

Carbon Black.



CB Response User Guide

Server/Cloud Version: 6.5
Document Date: August 2019

Copyrights and Notices

Copyright ©2011-2019 Carbon Black, Inc. All rights reserved. This product may be covered under one or more patents pending. CB Response is a registered trademark of Carbon Black, Inc. in the United States and other countries. Any other trademarks and product names used herein may be the trademarks of their respective owners.

This document is for use by authorized licensees of Carbon Black's products. It contains the confidential and proprietary information of Carbon Black, 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 Carbon Black. Carbon Black 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 the Carbon Black products

THERE IS NO WARRANTY FOR THE PROGRAM, TO THE EXTENT PERMITTED BY APPLICABLE LAW EXCEPT WHEN OTHERWISE STATED IN WRITING BY CARBON BLACK. 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.

Carbon Black acknowledges the use of the following third-party software in the CB Response software product:

- Antr 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) Marcin 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 Leibusky
- 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.
- libfreemage.a - Freemage 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 psychogreen - 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.
- Supervisor - Supervisor is Copyright (c) 2006–2015 Agendaless Consulting and Contributors.
- Switchery - Copyright (c) 2013–2014 Alexander Petkov
- Toastr - Copyright (c) 2012 Hans Fjällemark & 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.

Carbon Black, Inc.

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: 6.5

Document Revision Date: August 14, 2019

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	<p>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.</p> <ul style="list-style-type: none"> • 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	<p>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.</p>
CB Response Server	<p>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.</p>
CB Threat Intel Feeds	<p>Pre-configured threat intelligence feeds. These feeds contain threat intelligence data. These feeds come from various sources:</p> <ul style="list-style-type: none"> • Carbon Black • Our MSSP/IR partners • Our customers • Open-source <p>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 14, "Threat Intelligence Feeds."</p>

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)	<p>CB Response sensors constantly monitor your computers for IOCs and send alerts to the CB Response console when detected.</p> <p>Queries are dynamic indicators that look at behaviors that are continuously recorded by sensors on endpoints and centrally recorded for analysis.</p> <p>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.</p>
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 52 or [“Managing User Accounts for Cloud Servers”](#) on page 70 for more information about user roles and privileges.

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

Chapter	Description
1 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.
2 Getting Started	<p>Explains how to log in and out of CB Response (on-premises) and introduces the CB Response main menu.</p> <p>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.</p>

Chapter	Description
3 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.
4 Managing User Accounts for Cloud Servers	Describes how to manage access to the CB Response Cloud console for users and teams of users.
5 Installing Sensors	Describes installing and upgrading sensors on Windows, macOS, and Linux systems.
6 Managing Sensors	Provides an overview of how sensors work, the information that they provide, and how to modify their configuration.
7 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.
8 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.
9 Responding to Endpoint Incidents	Describes CB Response features for incident response—endpoint isolation, Live Response, and process hash banning.
10 Process Search and Analysis	Describes how to perform detailed process searches and in-depth analysis of the processes in search results.
11 Binary Search and Analysis	Explains how to search for and analyze binary metadata.
12 Advanced Search Queries	Describes CB Response query syntax and how to construct advanced queries to search for processes and binaries.
13 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.
14 Creating and Using Investigations	Describes how to work with investigations, which provide a way to group data for reporting, compliance, or retention purposes.
15 Watchlists	Describes creating and using watchlists, which are saved searches that are visible to all users.
16 Console and Email Alerts	Describes creating and managing CB Response alerts, which can be displayed in the console and also sent through email.

Chapter	Description
17 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.
A Sensor Parity	Contains two tables that outline availability of features on sensor systems in different operating systems.
B 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>

Email: support@carbonblack.com

Phone: 877.248.9098

Fax: 617.393.7499

Reporting Problems

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)
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
Reporting Problems	10
1 CB Response Overview	24
What is CB Response?	25
System Architecture	27
CB Response Cloud	28
CB Response On-Premises	28
Data Flow Diagrams	30
CB Response Workflow Overview	33
CB Response APIs	34
2 Getting Started	35
Logging In (On-Premises)	36
Logging In and Configuring Two-Factor Authentication (Cloud Only)	36
Logging In for the First Time from an Email Invitation (Cloud)	37
Configuring Two-Factor Authentication (Cloud)	38
Logging in After Initial Login (Cloud)	43
Enabling/Disabling Two-Factor Authentication (Cloud)	45
Logging Out (Cloud and On-Premises)	45
CB Response Console Controls	46
Navigation Bar	46
Username Menu	48
EU Data Sharing Banner	49
Notifications	50
Help: User Guide and Customer Support	50
3 Managing User Accounts for On-Premise Servers	52
Overview of User Management (On-Premises)	53
Managing User Access with Teams	54
Role-based Privileges for Teams	54
Analyst & Viewer Access by Feature	55
Adding Enhanced Permissions for Analysts	57
User/Team Permissions Example	59
Creating Teams	60
Modifying Teams	61
Deleting Teams	62

