CB Response server to do background collection of diagnostics data such as application logs and configuration files to facilitate troubleshooting with Carbon Black Customer Support. Requires **Enable Performance Statistics** to be enabled.
   - **Allow Upload of Sensor Diagnostics Data** – This setting determines whether the CB Response server can upload diagnostics data gathered from deployed endpoint sensors to Carbon Black for troubleshooting. Options are as follows:

     -Disabled – No sensor diagnostics data can be uploaded to Carbon Black.

     -Manual – You can manually upload sensor diagnostics data by using a utility that is installed on the sensor.

     -Automatic – Sensor diagnostics data is automatically uploaded when fault conditions are detected on the sensor.

> **Note**
>
> **Manual** or **Automatic** upload of sensor diagnostics data requires the **Enable Performance Statistics** option to be selected.

To enable data sharing with Carbon Black threat intelligence feed partners:

1. On the **Sharing** page, scroll to **Endpoint Activity Sharing**:

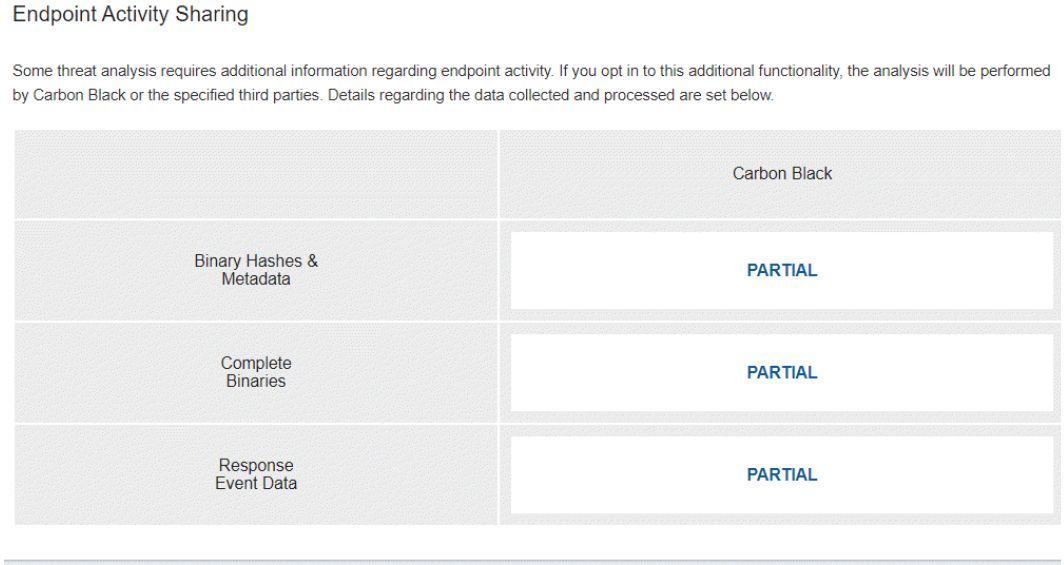

> **Note**
>
> To enable data sharing with feed partners in the Carbon Black Alliance, **Enable Alliance Communication** must first be selected.

2. Decide whether to share the following types of data with Carbon Black or CB Inspection (**Complete Binaries** only):
   - **Binary Hashes & Metadata**
   - **Complete Binaries**
   - **Response Event Data**
3. Click the current setting (**Enabled**, **Disabled**, or **Partial**) to specify the default sharing setting for sensor groups.

   In the resulting **Share** dialog box, read the **Summary**, **Data Shared**, and **Privacy** sections carefully before making further selections.

4. The following options are available at the bottom of each **Share** dialog box:
   - Select **Enable (Share from ALL Groups)** to share data from endpoints in all sensor groups.
   - Select **Disable (Do Not Share from ANY Groups)** to disable data sharing from endpoints in all sensor groups.
   - Select **Partial (Share from SOME Groups)** to share data from endpoints in some sensor groups. Use the arrows between the **SHARE FROM** and **DO NOT SHARE FROM** windows to choose the groups that are allowed to share data.

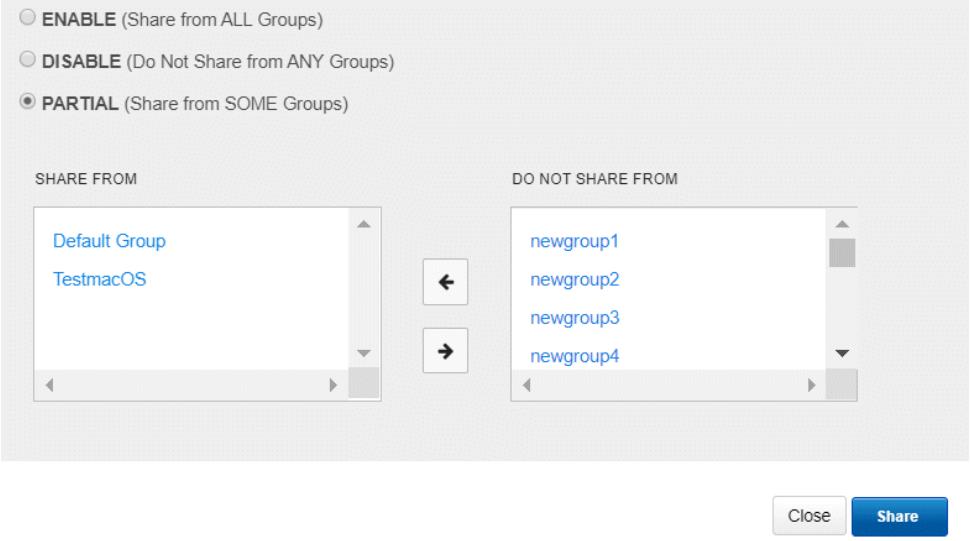

5. Click **Share** to begin sharing the data.

Enabling, Disabling, and Configuring a Feed

Each feed that is available on a CB Response server is represented by a panel on the **Threat Intelligence Feeds** page. The panel provides information about the feed and allows you to enable, disable, and configure it.

To enable and configure a threat intelligence feed:

1. In the navigation bar, select **Threat Intelligence**.
2. Locate the feed and click **Enabled**.
3. Configure the feed using the following options and controls:

Field/Menu	Description
More info	Link to the feed provider's website for technical information about the feed and general information about the provider and its products.

Field/Menu	Description
★★★☆☆ **(Rating)**	Rating of this threat intelligence feed by the community of CB Response users. The default rating is three stars. You can click a star to modify the rating of this feed on your server. The rating affects the severity assigned to alerts coming from this feed, which can affect the order of alerts if they are sorted by severity.
Enabled	If selected, the threat intelligence feed is enabled; otherwise it is disabled. **Note**: Most feeds also require that you select **Enable Alliance Communication** on the **Sharing** page. Also, feeds that upload data from your server require that you opt into hash sharing with that feed. See "Data Sharing Settings" on page 248.
Email Me on Hit	IOCs from this feed that reference a process or binary that is recorded on this CB Response server cause an email alert to be sent to the logged-in console user. See "Enabling Email Alerts" on page 294. Only Global Administrators or cloud Administrators can change this setting.
Notifications menu	• **Create Alert** – Indicators from this feed that reference a process or binary that is recorded on this CB Response server cause a console alert. See "Enabling Console Alerts" on page 286. • **Log to Syslog** – IOCs from this feed that reference a process or binary that is recorded on this CB Response server are included in Syslog output from this CB Response server. See the *CB Response Integration Guide* on the Carbon Black User Exchange for details on configuring SYSLOG output. Only Global Administrators or cloud Administrators can change these settings.
Process Matches	Link to the Process Search page with the results of a search that shows each process that matches IOCs from this feed. See Chapter 11, "Process Search and Analysis,".
Binary Matches	Link to the Binary Search page with the search results showing each binary that matches IOCs from this feed. See Chapter 12, "Binary Search and Analysis,".
Threat Reports	Link to the Threat Reports search page filtered to show any Threat Reports from this feed. See "Searching for Threat Reports" on page 255.
Actions menu	The **Actions** menu includes the following commands: • **Create Watchlist** – Creates a Watchlist, which is a saved search whose results are processes or binaries that match IOCs that this feed reports. • **Incremental Sync** – Adds report data from this feed that has been observed since the previous synchronization. • **Full Sync** – Rewrites all report data from this feed.

To disable a threat intelligence feed:

1. On the navigation bar, click **Threat Intelligence** and identify the feed to disable.
2. Deselect the **Enabled** check box within that feed panel.

If you disable a feed, its reports remain on the server and incoming data will be tagged against the locally existing IOCs that it reported. However, when you disable a feed:

- Reports from these feeds about IOCs will not be downloaded for scanning.
- For feeds that require data to be sent to them, new binary MD5s from your sensors will not be sent.

On-Demand Feeds from CB Threat Intel

CB Threat Intel provides a rich variety of intelligence and capabilities about files, domain names, IP addresses, and associated patterns of compromise, including IOCs, reputation, and attack classification.

Examples of these intelligence types include:

- Trust and threat ratings
- Domain/IP reputation and context
- Icon matching to help identify threats masquerading as files of another type
- Detection feeds of behavioral patterns of compromise

Some of this intelligence can be enabled or disabled through feeds listed on the Threat Intelligence Feeds page, and this information is added to process data as soon as the feed is received.

Other intelligence is made available to the CB Response server when a process, pattern, or other IOC that is part of the CB Threat Intel database is viewed on the Process Analysis page. The information in these on-demand feeds includes:

- **Damballa malware classification and context** – CB Threat Intel provides an enhanced network-to-endpoint attack classification through its integration with Damballa's threat intelligence on malicious destinations, advanced threat actor groups, and command-and-control communications. This information is added to attack classifications when a Process Analysis page containing a relevant domain name is displayed.
- **Geolocation information for network connections** – The location of addresses identified in both inbound and outbound connections is provided.

- **Domain and IP reputation** – CB Threat Intel computes a reputation score for domains using various inputs, information, and algorithms inside the cloud. This reputation score is displayed for domain names for which a score is available.

Note

For on-demand feed information to become available and displayed for a process, the sensor group on which the process was reported must be configured to send relevant data to the CB Threat Intel for analysis. This requires explicitly opting in to share CB Response events with CB Threat Intel. This is not enabled by default; you can enable it in the **Response Event Data** row in the **Endpoint Activity Sharing** section of the Sharing page. See “Data Sharing Settings” on page 248.

Creating and Adding New Feeds

You can create and add new threat intelligence feeds to a CB Response server. A feed can be created in any language that allows for building JSON, or you can build it by hand. One way to build a feed is to use the Carbon Black Feeds API (CBFAPI), found on github at:

https://github.com/carbonblack/cbfeeds

The CBFAPI is a collection of documentation, example scripts, and a helper library to help create and validate CB Response feeds. Regardless of how a feed is created, the feed file must match the feed structure (or schema) that the "Feed Structure" section of the CBFAPI documentation defines.

You have several options about the specification you provide when adding a new feed to a CB Response server. The minimum requirement is that you provide a URL to the feed.

To add a new threat intelligence feed to the CB Response server:

1. Confirm that the feed you have created follows the “Feed Structure” instructions in the CBFAPI documentation on github.
2. On the navigation bar, select **Threat Intelligence**.
3. On the Threat Intelligence Feeds page, click **Add New Feed**.
4. In the **Edit Alliance Feed** dialog box, do one of the following:
   - To add a feed from a URL, click the **Add from URL** tab and complete the following settings:

Field	Description
Feed URL	Enter the URL for the feed itself that will be providing IOC reports.
Use Proxy	Select this option to use a proxy for the feed URL. The configuration for this proxy must be configured in advance by Carbon Black Technical Support.
Validate Server Cert	Select this option to require a validation check on the feed server’s certificate.

Field	Description
Show/Hide Feed Server Authentication Options	If the server that is providing the feed requires authentication, click the **Show Server Authentication Options** link and provide the following authentication information: • **Username** • **Password** • **Public Cert** • **Private Key**

- To manually add a feed, click the **Add Manually** tab and complete the following settings:

Field	Description
Name	Enter the feed name to appear in the panel.
Feed URL	Enter the URL that the CB Response server will use to sync the data in the feed.
Provider URL	Enter the URL to the page to open when the user clicks **More Info** on the feed panel.
Summary	Enter the text that will appear in the panel to describe this feed.
Use Proxy	If the CB Response server must access the feed URL through a proxy, the proxy is added in the proxy field.
Validate Server Cert	Indicates if the CB Response server should validate the Feed Server certificate.
Show/Hide Feed Server Authentication Options	If the server providing the feed requires authentication, click the **Show Server Authentication Options** link and provide the following authentication information: • **Username** • **Password** • **Public Cert** • **Private Key**

5. Click **Save**.

   If the settings you entered provide access to a feed server, the new feed appears on the Threat Intelligence Feeds page.

Searching for Threat Reports

You might want to obtain more information on the report types that are provided by a particular feed, or you might want to explore specific reports.

Suppose you want to filter out a high volume of uninteresting reports from a feed that you otherwise you find useful. You can search for those reports on the Threat Intelligence Feeds page and mark them to be ignored in the future.

You can also search for all reports or perform a search on a page that is pre-filtered for one feed.

To open the Search Threat Reports page (unfiltered):

1. On the navigation bar, click **Threat Intelligence** and click **Threat Report Search**.

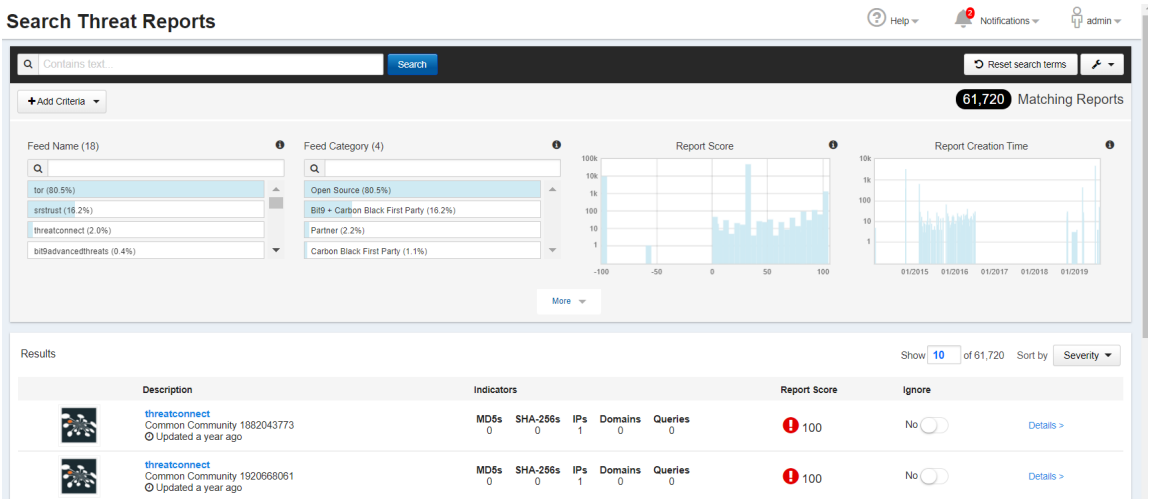

2. In the Search Threat Reports page, enter search criteria to search for the reports in which you are interested. See “Threat Report Searches and Results” on page 256.