Creating User Accounts for On-Premise Servers	63
Changing Passwords	65
Resetting API Tokens	66
Deleting User Accounts	67
Viewing User Activity	68
User Activity API Audit Logging	69
4 Managing User Accounts for Cloud Servers	70
Overview of Cloud User Management	71
Creating Cloud User Accounts	72
Inviting a New or Existing User to Access a Cloud Server	72
Activating an Account from an Invitation	74
Accessing Authorized Servers	76
User Account Lockout	76
Unlocking an Account	76
Viewing and Modifying Cloud User Accounts	77
Changing Security Settings, Email Address or Full Name	77
Changing Administrator / User Status	79
Resetting API Tokens	80
Viewing User Activity	81
Removing a User Account	81
5 Installing Sensors	83
Overview of Sensor Installation	84
Supported Operating Systems and Versions	84
Installing Sensors on Windows	84
HTTP Proxy Support in Windows Sensors	85
Uninstalling Windows Sensors	86
Installing Sensors on macOS Systems	87
Upgrading Sensors on macOS	88
Uninstalling Sensors on macOS	88
Installing Sensors on Linux Systems	88
Upgrading Sensors on Linux	89
Uninstalling Sensors on Linux	89
Upgrading Sensors	90
Uninstalling Sensors via the Console	90
Obtaining New Sensor Installation Packages	91
6 Managing Sensors	92
Overview of Sensor Management	93
Monitoring Sensor Status and Activity	94
The Sensors Page	94
Searching for Sensors	96
Exporting Sensor Data	97
Sensor Actions	97
Monitoring Sensor and Server Information	97
Viewing Sensor Details	100
Sensor Details Heading and Options	101

Sensor Vitals	101
EP Agent	102
Computer Vitals	103
Teams	103
Configuration	103
Sensor Activity	104
Sensor Data Queued - Historical View	104
Sensor Comm Failures	104
Sensor Driver Diagnostics	104
Reducing the Impact of Netconn Data Collection (Windows)	106
Sensor Event Diagnostics	106
Sensor Component Status	106
Sensor Resource Status	106
Sensor Upgrade Status	106
7 Sensor Groups	107
Overview of Sensor Groups	108
Create or Edit a Sensor Group	108
General Settings	110
Sharing Settings	110
Advanced Settings	113
Permissions Settings	115
Event Collection Settings	115
Exclusion Settings (OS X/macOS only)	115
Creating Exclusions	116
Upgrade Policy Settings	118
Moving Sensors to Another Group	118
Deleting Sensor Groups	119
8 Managing Certificates for Server-Sensor Communication	120
TLS Server Certificate Management Overview	121
Certificate Management Feature Summary	122
Server-Sensor Certificate Requirements	122
How CB Response Supports Multiple Certificates	123
Using Multiple Active Certificates in a Cluster	125
Managing Certificates on the Server	126
Viewing Certificate Information in the Console	126
Substituting a Legacy Certificate during Server Installation	127
Adding Certificates through the Console	128
Choosing a Validation Option	129
Changing the Expiration Notification Period	130
Deleting Certificates	131
Upgrades from Previous Server Releases	131
Assigning Certificates to Sensor Groups	132
Assigning different certificates to different sensor groups	132
Assigning a new certificate to all sensor groups	133
Sensor Support for Certificate Management	133
Upgrading to Sensors that Allow Certificate Management	134