To display a table of reports from one threat intelligence feed:

1. On the Threat Intelligence Feeds page, click the **Threat Reports** link at the bottom of the panel for the feed that contains reports to view.

   The Search Threat Reports page displays reports from the selected feed.

   You can further refine the search by using the available search options.

Threat Report Searches and Results

The Search Threat Reports page is divided into three major sections:

- The top section includes the following:
  - The **Search** field and button.
  - The **Add Criteria** button, which opens a Search Criteria page.
  - The **Reset search terms** button, which resets the search and removes any search criteria you have added.
  - The **Actions** menu, which applies to the entire page.
- The middle section contains a series of filters that include the following:
  - **Feed Name** – A list of the short names (for example, “nvd” for National Vulnerability Database) of each feed that has produced a report, and the percentage of all reports that have been produced by each feed.
  - **Feed Category** – A list of feed categories and the percentage of all reports that each feed category produces. Categories can include:

    Open Source – For example, Tor or Malware Domain List.
    Partner – A member of the CB Threat Intel Partners.

CB Response first party – Feeds supplied directly from CB Protection or CB Response products or services.

- **Report Score** – A graph of the number of reports at different score levels.
- **Report Creation Time** – A graph of the number of reports by creation date.

- The **Reports** table shows details for reports that match the search criteria. You can sort the reports by severity, most recently updated, or most recently added.

The Search Threat Reports page presents the following report data:

Column	Description
Description	This column includes: • The name of the feed that provided the report • The name of the specific report • The time elapsed since the report was received
Indicators	The column includes the number of certain elements in the report that were identified as threats: • MD5s – the number of suspicious files matching the MD5 hash • SHA-256s – the number of suspicious files matching the SHA-256 hash • IPs – the number of suspicious IP addresses • Domains – the number of suspicious domains • Queries – the number of queries in the report; depending on the feed, this value might be empty.
Report Score	The threat score of this report. Report scores range from minus 100 to 100, with lower scores indicating a lower threat and higher scores indicating a higher threat. Threat scores are used in the calculation of alert severity.
Ignore	Ignore any future instances of this report, so that they do not trigger alerts. See “Ignoring Future Reports” on page 259.
Details link	Opens a **Threat Report Details** page for the report in this row. See “Threat Report Details” on page 258.

Threat Report Details

Click the **Details** link in the far right column of a report in the Threat Report Search Reports table to see details for that threat report, if available.

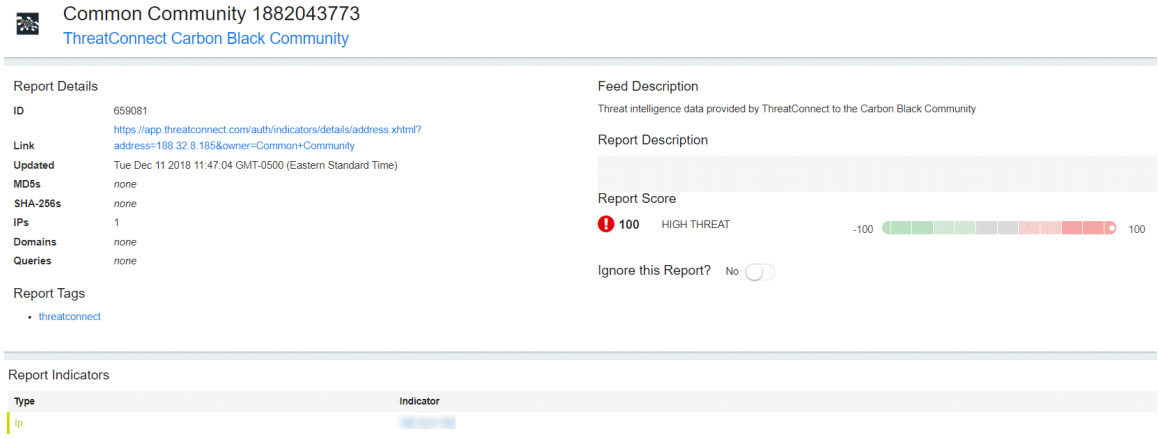

The information on the Threat Report Details page varies depending on the feed source and type of indicator. The following table describes the fields on this page.

Field	Description
Title	The feed name and the unique ID of the report.
Report Details	This section includes: • ID – the unique ID of the report • Link – if available, a link to the report on the website of the feed source • Updated – when the report was last updated • MD5s – the number of suspicious MD5s • SHA-256s – the number of suspicious SHA-256s • IPs – the number of suspicious IP addresses • Domains – the number of suspicious domains • Queries – the number of queries in the report.
Report Tags	One or more descriptive strings from the feed provider to help explain what the report is about. For example, tags can describe a specific threat, a threat category, a targeted industry, a known threat actor, or geographic information. Not all reports have tags.
Feed Description	The description of the feed given by the provider.
Report Description	The description of this report from the feed provider.
Report Score	The threat score of this report. Report scores range from minus 100 to 100, with lower scores meaning lower threat and higher scores meaning higher threat. Threat scores are one factor in the calculation of Alert severity.

Field	Description
Ignore this Report?	Ignore any future instances of this report so that they do not trigger alerts.
Report Indicators	A table of indicators that the report references (IPs, MD5s, SHA-256s, domains, queries). If the Type is MD5, clicking the indicator name links to the Binary Search page for that MD5.

Ignoring Future Reports

Feeds use a variety of criteria to decide whether a file or site is a threat, and you might not agree with the threat level that is indicated by all of the reports generated by certain feeds. When you review reports and determine that a report is not reporting an actual threat, you can ignore any future instances of reports by the same name.

To ignore reports, use one of the following options:

- On the Search Threat Reports page, on the **Actions** menu, click **Ignore all reports matching this search**.
- In the results table on the Search Threat Reports page, you can set the **Ignore** button for one report to **Yes**.
- On the Threat Report Details page, you can set the **Ignore this Report** button for one report to **Yes**.

You can also mark events to be ignored using the Triage Alerts page. See “Ignoring Future Events for False Positive Alerts” on page 293.

Chapter 16

Configuring the Event Forwarder

This chapter explains how to configure the Event Forwarder through the CB Response console.

Sections

Topic	Page
Overview of the Event Forwarder	261
Configuring the Event Forwarder in the Console	261

Overview of the Event Forwarder

The CB Response Event Forwarder is a standalone service that can export events (both watchlist/feed hits and raw endpoint events, if configured) from the CB Response enterprise bus in a normalized JSON or LEEF format. The events can be saved to a file, delivered to a network service, or automatically archived to an Amazon AWS S3 bucket. These events can be consumed by any external system that accepts JSON or LEEF, including Splunk and IBM QRadar.

The list of events to collect is configurable. By default, all feed and watchlist hits, alerts, binary notifications, and raw sensor events are exported into JSON. The configuration file for the connector is stored in `/etc/cb/integrations/event-forwarder/cb-event-forwarder.conf`.

For details on installing and manually configuring the Event Forwarder, see https://github.com/carbonblack/cb-event-forwarder.

With the release of CB Response 7.1.0 server, admins can customize the Event Forwarder from directly within the CB Response console.

By default, this feature is enabled for cloud instances, and disabled for on-prem deployments. You can enable the feature for on-prem by adding `EventForwarderEnabled=true` in `cb.conf` and restarting services.

Configuring the Event Forwarder in the Console

To view the Event Forwarder Settings page:

- On the navigation bar, click **Event Forwarder**.

  The Event Forwarder Settings page consists of four sections:

  - **Edit and status**: Allows you to edit and save or cancel changes to the configuration, displays the service status, and lets you stop/start the service.
  - **Events**: Identifies the events that will be forwarded.
  - **Output**: Configures the format and destination for the output.
  - **Certificates**: Identifies certificates to use for validation.

Edit and Status

The **Edit and Status** panel lets you edit the Event Forward configuration, view the current status of the Event Forwarder (**Service running** or **Service stopped**), and stop or start the service.

You do not have to stop the service to configure the settings in the Event Forwarder; however, you must stop and restart the service for saved changes to take affect.

Note

You must have set up the receiving service and credentials before you configure Event Forwarder for the first time. See https://github.com/carbonblack/cb-event-forwarder.

You cannot save the configuration until after you have established a valid configuration in the **Output** section.

Carbon Black validates the connection as soon as you click **Save**; therefore, it is important that the connection is viable before you set up forwarded events. If the connection is not viable, the configuration is not saved.

To configure the Event Forwarder (recommended sequence):

1. Click **Edit** at the top of the page and configure your output in the **Output** panel.
2. Click **Save**. If the connection is viable, the configuration is saved and you can proceed.
3. Click **Edit** and configure the events to be forwarded.
4. Click **Save**.
5. Stop and restart the service.

You can click **Cancel** to revert to the previous settings.

Events

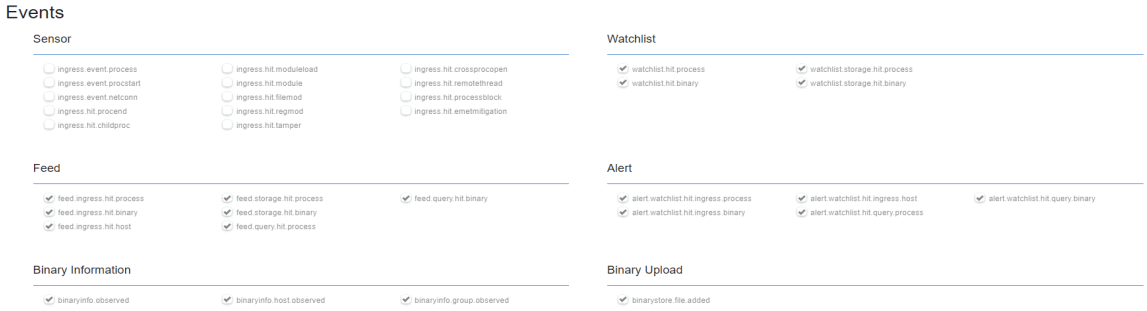

In the **Events** panel, you can configure which events are forwarded.

To configure forwarded events:

1. Click **Edit** at the top of the page.
2. Select the items to be forwarded by checking the checkboxes next to each item. To deselect an item, uncheck the checkbox.
3. If you have configured the **Output** settings, click **Save**.

Output

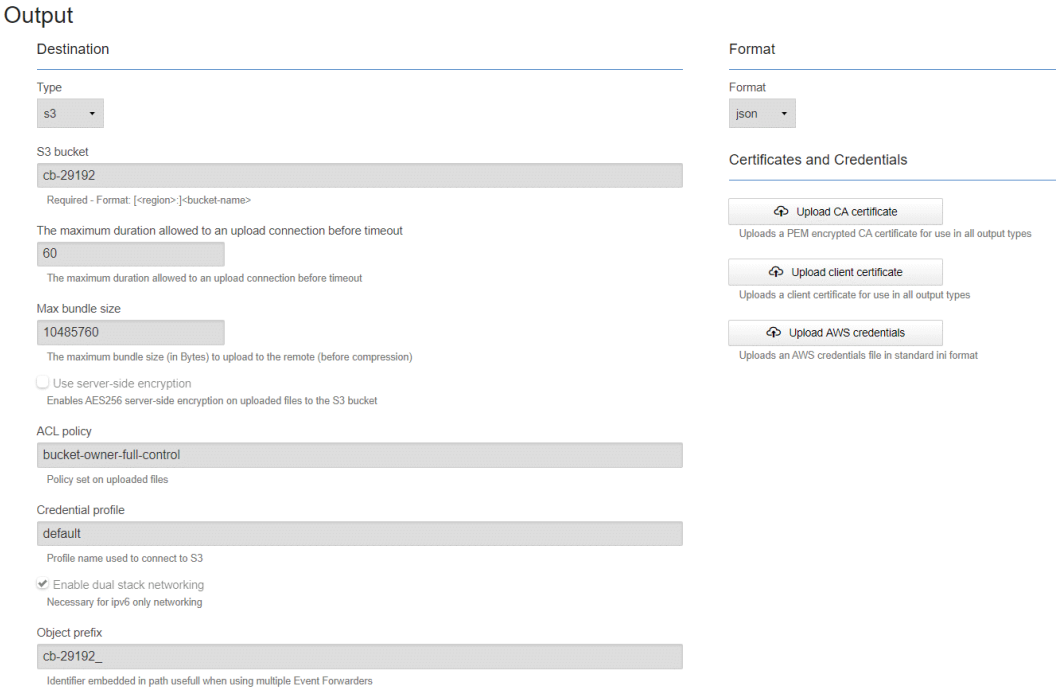

The output can be in either **LEEF** or **JSON** format, which you select in the **Format** dropdown menu. The default format is JSON.

You can select from the following types of destinations:

- **Splunk** (the default destination)
- **S3**
- **HTTP**
- **Syslog**

The required output parameters depend on the destination type.

Splunk

The Splunk destination type requires the following information:

Parameter	Description
Splunk HEC Endpoint	Required URL of the Splunk destination endpoint; for example, `http://www.example.com`.
HEC Token	Required token for HEC authorization.
Server Common Name	Optional. Common name (CN) of the destination server.
Send Timeout	Optional. Maximum duration of an upload connection. The default value is 60 seconds.
Upload Empty Files	Optional. Determines whether zero byte length files are uploaded. The default setting is false.

Parameter	Description
Max Bundle Size	Optional. The maximum bundle size (in bytes) to upload to the remote destination before compression is applied. The default value is 10MB.
Certificates and Credentials	Optional. Determines whether an uploaded certificate (identified by type) is required for the connection. Also allows you to specify AWS credentials.

S3

The S3 destination type requires the following information:

Parameter	Description
S3 Bucket	Required. The name of the S3 bucket to receive the output. The format must be `[<region>:]<bucket-name>`.
Send Timeout	Optional. Maximum duration of an upload connection. The default value is 60 seconds.
Upload Empty Files	Optional. Determines whether zero byte length files are uploaded. The default setting is false.
Max Bundle Size	Optional. The maximum bundle size (in bytes) to upload to the remote destination before compression is applied. The default value is 10MB.
Server-Side Encryption	Optional. Identifies the type of encryption that is required on the server. The default type is AES256.
ACL Policy	Required. Settings can be READ, WRITE, READ_ACP, WRITE_ACP, or FULL_CONTROL. These are typical permissions: the _ACP permissions also allow read/write to the ACL of the bucket itself. See Access Control List (ACL) Overview. We recommend that you use the default policy, which is “bucket-owner-full-control”.
Credential Profile	Required. The profile name that is used to connect to S3 as defined in the **Certificates and Credentials** section of this page.
Dual Stack Networking	Optional. If unchecked, this setting indicates that this connection will consist of IPv4 addresses only. Enable this setting for IPv6 connections only.
Object Prefix	Optional. Embedded identifier that is useful for multiple Event Forwarders. The object prefix helps distinguish between output from multiple Event Forwarder instances, each of which has a distinct prefix.

HTTP

The HTTP destination type requires the following information:

Parameter	Description
Host Address	Required. URL of the HTTP destination endpoint.
Server Common Name	Optional. Common name (CN) of the destination server.
Send Timeout	Optional. Maximum duration of an upload connection. The default value is 60 seconds.
Upload Empty Files	Optional. Determines whether zero byte length files are uploaded. The default setting is false.
Max Bundle Size	Optional. The maximum bundle size (in bytes) to upload to the remote destination before compression is applied. The default value is 10MB.
Upload Template	Required. Template for formatting the output messages. The required template format to enter in this field is: `{"filename": "{{.FileName}}", "service": "carbonblack", "alerts":[{{range .Events}}{{.EventText}}{{end}}]}`.
Content-type HTTP Header	Required. The HTTP header helps the HTTP client consume the output properly. For example: `application/json, application/text, application/xml`. We recommend that you use the default value for this setting.
HTTP Authorization Token	Optional. Token to communicate with the remote destination.
OAUTH JWT - Client Email	Required if using OAUTH authentication. Oauth client identifier for communicating with the configured OAuth provider.
OAUTH JWT - Private Key	Required if using OAUTH authentication. PEM-encoded private key to sign the JWT payloads.
OAUTH JWT - Token URL	Required if using OAUTH authentication. The endpoint that is required to complete the 2-legged JWT flow.
OAUTH JWT - Private Key ID	Optional. A hint that indicates which key is being used.
OAUTH JWT - Permission Scopes	Optional. A comma-delimited list of requested permission scopes.
Send Events as Binary	Optional. If enabled, event JSON is sent in a byte array field instead of plain text.
Use Compression	Optional. If enabled, compresses the HTTP payload before uploading.
Certificates	Optional. Determines whether an uploaded certificate is required for the connection.

Syslog

The Syslog destination type requires the following information:

Parameter	Description
Syslog Destination	Required. The Syslog destination. The format is: `<protocol>://<fqdn>[:<port>]`
Server Common Name	Optional. Common name (CN) of the destination server.
Certificates	Optional. Determines whether an uploaded certificate is required for the connection.

Certificates and Credentials

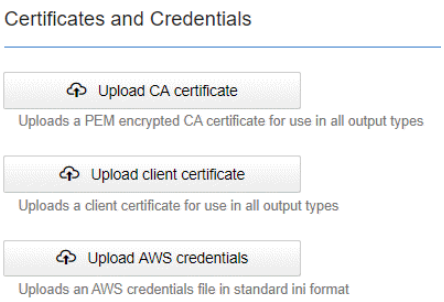

You can optionally upload certificates and AWS credentials to validate connections.

To upload a certificate:

1. Click the button for the certificate or credential to upload.
2. In the dialog box, specify the file to upload. Click **Upload**.

To use an uploaded certificate or credential, the corresponding option must be selected in the **Output** section (where it is applicable for the destination type).

After you have finished configuring the Event Forwarder, click **Save** at the top of the page. Saving the configuration immediately updates the configuration file on the server disk to reflect the changes. You must manually stop and restart the service.

Chapter 17

Creating and Using Investigations

This chapter describes how to work with investigations. Investigations allow you to group data for reporting, compliance, or retention purposes.

Sections

Topic	Page
Overview of Investigations	268
Creating Investigations	271
Creating Investigations	271
Adding Events to Investigations	271
Removing Events from Investigations	272
Adding Custom Events to Investigations	272
Deleting Investigations	272

Overview of Investigations

Investigations are collections of process events that share a common focus. They can include details and notes, and provide a way to group data for reporting, compliance, or retention purposes. Investigations are not particular to any user, so all investigations are available to each CB Response administrator.

It is a best practice to start an investigation whenever you begin any type of search — for example, after you discover a suspicious indicator and start searching for related process activity on your hosts.

You can create an investigation to keep an ongoing record of the scope and effects of the threat, so that you can stop it before it causes damage. There is no cost involved in creating an investigation, and if you tag process events during your search, you have a built-in record of the steps that provided the end result.

A default investigation comes with the server installation and is always available to collect any data that you tag. The default investigation cannot be deleted, so it is best used as a repository for data that interests you but does not warrant a dedicated investigation of its own.

The first time that you open the Investigations page, the default investigation appears.

Viewing Investigations

To view investigations:

- On the navigation bar, click **Investigations**.

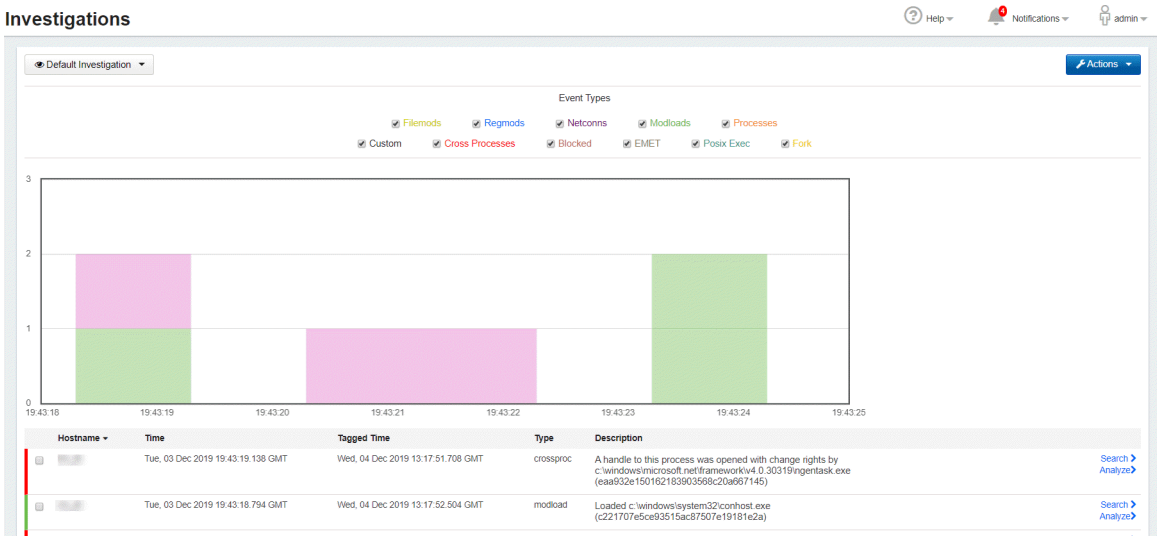

Investigations Menu Bar

The menu bar contains:

- A drop-down list of investigations. **Default Investigation** is the default.
- An **Actions** menu with the following options:

- **Remove Events** – See “Removing Events from Investigations” on page 272.
- **Add Custom Event** – See “Adding Custom Events to Investigations” on page 272.
- **Add Investigation** – See “Creating Investigations” on page 271.
- **Delete Investigation** – See “Deleting Investigations” on page 272.
- **Export timeline to PNG** – Exports data from the graph to a `.png` file and downloads it to your computer.
- **Export events to CSV** – Exports data to a .csv file and downloads it to your computer.

Event Types

Select and deselect checkboxes next to the event types to sort the events that display in the timeline and table. Only selected events will appear.

For detailed information on event types, see “Process Event Details” on page 193.

The following event types appear:

- **Filemods** – The number of files that were modified by process executions. Color-coded as yellow.
- **Regmods** – The number of Windows registry modifications that were made by processes executions. Color-coded as blue.
- **Netconns** – The number of network connections that process executions either attempted or established. Color-coded as purple.
- **Modloads** – The number of modules that were loaded by process executions. Color-coded as green.
- **Processes/Child Processes** – The number of child processes that were generated from process executions. Color-coded as orange.
- **Custom** – A custom event that you can create using the **Add Custom Event** option in the **Actions** menu. Color-coded as black.
- **Cross Processes** – (Windows only) A process that crosses the security boundary of another process. Color-coded as red.
- **Blocked** – Represents events that are related to the Ban Hash functionality. If an admin bans a hash and the sensor sees that binary and tries to stop it (already running) or prohibits it from running (blocks it), then the sensor generates a Blocked event. Color-coded as brown.
- **EMET** – Represents a specific type of event that deals with Microsoft’s Enhanced Mitigation Experience Toolkit (EMET) software. Color-coded as gray.
- **Posix_Exec** – (OS X and Linux only) The instance’s process that is loaded and the new binary image. Color-coded as green.
- **Fork** – (OS X and Linux only) The instance’s parent process, forked with a different Process ID (PID). Color-coded as yellow orange.

Bar Graph

The bar graph contains a timeline of the events that are tagged for the investigation. The events appear in color-coded bars (according to the event types). Events are stacked when they occur at the same time.

The color coding indicates which events happen at which times. Hovering over the color indicators on the timeline produces pop-up text, which explains what the block of color represents. For example:

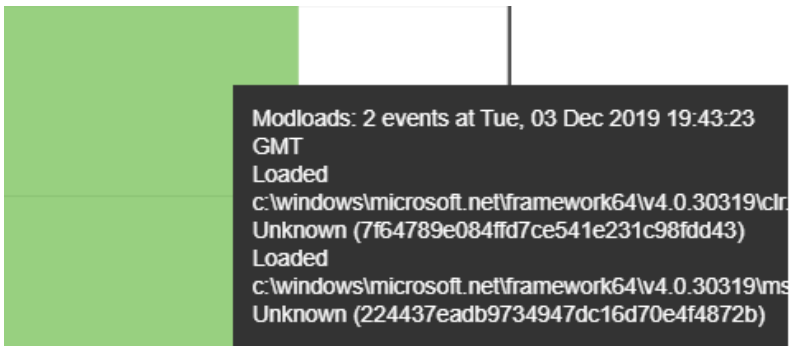

Events Table

The events table shows the events that are contained in the investigations. A colored bar appears on the left border of each row to indicates the event type.

Column	Description
Hostname	The name of the host on which the event occurred.
Time	The date and time that the event occurred.
Tagged Time	The time that the event was tagged for this investigation.
Type	The event type (filemod, regmod, netconn, modload, child process, fork, posix_exec, custom, crossproc, blocked, EMET).
Description	Description of the event; for example, paths to files and registry elements that were modified, signature status, and hash values.
Search	Opens the event in the Process Search page.
Analyze	Opens the event in the Process Analysis page.

Edit Event Description

When you hover over the description in a row, an **Edit** icon appears. Click the **Edit** icon to open the **Edit Event** window:

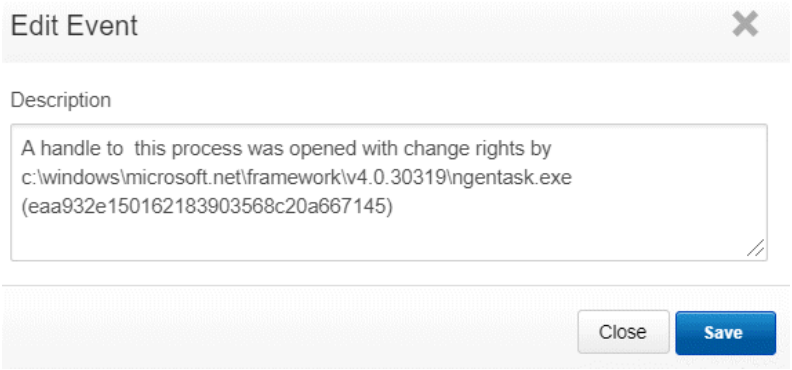

Use the **Edit Event** window to add context to the technical description or insights to share with the rest of your investigative team. Edits are visible within the investigation, but do not appear in the process execution data when viewed outside the Investigations page.

Child Processes

Rows that represent child processes contain a **Search** (magnifying glass) icon. This option displays a preview of the Process Analysis page for the child process.

Creating Investigations

You can create customized investigations and then add events to these, as discussed in “Adding Events to Investigations” on page 271.

To create an investigation:

1. On the navigation bar, click **Investigations**.
2. The **Default investigation** appears.
3. Click **Actions > Add Investigation**.
4. The **Add Investigation** window appears:
5. Enter a name for the investigation and click **Save**.

   Note: The name must be alpha-numeric. Special characters are not allowed.
6. The new investigation appears in the **Investigations** window, but is empty until you add events to it. See “Adding Events to Investigations” on page 271.

Adding Events to Investigations

To add events to investigations:

1. On the navigation bar, click **Process Search**.
2. Select a process from the **Process Search** page (see “Overview of Process Search” on page 173) and open the **Process Analysis** page (“Process Analysis Page” on page 184).

3. Scroll to the list of events. Click the **Tag** icon in an event row and select an investigation.

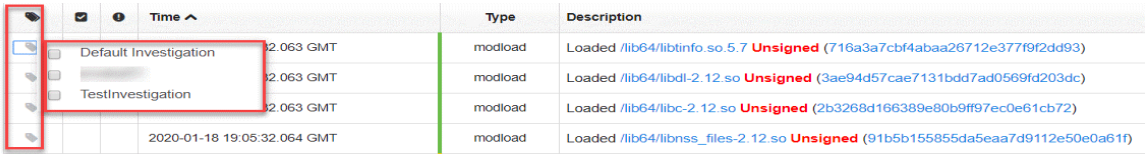

Removing Events from Investigations

When you remove an event from an investigation, it continues to exist in the system but is no longer included in the investigation.

To remove an event from an investigation:

1. On the navigation bar, click **Investigations**.
2. The **Default investigation** appears.
3. Use the **Investigation** drop-down menu in the top-left corner to open an investigation.
4. Click **Actions > Remove Events**.

Adding Custom Events to Investigations

You can create a custom event that you can use to:

- Add a new event type to the system.
- Add a note that displays on its own line in the rows at the bottom of the **Investigations** page.

You can also specify time parameters for the event, so that it appears in a designated space in the timeline.

To create a custom event:

1. On the navigation bar, click **Investigations**.
2. Click **Actions > Add Custom Event**.
3. In the **Description** field, type a description for the event.
4. In **Start Time**, enter the date and time for the event.
5. Click **Save**.

Deleting Investigations

When you delete an investigation, only the grouping, tagging, and edited descriptions are deleted. It has no other effect on the process executions that were a part of the investigation, or on how those processes appear in other pages.

To delete investigations:

1. On the navigation bar, click **Investigations**.
2. Select the investigation to delete.
3. Click **Actions > Delete Investigation**.
4. Click **OK** to confirm the removal.

Chapter 18

Watchlists

This chapter describes creating, using and managing watchlists. Watchlists are saved searches that run periodically against the process or binary data in CB Response. Watchlists are visible to all users.

Sections

Topic	Page
Overview	275
Viewing Watchlists and their Results	275
Built-in and Community Watchlists	277
Creating Watchlists	278
Managing Watchlists	281
Editing Watchlists	283
Deleting Watchlists	284

Overview

Watchlists are named process or binary searches that the server runs periodically (approximately every 10 minutes) without user action. When those saved searches produce new results, the server notifies users about them in a configurable way.

First responders can use the Watchlists page to quickly see items that are potentially interesting. For example, the **Newly Executed Applications** watchlist gives you rapid access to a list of the latest applications that were executed. If there are known recent issues with any new applications, you can quickly scan the results of that watchlist to find potential problems.

For watchlists that are based on threat intelligence feeds, you can factor scoring into a saved search. These watchlists tag a process or binary that is found on one of your endpoints when the score from a specified feed matches a specified score or falls within a specified score range. The score is the rating that is used to calculate the severity that is assigned to IOCs from a feed.

Additional information about enabling and using watchlists in specific contexts appears in the following chapters:

- Chapter 10, “Responding to Endpoint Incidents”
- Chapter 11, “Process Search and Analysis”
- Chapter 12, “Binary Search and Analysis”
- Chapter 15, “Threat Intelligence Feeds”
- Chapter 19, “Console and Email Alerts”

Viewing Watchlists and their Results

On the navigation bar, click **Watchlists** to open the Watchlists page. On this page, names of existing watchlists appear in a table on the left. Details and results for one watchlist (by default the first one in the table) appear on the right. You can display the details and results of a different watchlist by clicking its name.

The Watchlists Page

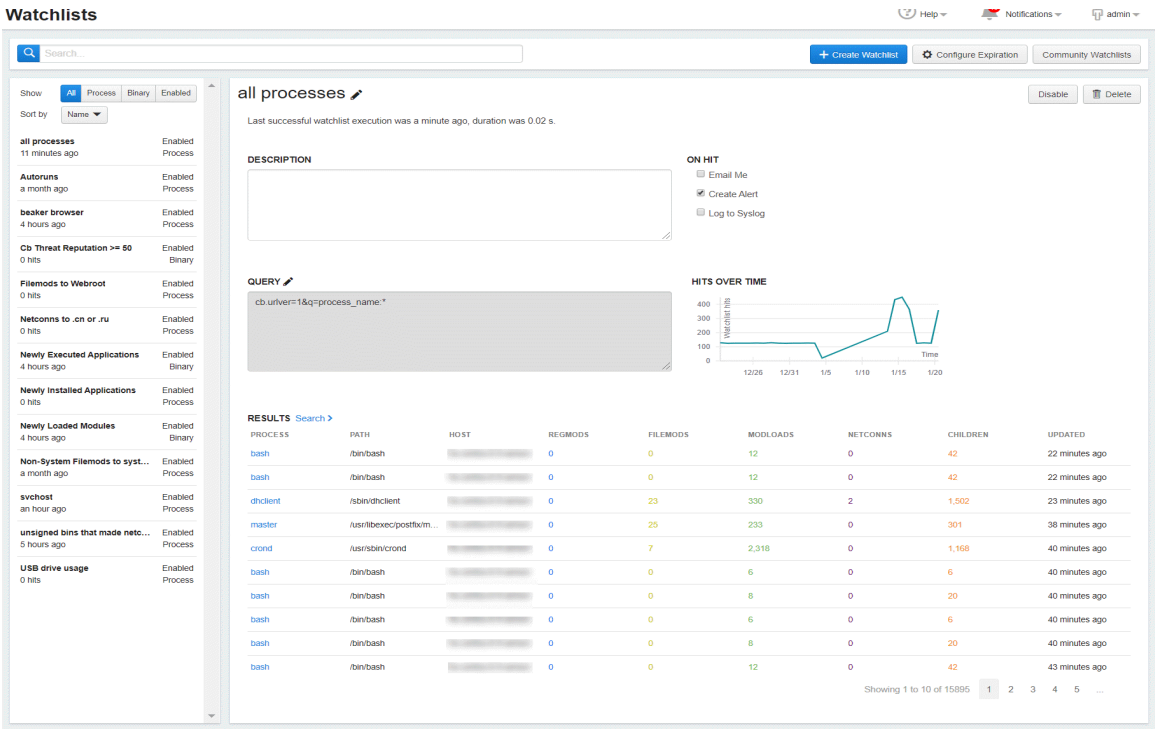

The left panel on the Watchlists page shows all available watchlists, their status and type, the number of hits, and either the time of their last run or another status message if they have not run recently. There are two tools for filtering these watchlists:

- At the top of the Watchlists page, use the **Search** box to search for watchlists by name.
- Immediately above the table of Watchlists on the left, filters and sorting controls can modify what is shown in the table. In the **Show** field, you can chose to show all watchlists, process watchlists only, binary watchlists only, or enabled watchlists only. In the **Sort by** field, you can sort by name, by the time the watchlist was created, by duration (how long it took the query to run), or by when each watchlist was most recently triggered. See “Managing Watchlists” on page 281 for how you can use these features to effectively manage watchlists.

The Watchlist Details Panel

The **Watchlist Details** panel on the right shows details for the currently selected watchlist. It includes the following information:

- Name and Description (if provided) of the watchlist
- If the most recent execution was successful, its time and duration; for unsuccessful executions, this line shows either timeout or error information. Typically watchlists are scheduled to run every 10 minutes, but if a previous watchlist session is still running, the next one will be delayed and try to start periodically (every 10 minutes).
- Query used to match events to the watchlist.

- The **On Hit** settings determine how (or if) you are notified when an event matches the query.
- A graph that shows the number of hits on this watchlist over time.
- Table of results showing details for each hit.

> **Note**
>
> For each watchlist that is run, the number of matching events that are tagged is limited to 100, even if more events actually match the watchlist. This limit prevents performance issues and eliminates the potential for excessive numbers of notification emails that are unlikely to add useful information.

Click the **Search** link to show query results in the context of the Process Search page or the Binary Search page.

The **Watchlist Details** panel also provides buttons in the top right to disable or delete the watchlist. When you click the **Disable** button, the watchlist is disabled and no longer runs. New results that match the search query do not result in any notification or record that they triggered a hit for the watchlist.

When you click the **Delete** button and confirm the deletion, the watchlist is permanently removed.

Built-in and Community Watchlists

Although you can create your own watchlists, CB Response provides access to two sources of pre-configured watchlists:

- The Watchlists page in the console includes a list of default watchlists.
- The Carbon Black User Exchange provides a forum for sharing watchlists.

The default watchlists on the Watchlists page include the following:

- **Autoruns**
- **Filemods to Webroot**
- **Netconns to .cn or .ru**
- **Newly Executed Applications**
- **Newly Installed Applications**
- **Newly Loaded Modules**
- **Non-System Filemods to system**
- **USB drive usage**
- **CB Threat Reputation**

In addition to the built-in Watchlists, in the top-right corner of the **Watchlists** page, you can click the **Community Watchlists** button to access Threat Research on the Carbon Black User Exchange.

Threat Research is a central portal where watchlist users can publish and discuss watchlists that might eventually be included as a feed from CB Threat Intel.

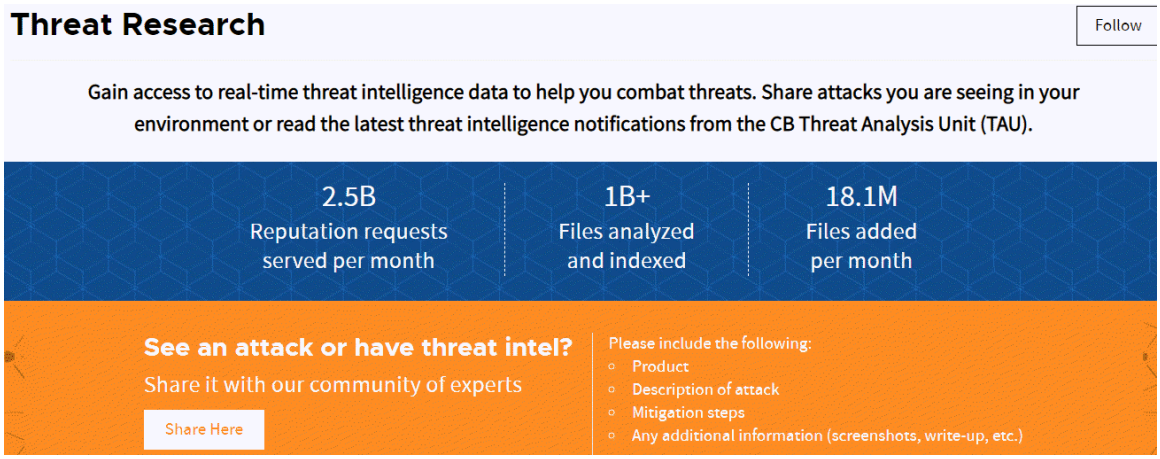

Creating Watchlists

You can create your own customized watchlists from the Watchlists, Process Search, Binary Search, or Threat Intelligence Feeds pages. The information you must provide varies, depending on the location from where you start.

To create watchlists from Process Search or Binary Search pages:

1. On the navigation bar, select either **Process Search or Binary Search**.
2. Enter the query for the processes or binaries for which to create a watchlist. The syntax should match a search box query in the Process or Binary Search pages.

   Caution: As with searches outside of a watchlist, use of leading wildcards is discouraged because of performance issues.

   You cannot edit several aspects of a watchlist search query, so examine the results carefully before proceeding. For more information on editing queries, see “Editing Watchlists” on page 283.

   For more information on performing searches, see:

   - Chapter 11, “Process Search and Analysis,”
   - Chapter 12, “Binary Search and Analysis,”
   - Chapter 13, “Advanced Search Queries.”

   If you are using multiple MD5 or SHA-256 hash values for search criteria to create a watchlist, you must enclose the values in parentheses ().

   For example:

   ```
   (md5:45cc061d9581e52f008e90e81da2cfd9
   md5:829e4805b0e12b383ee09abdc9e2dc3c
   md5:ac9fa2ba34225342a8897930503ae12f
   md5:5f7eaaf5d10e2a715d5e305ac992b2a7)
   ```

   If you do not enclose the list in parentheses, the only value that is tagged for the watchlist is the last value in the list.
3. On the Process Search page, click **Create Watchlist** or, on the Binary Search page, click **Create Watchlist** from the **Action** menu. See “To create watchlists from the Watchlists page:” on page 279 for details on filling out the form.

To create watchlists from the Watchlists page:

1. On the navigation bar, click **Watchlists**.
2. Click the **Create Watchlist** button.

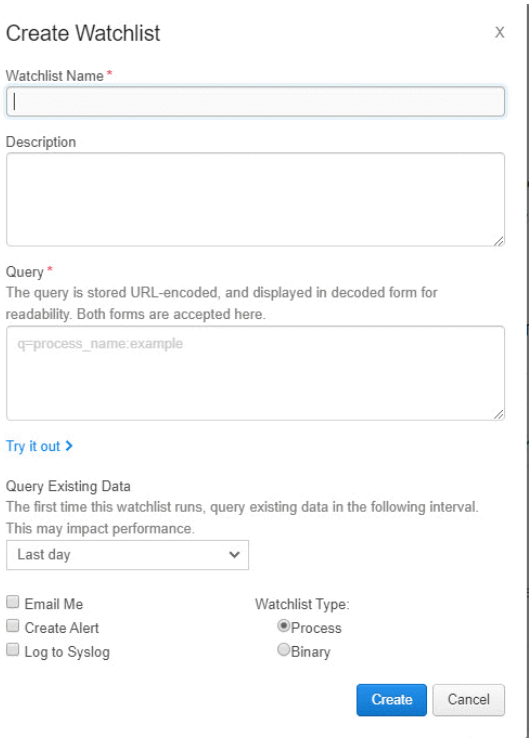

3. **Watchlist Name:** Enter a meaningful name for the watchlist.
4. **Description** (optional): Provide the purpose of the watchlist.
5. **Query**: The query that is currently open, if any.
6. **Query Existing Data** dropdown menu: Define the time period for which existing data is queried on the first run of the watchlist. The longer the timeframe that is selected, the longer it will take the query to run directly after this watchlist is created. A longer time can also stress other product services, such as process search, while the watchlist is running. After the watchlist has run one time, it will run on new data in 10 minute intervals thereafter.
7. **Email Me**: Check the checkbox to receive email notifications for matching hits.
8. **Create Alert**: Check the checkbox to send an alert when conditions matching the watchlist occur. Triggered alerts are reported in the Alert Dashboard page and the Triage Alerts page. For more information on alerts, see Chapter 19, “Console and Email Alerts.”
9. **Log to Syslog**: Check the checkbox to log all hits `syslog`. Syslogs are written to `/var/log/cb/notifications/`. In this case, the log filenames have the format `cb-notifications-<watchlist ID>.log`.
10. **Watchlist Type**: Identify the type as **Process** or **Binary**.
11. Click **Create**.

To create watchlists from the Threat Intelligence Feeds page:

1. On the navigation bar, click **Threat Intelligence**.
2. Select the feed for which to create a watchlist.
3. From the **Actions** menu, click **Create Watchlist**:

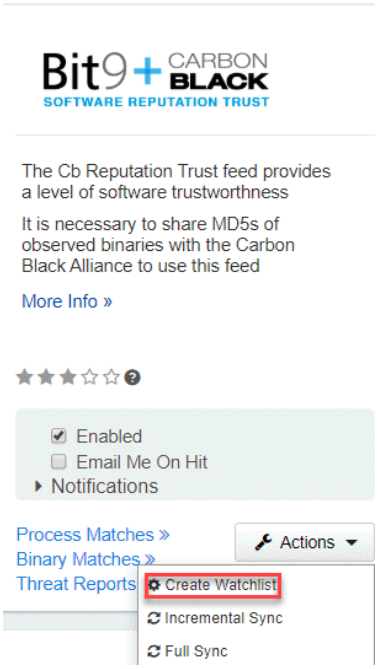

4. The **Add Watchlist** window opens.

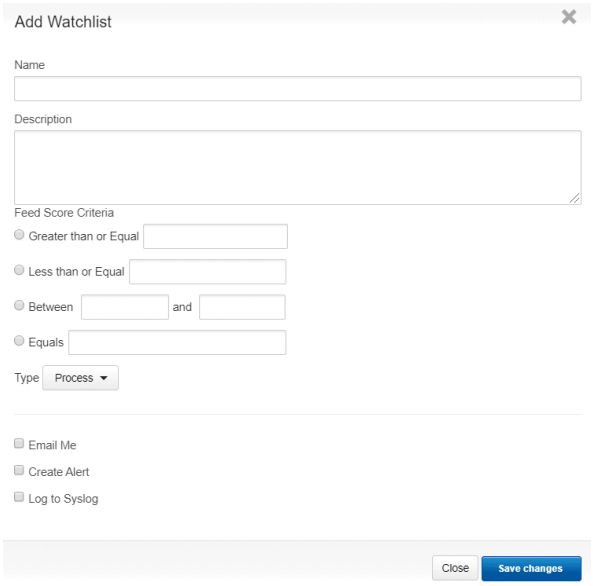

5. In the **Name** field, enter a meaningful name for the watchlist.
6. (Optional) Provide a description to give additional details about the watchlist, such as its purpose.
7. In the **Feed Score Criteria** section, use the various fields to enter the score criteria for the severity of IOCs to track.

8. On the **Type** dropdown menu, click **Process** or **Binary**.
9. Select the **Email Me** check box to receive email notifications when hits occur.
10. Select the **Create Alert** check box to send an alert when conditions matching the watchlist occur. Triggered alerts are reported in the Alert Dashboard page and the Triage Alerts page. For more information on alerts, see Chapter 19, “Console and Email Alerts.”
11. Select the **Log to Syslog** check box to log all matching hits to `syslog`. Syslogs are written to `/var/log/cb/notifications/`. For watchlists, log filenames have the form `cb-notifications-<watchlist ID>.log`.
12. Click **Save changes**.

Managing Watchlists

Watchlists can provide you with valuable information about conditions that matter in your environment. You might need to fine-tune watchlists for your environment, based on their performance and the quality of the information they provide.

You can monitor the status of a watchlist to see whether and when it has executed, and whether there are any error conditions associated with the watchlist. If you find that the watchlist is not performing as expected, you can edit, disable, or delete it.

Watchlist Status

Watchlists show the following status in the table view:

- **Queued** – This appears when a watchlist was recently created and is waiting to be executed.
- **Timeout** – This appears when a watchlist does not execute successfully (or generate an error) after two minutes. A timed-out watchlist will be re-tried, but will only be run on events that appeared between its failed execution and the retry time. See “Slow or Error-producing Watchlists” for more information.
- **Expired** – This appears when the watchlist has not had any hits in the specified period. See “Watchlist Expiration” for more information.
- **Error** – This appears when an error happens during watchlist execution and indicates that the watchlist did not execute successfully. See “Slow or Error-producing Watchlists” for more information. If you are unable to resolve an error condition, consider contacting Carbon Black Support.

In the **Watchlist Details** panel, descriptive messages appear if the last execution of the watchlist resulted in an error or a timeout. For successful executions, the details panel shows the following:

- **Last execution** – The time of the last successful execution.
- **Duration** – The duration required to complete execution. See “Slow or Error-producing Watchlists”.

Watchlist Expiration

You can configure CB Response to notify you when a watchlist has not received any hits over a specified period of time. This might be a sign that the watchlist is not useful and can be deleted, or perhaps that the query in the watchlist must be modified to be effective. Watchlist expiration is informational only. When a watchlist expires, you are prompted to take action on it, but it is still fully functional unless you delete or disable it.

A single watchlist expiration configuration applies to all watchlists.

To configure watchlist expiration:

1. On the navigation bar, click **Watchlists**.
2. Click **Configure Expiration**.

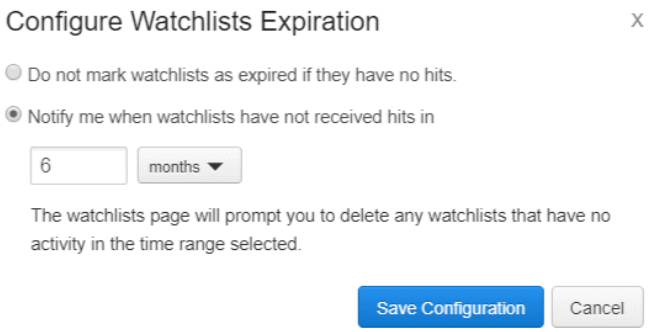

3. By default, watchlists are marked as “Expired” if they have received no hits over a six-month period. To use a different time period for expiration or to reconfigure a watchlist page that had expiration turned off:
   a. Make sure the **Notify me when watchlists have not received hits in** radio button is selected.
   b. Enter a number in the box and choose the units (days, months, or years) from the menu.
4. If you do not want any watchlists to be designated as expired, click the radio button that reads **Do not mark watchlists as expired if they have no hits**.
5. Click **Save Configuration**.

Slow or Error-producing Watchlists

Temporary conditions might cause a watchlist to timeout or fail with an error message. However, if a watchlist continues to fail, you might need to investigate it and consider modifying the query or deleting the watchlist.

You can identify slow or error-producing watchlists on the watchlist table by using the **Duration** choice on the **Sort by** menu.

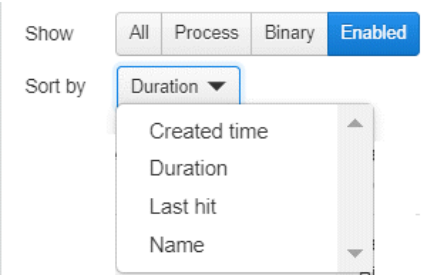

This produces the following results:

- Watchlists that have not executed successfully, including disabled, queued, errored out or timed out watchlists, appear first. Because you are not usually interested in disabled watchlists, consider clicking the **Enabled** tab to eliminate disabled watchlists from your results.
- After the non-executed watchlists, watchlists that have been executed successfully are listed, beginning with the slowest (longest duration) watchlist and then in descending order of duration.

Duration, timeout and error status is also displayed underneath the watchlist name in the **Watchlist Details** panel.

After you identify a problematic watchlist, you can examine its **Query** field or **Feed Score Criteria** to see whether there are any obvious issues, such as leading wildcards in the query. Chapter 13, “Advanced Search Queries,” includes guidelines for creating queries, including query usage that could cause difficulties.

If you are unable to modify a watchlist in a way that produces efficient, successful performance, you can contact Carbon Black Support for further troubleshooting.

Editing Watchlists

You can edit a watchlist in the **Watchlist Details** panel of the **Watchlists** page. For most changes, the underlying ID that uniquely identifies the watchlist remains the same. However, if you edit the watchlist search query, it effectively becomes a new watchlist.

To edit watchlists:

1. On the navigation bar, click **Watchlists**.
2. In the left panel, select the watchlist to edit. Its details appear in the right panel.

3. You can edit the following attributes of the watchlist:
   a. To change the name of the watchlist, click the pencil icon next to the name at the top of the page.
   b. To edit the watchlist query, click the pencil icon for the **Query** box. In the **Edit Watchlist Query** dialog box, modify the query and then click **Save Changes**.

      Note: Saving a modified watchlist query overwrites the watchlist ID even if the watchlist name is the same. Therefore, any references to the older version of the watchlist, such as in alerts or through the API, are no longer connected.
   c. To disable the watchlist, click **Disable**. To enable it, click **Enable**.
   d. To receive email notifications when there are hits that match your search, select **Email Me**. Deselect the checkbox to stop receiving email notifications.
   e. To send an alert when conditions matching the watchlist occur, select **Create Alert**. Deselect the checkbox to stop sending alerts.
   f. To log all hits that match the search to `syslog`, select **Log to Syslog**. Syslogs are written to `/var/log/cb/notifications/`. In this case, the log filenames have the form `cb-notifications-<watchlist ID>.log`.

Deleting Watchlists

You delete watchlists using controls on the details panel for that watchlist on the **Watchlists** page.

To delete a watchlist:

1. On the navigation bar, click **Watchlists**.
2. In the left panel, select the watchlist to delete. Its details appear in the right panel.
3. In the top-right corner, click **Delete** and click **OK** to confirm the deletion.

Chapter 19

Console and Email Alerts

This chapter explains how to create and manage CB Response alerts on the console. Alerts can be triggered based on watchlist or CB Threat Intel feed events. This chapter also explains how to enable email alerts to report these events.

Sections

Topic	Page
Overview of Alerts	286
Enabling Console Alerts	286
Viewing Alert Activity on the HUD Page	287
Managing Alerts on the Triage Alerts Page	289
Enabling Email Alerts	294

Overview of Alerts

You can create alerts to indicate in the CB Response console when suspicious or malicious activity appears on your endpoints. Alerts are available for two types of events:

- **Watchlist hits** – Watchlists can be configured to send an alert when conditions matching the watchlist occur. See Chapter 18, “Watchlists,”.
- **Threat intelligence feed hits** – Threat intelligence feeds can be configured to send an alert when that feed reports an IOC. See Chapter 15, “Threat Intelligence Feeds,”.

Triggered alerts are reported in two locations in the CB Response console:

- The HUD page contains a summary that shows the number of unresolved alerts, the number of hosts that have unresolved alerts, and other alert-related data, including the alerts for each host. See “Viewing Alert Activity on the HUD Page” on page 287.
- The **Triage Alerts** page contains more details about triggered alerts and provides a filter and search interface to find alerts that match different criteria. It also allows you to manage the alert workflow, marking the status of each alert from its initial triggering to its resolution. See “Managing Alerts on the Triage Alerts Page” on page 289.

You can configure watchlists and threat intelligence feeds to send email alerts when there is a “hit” on data from a CB Response sensor that matches the watchlist or feed. You can enable email alerts in addition to or instead of the CB Response console-based alerts. See “Enabling Email Alerts” on page 294 for more information.

Enabling Console Alerts

You can enable alerts for any watchlist or threat intelligence feed. Consider how many hits you are likely to receive when you enable alerts. Some watchlists or feeds might generate too many hits to be useful, making it more difficult to identify significant alerts. Ideally, an alert should get your attention for issues that you need to follow up on. No alerts are enabled by default.

Watchlist Alerts

Watchlists are user-created, custom, saved searches that are based on process search, binary search, or feed results. You can use watchlists to monitor endpoints for detected IOCs. You can also select the most important watchlists to monitor and add console alerts. Then, you can then view and manage these key watchlist and feed hits in the Triage Alerts page.

To enable console alerts for a watchlist:

1. On the navigation bar, click **Watchlists**.
2. In the left panel of the **Watchlists** page, select the watchlist for which to create an alert. Use the **Search** box at the top of the panel to locate a watchlist that does not immediately display.
3. In the right panel, click the **Enable b**utton if the watchlist is disabled, and select the **Create Alert** check box.

   The watchlist will begin generating alerts.

Threat Intelligence Feed Alerts

Threat intelligence feeds provide information that helps you identify malware and its sources. CB Response integrates with third-party and internal feeds (such as the CB Threat Intel Reputation and CB Response Tamper Detection) that identify hosts on which tamper attempts have occurred.

Adding a CB Response console alert to a feed allows you to highlight hits matching reported malware from a specific source. You can then view and manage high-importance feed and watchlist hits in one place, on the **Triage Alerts** page.

Make sure you understand the volume of reports that you will receive from any feed before enabling alerts for it. Among other things, read the description of a feed on the **Threat Intelligence Feeds** page. Some feeds include a specific recommendation *not* to enable alerts, because of the report volume or percentage of false positives that can occur.

To enable console alerts for a threat intelligence feed:

1. On the navigation bar, click **Threat Intelligence**.
2. Select the **Notifications** dropdown menu and select the **Create Alert** check box for each feed panel to enable console alerts.

To disable console alerts for a threat intelligence feed:

1. On the navigation bar, click **Threat Intelligence**.
2. Select the **Notifications** dropdown menu, and uncheck the **Create Alert** check box for each feed panel to disable console alerts.

Viewing Alert Activity on the HUD Page

The HUD page is a customizable page that provides a summary of alerts on hosts that report to your CB Response server. See “Using the Head-Up Display Page” on page 296.

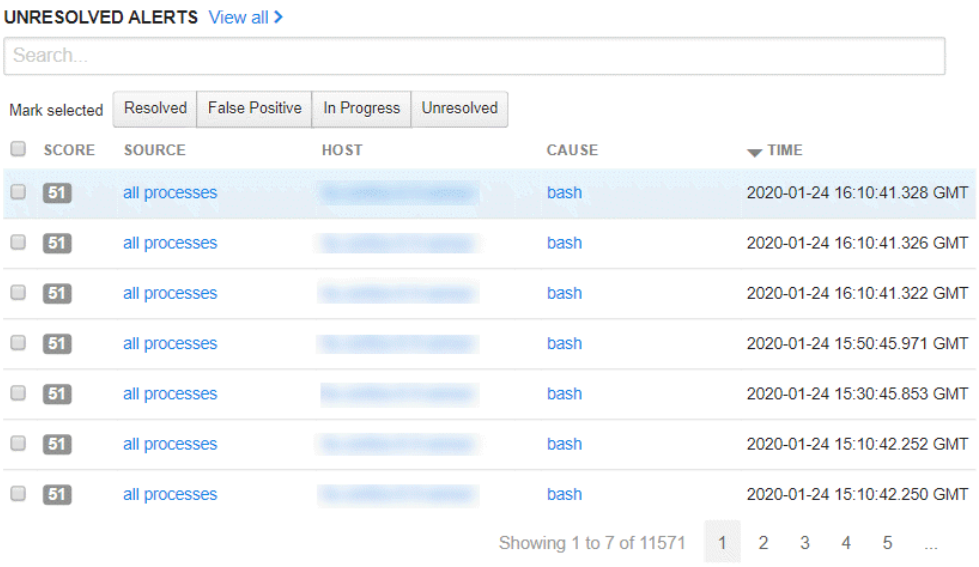

By default, the **Unresolved Alerts** panel displays all unresolved alerts for a sensor. You can also display resolved, false positive, and in- progress alerts by clicking a button at the top of the **Unresolved Alerts** panel:

- **Resolved**
- **False Positive**
- **In Progress**
- **Unresolved**

Note

You can enlarge the **Unresolved Alerts** panel to display more details by holding your left mouse button down on the bottom-right expansion icon and dragging the panel to the desired size.

The **Unresolved Alerts** panel contains these columns:

Note

Some columns in this panel are sortable, such as the **Score** and **Time** columns. You can determine if columns are sortable by hovering your cursor over the column name; sortable column names will turn black and your cursor will change to a hand icon. An arrow appears, indicating the sort direction (ascending/descending).

Pane	Description
Score	Displays the alert severity, where 100 is a severe threat and 1 is not a threat.
Source	Displays the feed that is associated with the alert, such as threat intelligence and watchlist feeds. Clicking a link in this column opens the associated page.
Host	Displays the host that is associated with the alert. Clicking a link in this column opens the Sensors page.
Cause	When the alert is caused by a binary, this column displays the binary's MD5 hash. Clicking on this link takes you to the Search Binaries page. When the alert is caused by a process, this column displays the process name. Clicking on this link takes you to the Search Processes page.
Time	Displays the time when the alert occurred.

The **Unresolved Alerts** panel also contains a **View all** link in the top-right corner. Clicking this link displays the Triage Alerts page.

Managing Alerts on the Triage Alerts Page

When an alert is received that indicates suspicious or malicious activity, incident responders must:

- Determine the seriousness of the alert.
- Determine whether the alert indicates a sufficiently severe threat.
- Find a way to resolve a serious threat.

These tasks might involve using the following CB Response features:

- Endpoint Isolation
- Live Response
- Banning

It might also require using other tools.

Given the high volume of threat reports, it is critical to prioritize, investigate, and keep track of alert statuses. After an alert is resolved, it should be removed from the list of threats requiring attention so that ongoing threats can be addressed.

The Triage Alerts page provides features for alert management. It includes search and filtering capabilities for locating specific alerts or alert types. It also allows you change alert status.

To open the Triage Alerts page:

On the navigation bar, click **Triage Alerts**.

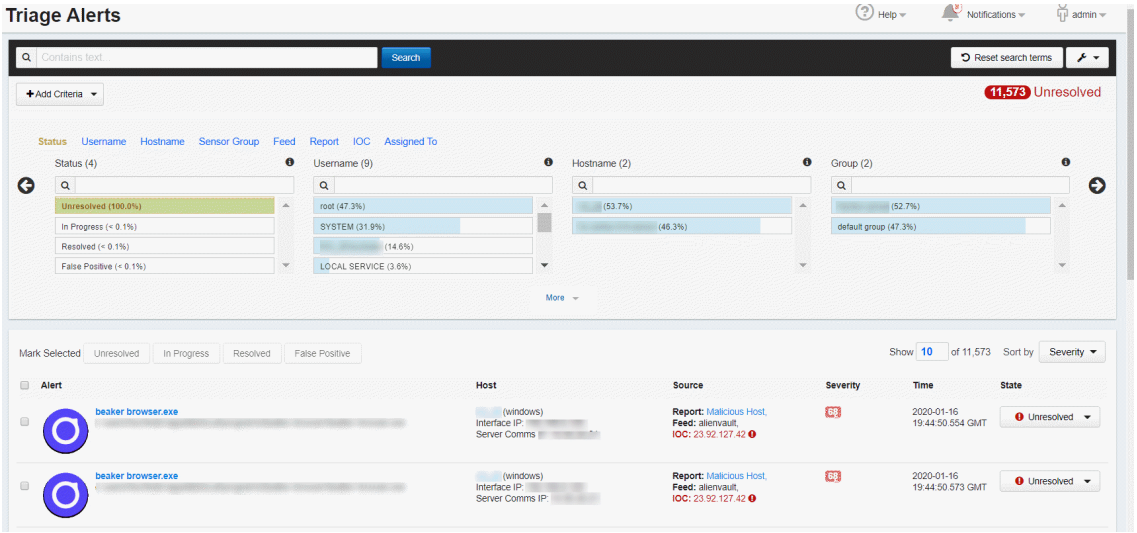

> **Note**
>
> You also can navigate to the Triage Alerts page from the HUD page by clicking **View all** in the **Unresolved Alerts** panel. See "Viewing Alert Activity on the HUD Page" on page 287.

The Triage Alerts page is divided into three major sections:

- The top section includes the **Search** field and button, **Add Criteria** button, **Reset search items** button, and **Actions** menu.
- The middle section contains filters that are category-specific lists (**Status**, **Username**, and so on). These filters show the percentage of alerts that match different values in each category, and allow you to filter the view to show alerts that match values.
- The bottom section contains the **Alerts** table, which contains details for alerts that match the search criteria that is entered in the first two sections.

Displaying the Report Name

The middle section of the Triage Alerts page lets you filter by various criteria, including **Reports**. By default, the **Reports** display shows the report ID (for example, dbe2eab5-3829-45df-b6c4-3dfb7a215d69). You can change the display to show the report name (for example, “PowerShell executed with encoded instructions”).

To change the display, you must change a setting in the `cb.conf` file. The default value of this setting is `False`: use this feature with caution because additional memory will be used in proportion to the number of reports on your server.

To enable the display of report names:

1. On the CB Response server, open `/etc/cb/cb.conf` for editing.
2. Set `FeedHitLoadReportTitles=True.`
3. Set the number of characters (from -1 to 80) for the report name in the `FeedHitMaxReportTitleLength` field. The default (and maximum) number of characters is 80. A value of -1 keeps the report name from being truncated in bus events, syslog, and email notifications.

```
FeedHitLoadReportTitles=True
FeedHitMaxReportTitleLength=80
```

4. Restart cb-enterprise services.

After you have changed the `cb.conf` setting and restarted cb-enterprise services, the report names are populated in the following places:

- In the Triage Alerts page **Records** facet.
- Bus events.
- Syslog notifications.
- Email notifications. Both report ID and report name are displayed in the email. If the feature is turned off, the report name is displayed as “Unknown”.

Reviewing Alerts

Each row in the **Alerts** table shows the description and data for an individual alert. The description and data that appears can vary depending on a variety of factors, including:

- The source and type of the alert.
- Whether the binary for a process has been signed.
- Whether a binary is “Trusted” by the CB Response Alliance.

The **Alerts** table has several tools for adjusting the table display:

- **Sort order** – You can sort the **Alerts** table using the **Sort By** button in the top-right corner of the **Alerts** table. You can sort by:
  - **Severity** (default)
  - **Most Recent**
  - **Least Recent**
  - **Alert Name Ascending**
  - **Alert Name Descending**
- **Page navigator** – You can use the page navigation bar in the bottom-right corner of the **Alerts** table to move between pages in tables views that do not fit on a single page.

Alerts Table Data

The row for each alert in the table shows the following columns:

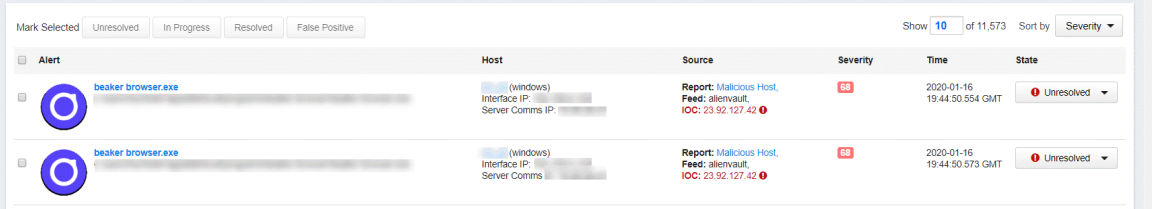

- **Alert** – Contains the following details:
  - An icon that represents the process or binary that caused the threat alert, if available. If there is no special icon for this binary, a generic file icon is used **Note:** Tamper alerts show what feed is triggered; the icon is of the host type.
  - The directory path where the process or binary is installed.
  - If this is a binary, the blue process link takes you to the Binary Details page. If this is a process, the blue process link takes you to the Process Analysis page.
- **Host** – Shows host details and a link to the Sensor Details page.
- **Source** – The watchlist or feed that triggered the alert with a link to that watchlist or feed.
- **Severity** – The severity score of the alert that CB Response produces based on underlying alert data. Click the **Severity** number to show additional details by which the severity is calculated:
  - **Feed rating**
  - **Report scores**
  - **Confidence**
  - **Criticality**

  The **Severity** numbers are color-coded, with red being severe threats and green being low threats. A score of 100 represents the most severe alerts.
- **Time** – The time when the alert was triggered.
- **State** – The alert state, which includes:

- **Mark as Resolved** – Select this when the alert is resolved.
- **Mark as Unresolved** – By default, only unresolved alerts are displayed.
- **Mark as In-Progress** – Select this for alerts that have resolutions in progress.
- **Mark as False Positive** – Select this for alerts that were not true threats.

See "Managing Alert Status" on page 292.

Managing Alert Status

You can change the status of individual alerts or all alerts in the current view. Changing alert status is strictly for alert management purposes. It helps you organize alerts that need attention, are being investigated, have been resolved, or are false positives.

Change an alert status to indicate what you are doing or have done based on your review of an alert. An alert status has no effect on the actual issue that caused the alert.

In the **Alerts** table, the far-right column includes an icon representing the current alert status and a drop-down list for changing that status.

The following table describes the alert status options and shows their icons.

To change the status of all alerts matching a search and/or filter:

1. On the navigation bar, click **Triage Alerts**.
2. On the Triage Alerts page, enter the search string and/or filter criteria for the alerts that have a status to change.
3. From the **Actions** menu (located in the top-right corner of the **Triage Alerts** page), select the **Mark all** menu option for the status to assign.
4. Click **OK** in the confirmation window to change the status of all of the alerts on the page.

Note

When using the **Mark all** commands, be sure that you want to change all of the alerts matching the current filter and search, including those on pages that are not displayed. After you change the status, there is no "undo" command. Be especially careful about changing alert statuses when the view is unfiltered (showing all alerts).

To change the status of a single alert:

1. On the navigation bar, click Triage Alerts.
2. In the **Alerts** table, select the check box to the left of the alert that has a status that you want to change.
3. From the **Actions** drop-down list (located in the top-right corner of the **Triage Alerts** page), select the appropriate option for the status you want to assign.
4. Click **OK** in the confirmation window to change the status of the selected alert.

Note

Changed alerts will disappear from the current view if you have filtered the page for a different status.

Ignoring Future Events for False Positive Alerts

Feeds use a variety of criteria to determine if a file or event is a threat, and you might not agree with all of the alerts that are generated by certain feeds. When you review alerts and determine that an alert is not reporting an actual threat, you can mark that alert as a "false positive", so you can eliminate it from the list of alerts that require your attention.

CB Response also provides a feature that lets you ignore future instances of a false positive alert from a threat feed. You can choose to ignore an individual alert, or specify that all alerts matching your search criteria should be ignored in the future.

> **Note**
>
> Only threat feed alerts can be designated as alerts to ignore. Alerts from watchlist matches are always triggered, since watchlists are assumed to use criteria that your CB Response users specifically chooses.

To ignore the triggering event for an alert:

1. On the navigation bar, click **Triage Alerts**.
2. In the **Alerts** table, select the check box to the left of the alert.
3. In the **Status** column, select **Mark as False Positive** in the dropdown menu.

4. In the **Mark as Resolved False Positive** dialog, you can ignore future events from this report by moving the slider button to **Yes**. Click **Resolve**.

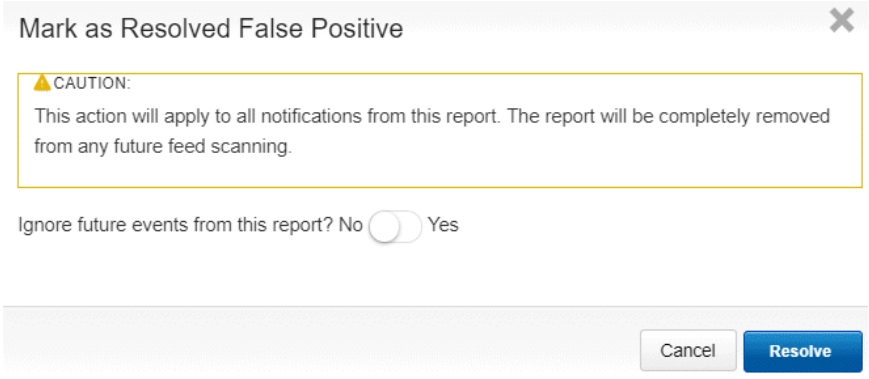

Marking events from multiple alerts to be ignored involves searching for the alerts you want to ignore, confirming that the results are what you expect, and then making a bulk resolution.

Enabling Email Alerts

You can enable email alerts to report events that trigger watchlist and threat intelligence feed alerts. This feature informs you of events of interest, even when you are not logged into the CB Response console. You can then go to the console to investigate and resolve the alert. The email alerts feature is enabled on a per-console user basis.

Configuring an Email Server

Before enabling email alerts for specific watchlists or feeds, you should decide which email server to use. You can:

- Use your own email server.
- Use the CB Response external email server.
- Opt out of email alerts.

If you use the CB Response external email server, the following information is sent through the server and stored by CB Response:

- Your server ID
- The time of the email
- The name of the watchlist or feed are that triggered the hit

> **Important**
>
> Carbon Black strongly recommends that you use your own email server, because email sent through the CB Response external email server is sent over the Internet in clear text.

To configure an email server for alerts:

1. Log in to CB Response as a Global Administrator (for on-premises installations) or an Administrator (for the cloud).
2. In the top-right corner of the CB Response console, click ***Username*** **> Settings**.
3. On the **Settings** page, click **E-Mail**.

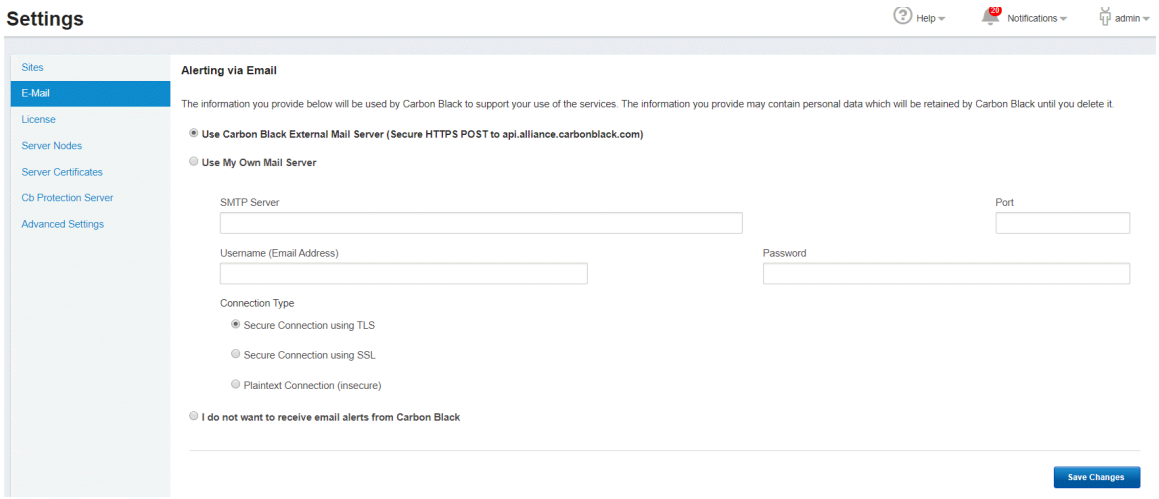

All email alerts for all console users will use these settings.

Enabling Specific Email Alerts

After you have configured an email server, any watchlist or feed can be configured to send email alerts when it gets a hit on a CB Response sensor.

You can turn on/off email alerts for individual watchlists and feeds as needed (for example, if you find that a watchlist or feed is creating too much email traffic). Email alerts for any specific watchlist or feed are enabled on a per-user basis.

> **Note**
>
> If you have upgraded from a previous release of CB Response, any email alerts that you had enabled for watchlists and threat intelligence feeds remain enabled after the upgrade.

To enable email alerts for a watchlist:

1. On the navigation bar, click **Watchlists**.
2. Select the watchlist for which to enable email alerts. If the watchlist name is not visible or you are not sure what the name is, use the **Search** field.
3. Confirm that the watchlist is enabled.
4. Check the **Email Me** check box.

To enable email alerts for a threat intelligence feed:

1. On the navigation bar, click **Threat Intelligence**.
2. To activate email alerts for a feed, select the **Email Me on Hit** check box.

Chapter 20

Using the Head-Up Display Page

This chapter explains how to use the Head-Up Display (HUD) page, which is a customizable dashboard in the CB Response console.

Sections

Topic	Page
Overview of HUD	297
Endpoint Hygiene Panel	298
Event Monitor Panel	298
Query Duration Panel	299
Resolution Time Panel	300
Saved Searches Panel	300
Sensors Panel	300
Unresolved Alerts Panel	302

Overview of HUD

The HUD page is a customizable page that provides a summary of alerts on hosts that report to your CB Response server. It provides a quick reference view that includes details on the following:

- Endpoint hygiene
- Event monitoring
- Query duration time
- Resolution time
- Saved searches
- Sensors involved in alert activity
- Unresolved alerts

Viewing the HUD Page

To view the HUD page:

- In the navigation bar, click the logo at the top left corner of the console.

Customizing the HUD Page

To reposition HUD panels:

- Hold down your left mouse button on a panel and drag the panel to the desired location.

To resize HUD panels:

- Hold down your left mouse button on the bottom-right resizing icon and drag the panel to the desired size. Note that larger panels display more details.

As you customize the HUD layout, the layout is automatically saved for future use per device.

Sortable Columns

Some columns are sortable. You can determine if a column is sortable by hovering your cursor over the column name. Sortable column names change colors and your cursor changes to a hand icon. An arrow appears that indicates the sort direction (ascending or descending).

The following sections describe **HUD** page panels. However, the **Unresolved Alerts** panel is described in “Viewing Alert Activity on the HUD Page” on page 287.

Endpoint Hygiene Panel

The **Endpoint Hygiene** panel shows the daily percentage of hosts that are reporting suspect processes over a 30-day period. This percentage is based on two values that CB Response records:

- The total number of active hosts in the network.
- The number of hosts that have one or more bad processes.

The following image shows zero endpoint activity, which is desirable.

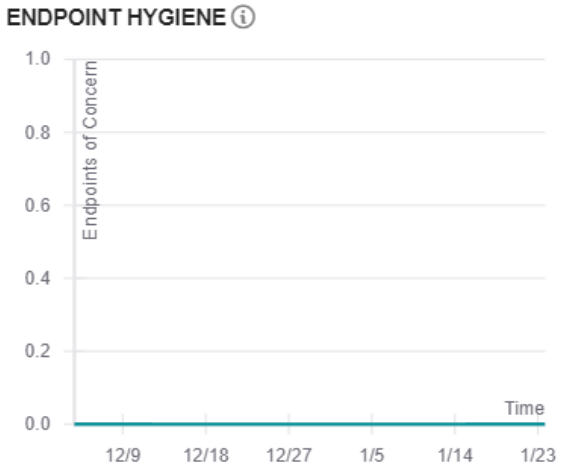

Event Monitor Panel

The **Event Monitor** panel provides a live feed of event activity. It updates every five seconds.

- Vertical bars indicate alert activity, such as resolving an alert or incoming alerts.
- Horizontal lines indicate watchlist activity.

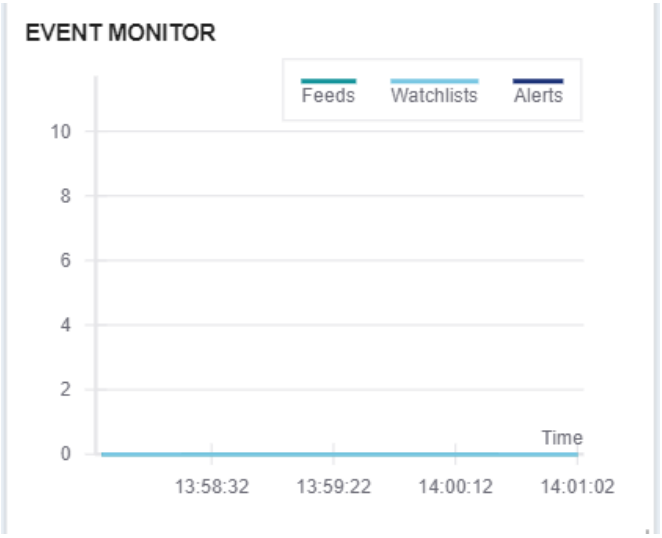

Query Duration Panel

Queries that take longer than a second to complete are presented in this panel. At a glance, you can see which queries are taking a long time to complete, and take action to improve query structures and efficiency.

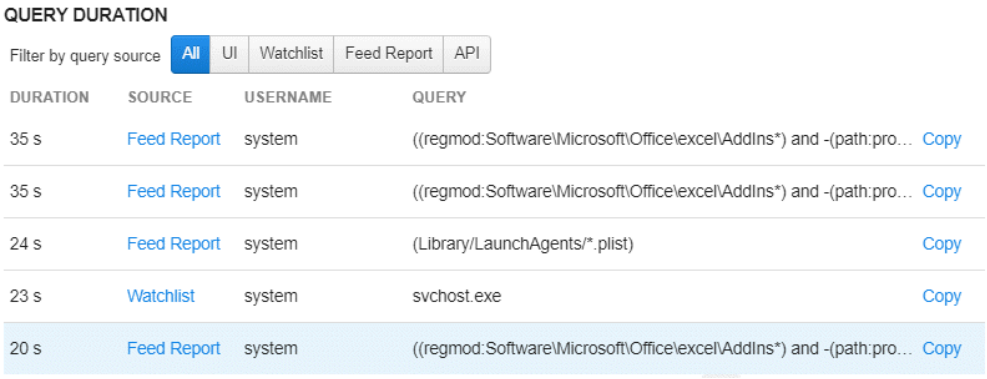

You can filter the displayed queries in the following ways:

- **All** – Displays all queries that take longer than a second to complete.
- **UI** – These slow queries are generated at the UI.
- **Watchlist** – Automated queries. Watchlist queries are created by CB Response users and run every 10 minutes.
- **Feed Report** – Automated queries that the threat research team generates. You cannot edit the queries, but you can ignore them.
- **API** – These queries are run via an API.

A user or script can run UI- or API-generated queries many times. If any query takes long enough to appear in the **Query Duration Panel**, multiple executions of that query add to the overall effect.

For queries that are too long to display in the panel, you can hover over the query to cause the entire query to display in the hover text. You can also click **Copy** to copy a query. This is useful for closely examining a complex slow-running query, and for editing a query to improve performance.

Resolution Time Panel

The **Resolution Time** panel contains a graph that displays the average time (in hours) between reporting and resolution of alerts on each day over a 30-day period.

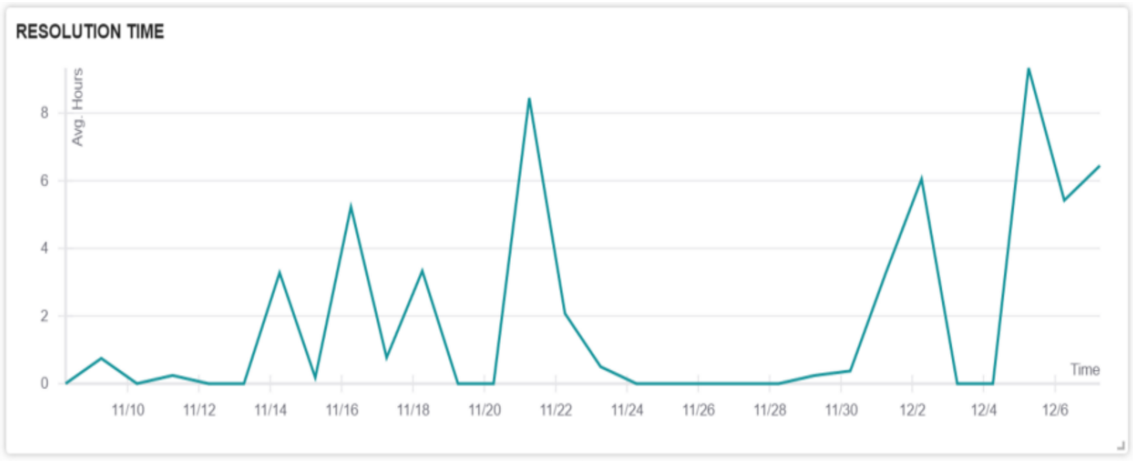

Saved Searches Panel

The **Saved Searches** panel provides a convenient list of your saved process searches. Click on a saved search to go to the Process Search page. See “Save Searches” on page 177.

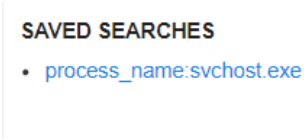

Sensors Panel

Use the **Sensors** panel to monitor and manage sensor hosts. See “Managing Sensors” on page 85.

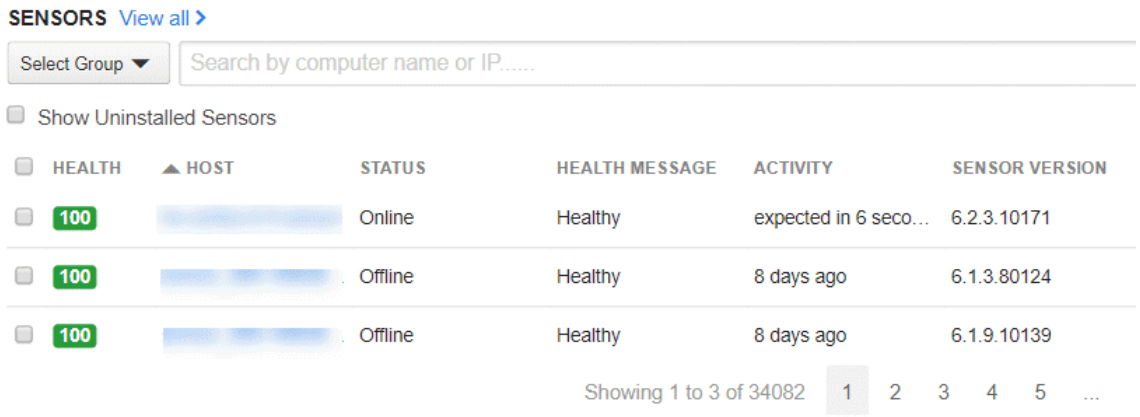

By default, uninstalled sensors do not display in this panel. Select the checkbox to show uninstalled sensors.

By default, all sensor hosts in your organization display; in this case, you cannot perform any actions on the displayed sensor hosts. You can select a group for which you have permissions and then perform the following actions on the hosts in that group:

- Sync
- Restart
- Uninstall; see “Installing Sensors” on page 76.
- Isolate; see “The summary information consists of the following data:” on page 186.
- Remove isolation

You can also search for a specific host by computer name or IP address.

The **Sensors** panel contains the following columns:

Pane	Description
Activity	Displays the time related to the sensor activity.
Health	Displays the sensor health score, where 100 is healthy. Lower numbers indicate problems. See “Sensor Health Score Messages” on page 311.
Health Message	Displays a health message that relates to the sensor health score. See “Sensor Health Score Messages” on page 311.
Host	Displays the name of the host on which the sensor is installed. Click the host name to go to the **Sensors** page for that host. See “Viewing Sensor Details” on page 93.
Sensor Version	Displays the sensor version.
Status	Indicates whether the sensor is online or offline, and whether the sensor is undergoing any activity. For example, if a sensor is online and syncing, the status displays **syncing-online**.

Unresolved Alerts Panel

By default, this panel displays all unresolved alerts for a sensor. You can also display resolved, false positive, and in-progress alerts.

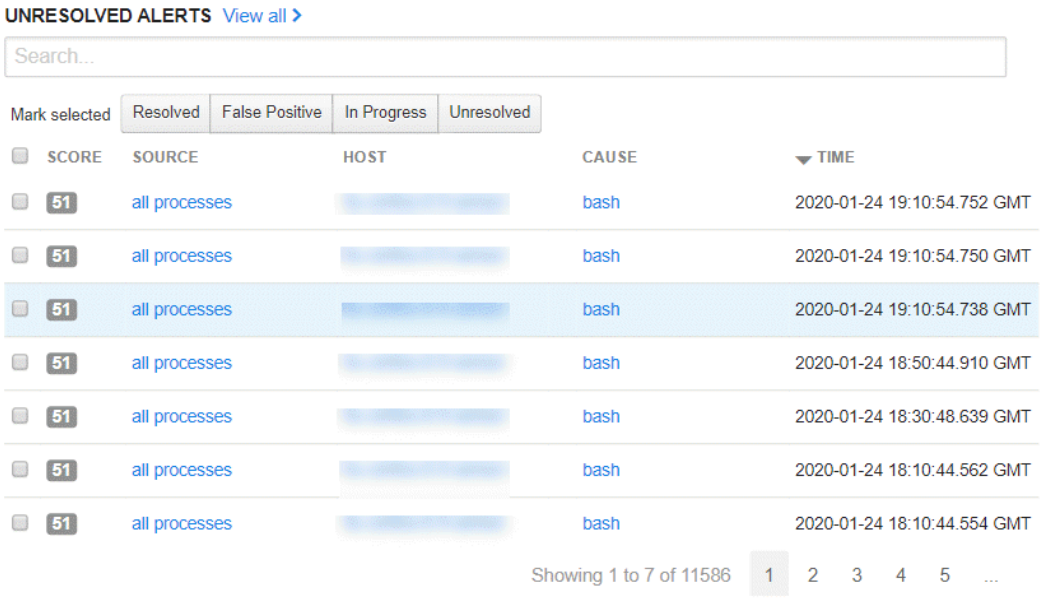

See “Viewing Alert Activity on the HUD Page” on page 287.

Chapter 21

Netconn Metadata

This chapter describes additional or recently added netconn metadata in CB Response. It specifically describes TLS fingerprinting.

Sections

Topic	Page
Overview of TLS Fingerprinting	304
How TLS Fingerprinting Works	304
CB Response Implementation of TLS Fingerprinting	305

Overview of TLS Fingerprinting

Transport Layer Security (TLS) fingerprinting is a platform-independent method for creating TLS fingerprints that can easily be shared for improved threat intelligence. TLS fingerprints are properties of a netconn event for TCP connectivity only.

JA3 and JA3S are TLS fingerprinting methods. JA3 fingerprints how a client application communicates over TLS, and JA3S fingerprints the server response. Combined, they create a fingerprint of the cryptographic negotiation between client and server.

JA3, when used in combination with JA3S, is a useful method to fingerprint a TLS negotiation between client and server. When used in conjunction with a process hash, even greater fidelity can be achieved. For example, some Peer-to-Peer (P2P) client connections can be tracked via TLS fingerprinting. This can be used to correlate an application if the binary and/or process metadata has been changed to avoid more direct forms of identification. Additionally, commodity malware variants often re-use cryptographic information, resulting in a common JA3 hash across families.

How TLS Fingerprinting Works

To initiate a TLS session, a client sends a TLS `Client Hello` packet following the TCP handshake. This packet, and the way in which it is generated, is dependent on packages and methods that are used when building the client application. The server responds with a TLS `Server Hello` packet that is based on server-side supported ciphers and configurations as well as details in the `Client Hello`.

Because TLS negotiations are transmitted in the clear, it is possible to fingerprint and potentially identify client applications using the details in the TLS `Client Hello` packet.

The JA3 method gathers the decimal values of the bytes for the following fields in the `Client Hello` packet:

- Version
- Accepted cipher suites
- List of extensions
- Elliptic curves
- Elliptic curve formats

It then concatenates those values together to create an MD5 hash (or unique fingerprint) that can enhance traditional cybersecurity approaches such as whitelisting, blacklisting, and searching for IOCs.

The JA3S method then gathers the decimal values of the bytes for the following fields in the `Server Hello` packet:

- Version
- Accepted cipher
- List of extensions

It concatenates these values in the same way as the `Client Hello` packet, resulting in an MD5 hash known as a JA3S fingerprint.

CB Response Implementation of TLS Fingerprinting

TLS fingerprinting is available with the 7.1.0 release of CB Response (for Windows 7.0.0 and higher sensors only). It provides additional endpoint telemetry that can be delivered to the CB Response server, and used for narrowing investigations of known malware by identifying known TLS fingerprints.

TLS fingerprints can be specified as IOCs in custom threat feeds. See “Threat Intelligence Feeds” on page 244.

TLS fingerprints can be used in the following ways:

Process Search

TLS fingerprints are searchable via Process Search. See “Overview of Process Search” on page 173. For example:

Process Analysis

TLS fingerprints appear in the Process Analysis page (under netconn events), and as quick filters. See “Process Analysis Page” on page 184 and “Process Event Filters” on page 191. For example:

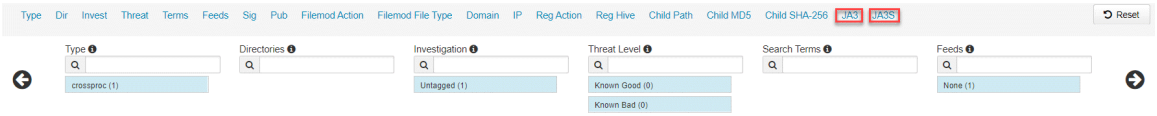

Watchlists

TLS fingerprints can be used in watchlists. See “Watchlists” on page 274. For example, to create a TLS fingerprint watchlist:

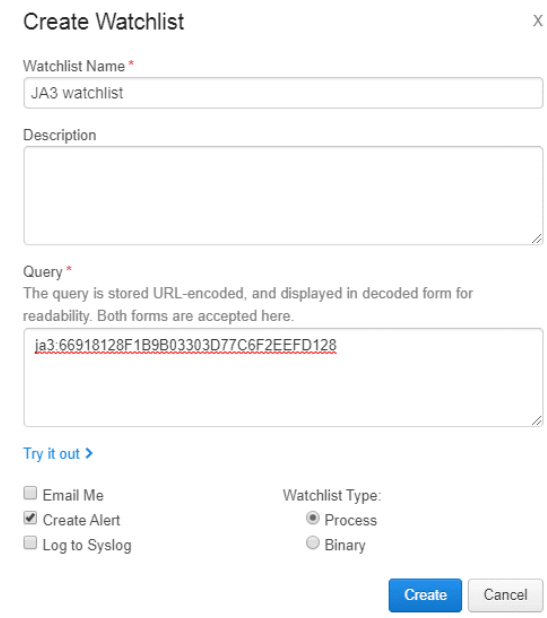

In addition, TLS fingerprints can trigger an alert, email or syslog event.

Appendix A

Sensor Parity

This appendix contains two tables that show which CB Response features or configurations are available for sensors on each operating system platform.

- Sensors are discussed in Chapter 6, “Managing Sensors.”
- Sensor groups are discussed in Chapter 7, “Sensor Groups.”

Sections

Topic	Page
Sensor Feature Support	308
Sensor Group Feature Support	310

Sensor Feature Support

The following table describes whether certain features are available on sensors:

Feature	Windows	Linux	OSX
Binaries (Collection)	Yes	Yes	Yes
Binary Info (Collection)	Yes	Yes	Yes
BinaryModule loads (Collection)	Yes	Yes	Yes
Carbon Black Live Response	Yes	Yes	Yes
Child Process events (Collection)	Yes	Yes	Yes
Compatibility Control	No	No	Yes
Cross Process events (Collection)	Yes	No	No
Retention Maximization	Yes	No	Yes
Diagnostics collection with SensorDiags	Yes	Yes	Yes
Disable sensor operation events	Yes	No	No
EMET events (Collection)	Yes	N/A	N/A
File modifications (Collection)	Yes	Yes	Yes[1]
Global VDI Support	Yes	Yes	Yes
Hash Banning	Yes	Yes[2]	Yes
Hash Banning Whitelist (restrictions)	Yes	No	No
Improved proxy support: WPAD & PAC files	Yes	No	No
Known DLLs (Dylib/Mac) Filtering	Yes	No	Yes
Network Connections (Collection)	Yes	Yes	Yes
Network Connections for IPv6 (Collection)	Yes	Yes	Yes
Network Isolation	Yes	Yes[2]	Yes
Non-Binary File Writes (Collection)	Yes	Yes	Yes
ODX Support	Yes	N/A	N/A
Process Information (Collection)	Yes	Yes	Yes
Process user context (Collection)	Yes	Yes	Yes
Proxy Support (unofficial support)	Yes	Yes	Yes
Registry modifications (Collection)	Yes	N/A	N/A
Server TLS certificate swapping	Yes[3]	No	Yes[3]

Feature	Windows	Linux	OSX
SHA256 hashes in events (Collection)	Yes[4]	No	Yes[4]
Support for FIPS	Yes	No	No
Tamper Detection	Yes	No	No
Tamper Protection	No	No	No
TLS JA3 and JA3S Fingerprinting	Yes	No	No

Notes

1 - The OS X sensor reports a file write event at the time a process opens the file. This event is based on the requested access mask. It is not based on actual writes. Even if the process does not write anything in the file, a file write event occurs.

2 - At the time of this server release, currently available eBPF-based sensors (for RHEL/CentOS 8.0 and SUSE 12&15) do not support isolation or banning.

3 - TLS cert swapping support is for sensor versions Windows 6.2.3-win and Mac 6.2.5-osx and above.

4 - SHA-256 sensor support begins with 6.2.x sensors for both Windows and macOS. Check the User Exchange or Carbon Black Support for any updates about other sensors that can generate this hash type.

SHA-256 hashes are reported in addition to MD5 hashes. They can be used to report information to the Event Forwarder (v3.4.0 or later) and are also displayed on relevant pages in the console. See https://github.com/carbonblack/cb-event-forwarder for information on installing and configuring the event forwarder.

Sensor Group Feature Support

The following table describes whether certain features can be configured for sensor groups:

Feature	Windows	Linux	OSX
Alerts Critical Level	Yes	Yes	Yes
Banning Settings	Yes	Yes	Yes
Binaries (Enable/Disable)	Yes	Yes	Yes
Binary Info (Enable/Disable)	Yes	Yes	Yes
BinaryModule loads (Enable/Disable)	Yes	Yes	Yes
Child Process events (Enable/Disable)	Yes	Yes	Yes
Cross Process events (Enable/Disable)	Yes	N/A	N/A
Retention Maximization (Enable/Disable)	Yes	No	Yes
EMET events (Enable/Disable)	Yes	N/A	N/A
File Modifications (Enable/Disable)	Yes	Yes	Yes
Known DLLs (Dylib/Mac) Filtering (Enable/Disable)	Yes	No	Yes
Network Connections (Enable/Disable)	Yes	Yes	Yes
Non-Binary File Writes (Enable/Disable)	Yes	No	No
Process Information (Enable/Disable)	Yes	Yes	Yes
Process user context (Enable/Disable)	Yes	Yes	Yes
Registry modifications (Enable/Disable)	Yes	N/A	N/A
Sensor Name	Yes	No	No
Sensor Network Throttling	Yes	No	Yes
Sensor Upgrade Policy	Yes	Yes	Yes
Sensor-side Max Disk Usage (%)	Yes	Yes	Yes
Sensor-side Max Disk Usage (MB)	Yes	Yes	Yes
Server TLS certificate swapping (choose cert)	Yes	No	Yes
Server TLS strict certificate validation	Yes	No	Yes
Tamper Level Settings	Yes	N/A	N/A
VDI Behavior Enabled	Yes	Yes	Yes

Appendix B

Sensor Health Score Messages

This appendix describes sensor health score messages that display on the Sensor Details page. Sensors are discussed in Chapter 6, “Managing Sensors.”

Sensor health scores are generated by using a variety of inputs. The default score for a sensor that is running without any issues is 100. Carbon Black subtracts points from this score for events that fall outside of the “healthy range”, based on severity. Sensor health score messages are provided in the console when the sensor is in an unhealthy state.

Health events are presented in priority order. If two events occur at the same time, the message for the higher priority event is presented, regardless of the severity. The sensor can only report one message at a time even when multiple messages occur. The last message type that is processed by the sensor is the one that is reported to the server.

The priority order for each sensor type is listed in the following applicable sections.

Sections

Topic	Page
Windows Health Events	312
macOS Health Events	316
Linux Health Events	319

Windows Health Events

Priority List

1. Driver and component failures
2. NetMon Stream driver failure
3. Service component failure
4. Memory usage
5. GDI handle count
6. Handle count
7. Disk space
8. Event loss
9. Event load

Driver and Component Failures

Cause

This alert occurs if Netmon, Svc component, or core drivers fail to load.

Impact

The sensor does not collect netconn events if the Netmon driver fails. The sensor can stop collecting one or more event types if Svc component fails. The sensor does not collect any events if the core driver fails.

Severity Scale

Driver failure	Health score	Message
Svc component	-25	Svc component failure
Netmon driver	-25	NetMon stream driver failure
Core driver	-100	Core driver failure

Remediation

Restart the failed service. For Netmon issues, a system reboot and re-installation of the network driver might be necessary if issues persist. Contact Support if issues continue.

Memory Usage

Cause

CB Response Sensor service memory usage has risen above expected values.

Impact

Excessive memory consumption can impact system and application performance.

Severity Scale

Memory (MB)	Health score	Message
> 50	-5	Elevated memory usage
> 100	-10	Elevated memory usage
> 200	-20	High memory usage
> 512	-25	Very high memory usage
> 1024	-50	Excessive high memory usage

Remediation

Restart service. Contact Support if issues continue.

GDI Handle Count

This metric records GDI handles usage from the sensor service. GDI handles are used in module extraction only.

Cause

CB Response Sensor service GDI handle usage has risen above normal values.

Severity Scale

GDI handles	Health score	Message
> 100	-5	High GDI handle count
> 500	-10	Very high GDI handle count
> 1000	-20	Excessive GDI handle count

Remediation

Analyze event collection to see if a specific event type is generating an excessive count. If these are non-binary file writes, this collection type can be often be turned off; see Turning off event-collection of Non-Binary file writes.

Handle Count

This metric records kernel handles usage from the sensor service. This metric does not include GDI (Graphics Device Interface) or user handles. Sensors that are running on Windows XP x86 do not report this metric.

Cause

CB Response Sensor service kernel handle usage has risen above normal values.

Severity Scale

Handles	Health score	Message
> 500	-5	Elevated handle count
> 1000	-10	High handle count
> 2000	-25	Very high handle count
> 4000	-50	Excessive handle count

Remediation

Analyze event collection to see if a specific event type is generating an excessive count. If these are non-binary file writes, this collection type can be often be turned off; see Turning off event-collection of Non-Binary file writes.

Disk Space

Cause

The free disk space on the volume where the CB Response sensor is installed has dropped below normal values. This metric does not consider available disk space on other system disks.

Impact

Data collection/usability impact.

Severity Scale

Disk space (MB)	Health score	Message
< 1024	-5	Low disk space
< 100	-25	Very low disk space
< 10	-50	Excessively low disk space

Remediation

Run utilities to clear disk space.

Event Loss

Cause

Events are dropped by the kernel driver due to high activity or component malfunction.

Note that this is calculated by the percentage of total kernel events that were dropped.

Impact

Potential impact to data collection.

Severity Scale

Event loss (%)	Health score	Message
1	-5	Elevated event loss
2	-10	High event loss
5	-20	Very high event loss
10	-50	Excessive event loss

Remediation

Restarting the service should resolve the issue.

Event Load

Cause

The number of outstanding raw kernel events to be processed has exceeded a threshold.

Note that Netconn events are handled in a separate driver.

Impact

Data collection/usability impact.

Severity Scale

Event queue depth	Health score	Message
> 512	-5	Elevated event load
>1024	-10	High event load
> 4096	-25	Excessive event load

Remediation

Analyze event collection to determine what is generating the event load. Consider disabling event collection on certain event types.

macOS Health Events

Priority List

1. Memory usage
2. Out of license
3. Upgrade issue
4. Proxy driver failure
5. Procmon driver
6. Netmon driver

Memory Usage

Cause

CB Response sensor service memory usage has risen above normal values.

Impact

System stability and performance can be impacted if abnormal memory usage persists.

Severity Scale

Memory (MB)	Health score	Message
> 100	-10	Elevated memory usage
> 250	-20	High memory usage
> 512	-25	Very high memory usage
> 1024	-50	Excessive memory usage

Remediation

Restart the service. Contact Support if issues continue.

Out of License

Cause

Server license is expired.

Impact

The sensor is currently unable to push data to the server. Event data is cached on the endpoint. Attempts to send data can cause elevated bandwidth consumption.

Severity Scale

Condition	Health score	Message
Expired license	-25	Out of License

Remediation

Apply updated license to the CB Response server.

Upgrade Issue

Cause

Probably due to an unapproved kext or a required restart at the endpoint to complete the upgrade.

Impact

Inoperable until the condition is resolved.

Severity Scale

Condition	Health score	Message
Upgrade incomplete	-75	Endpoint must be restarted to complete upgrade
Upgrade failed	-75	CB Response kernel extensions are not approved for load

Remediation

Check kext status and approve if necessary for upgrade failure condition. Contact Support if issues continue.

Proxy Driver Failure

Cause

Probable cause is an OS kernel major version mismatch with the proxy driver.

Impact

Sensor does not collect process events correctly because the proxy driver is not connected to OS sys tables.

Severity Scale

Condition	Health score	Message
Driver fails to load	-25	Proxy driver failure

Remediation

Validate that the kernel version is supported by the CB sensor. If the OS version is supported, restart the service. Contact Support if issues continue.

Procmon Driver

Cause

Issue with loading procmon (process monitoring) driver, or version mismatch from a failed upgrade.

Impact

The sensor might stop collecting one or more data collection types.

Severity Scale

Condition	Health score	Message
Version does not match sensor version	-37	Procmon driver version mismatch
Driver fails to load	-37	Procmon driver failure

Remediation

Restart the service. Contact Support if issues continue.

Netmon Driver

Cause

There is an issue with loading netmon (network monitoring) driver, or a version mismatch from a failed upgrade.

Impact

The sensor might stop collecting netconn events.

Severity Scale

Condition	Health score	Message
Version does not match sensor version	-37	Netmon driver version mismatch
Driver fails to load	-37	Netmon driver failure

Remediation

Restart the service. Contact Support if issues continue.

Linux Health Events

Priority List

1. Out of license
2. Failed to get event log stats
3. Driver failure
4. Memory usage

Out of License

Cause

Server license is expired.

Impact

The sensor is currently unable to push data to the server. Event data is cached on the endpoint. Attempts to send data can cause elevated bandwidth consumption.

Severity Scale

Condition	Health score	Message
Expired license	-25	Out of License

Remediation

Apply updated license to the CB Response server.

Failed to get Event log Stats

Cause

There is an issue loading a procmon driver, or there is a version mismatch from a failed upgrade.

Impact

The sensor cannot track the current event log.

Severity Scale

Condition	Health score	Message
Cannot determine current event log stats	-50	Failed to get Event log stats

Remediation

Restart CB daemon. Contact Support if issues continue.

Driver Failure

Cause

An issue occurred when loading a driver.

Impact

The sensor might stop collecting one or more data collection types.

Severity Scale

Condition	Health score	Message
Driver failure	-50	Driver failure

Remediation

Restart CB daemon. Contact Support if issues continue.

Memory Usage

Cause

CB Response Sensor daemon memory usage has risen above normal values.

Impact

System stability and performance can be impacted if abnormal memory usage persists.

Severity Scale

Memory (MB)	Health score	Message
> 75	-5	Elevated memory usage
> 100	-10	Elevated memory usage
> 250	-20	High memory usage
> 300	-25	Very high memory usage
> 450	-50	Excessive memory usage

Remediation

Restart CB daemon. Contact Support if issues continue.