9 Troubleshooting Sensors	135
Troubleshooting Windows Sensor Installations	136
Using Control Codes to Generate Logs of Diagnostic Data	136
Debugging Sensor Communications	138
Troubleshooting Linux Sensor Installations	139
General Logging	139
Installation Verification	140
Installation Failures	140
Sensor Communication History	141
Manual Sensor Daemon Start and Stop	141
Determine Server URL	141
Trigger an Immediate Checkin to the Server	141
Driver Debug Parameters	141
Daemon Debug Options	142
Determine Sensor Version	143
Trigger a Diagnostic Data Dump	143
Troubleshooting OSX Sensor Installations	143
Installation Verification	143
Installation Failures	144
Communications Logging	144
Manual Sensor Daemon Start and Stop	144
Determining Sensor Version	144
Determine Server URL	144
Trigger an Immediate Checkin to the Server	145
Trigger a Diagnostic Data Dump	145
Diagnostic Uploads Utility	145
Automatic Crash Data Upload	146
Manual Upload Option (Command Line Utility)	146
Enabling Sensor Diagnostics Uploads	147
File Transfer and Security	147
Data Collected by Sensor Diagnostics	148
10 Responding to Endpoint Incidents	149
Overview of Incident Response	150
Isolating an Endpoint	151
Isolation Exclusions	153
Using Live Response	154
Live Response Endpoint Sessions	154
Registry Access in Live Response	159
Detached Session Management Mode	161
Extending Live Response	162
Live Response Activity Logging and Downloads	162
Banning Process Hashes	163
Creating Process Hash Bans	164
Banning a List of Hashes	165
Managing and Monitoring Hash Bans	167
The Manage Banned Hashes Page	168
Monitoring Banning Events	170

Searching for Blocked Processes	170
Enabling Alerts and Syslog Output for Banning Events	172
Disabling a Hash Ban	173
Disabling or Restricting the Hash Ban Feature	174
Disabling Bans in a Sensor Group	174
11 Process Search and Analysis	175
Overview of Process Search	176
Time Filters	177
Search Filters	177
Enable/Disable Filters	178
Select Multiple Filter Rows	179
Filter Row Percentages	179
Information Icon	180
Filter Search Fields	180
Search Field	180
Saved Searches	180
Clear Preferences	181
Add Search Terms	181
Reset Search	183
Group By Process	183
Search Result Warnings	183
Get Comprehensive Results Button	184
Example Process Search	184
Managing High-Impact Queries	185
Responding to Blocked Searches	185
Process Search Settings in the Console	185
Process Search Settings in cb.conf	186
Results Table	187
Results Table Features	187
Results Table Row Details	188
Process Analysis Page	190
Process Analysis Features	191
Process Summary	192
Isolate Host	192
Go Live	193
Actions Menu	193
Interactive Process Tree	195
Process Execution Details	196
Binary Metadata	197
Feeds	198
On Demand Feeds	198
EMET Protections Enabled (Windows Only)	199
Process Event Filters	199
Event Timeline	202
Process Event Details	202
Process Event Types	206
Analysis Preview Page	209

12 Binary Search and Analysis	212
Overview of Binary Search	213
Entering Search Criteria	213
Additional Search Page Features	215
High-level Result Summaries	215
Related Metadata	217
Binary Search Results Table	217
Binary Preview	218
Binary Analysis	219
Binary Overview	221
Frequency Data	221
Feed Information	221
General Info	222
File Version Metadata	222
Digital Signature Metadata	223
Observed Paths	223
Observed Hosts and Sensor IDs	223
13 Advanced Search Queries	224
Query Syntax Details	225
Terms, Phrases, and Operators	225
Restrictions on Terms	226
Whitespace	226
Parenthesis	226
Negative Sign	226
Double Quotes	227
Leading Wildcards	227
Fields in Process and Binary Searches	227
Fields in Alert and Threat Report Searches	233
Field Types	235
domain	235
ipaddr	236
ipv6addr	236
text	236
count	236
datetime	237
keyword	237
md5	237
sha256	237
path	238
Wildcard Searches	238
Modload Path Searches	238
Regmod Path Searches	238
bool	239
sign	239
cmdline	239
Tokenization Rules	240
Tokenization Changes on Server Upgrade	241
Retention Maximization and cmdline Searches	242

Searching with Multiple (Bulk) Criteria	242
Searching with Binary Joins	244
Example Searches	245
Process Search Examples	245
Binary Search Examples	249
Threat Intelligence Search Examples	250
14 Threat Intelligence Feeds	251
Overview of Threat Intelligence Feeds	252
Threat Intelligence Feed Scores	253
Firewall Configuration for Feeds	253
Managing Threat Intelligence Feeds	253
Checking for New Threat Intelligence Feeds	255
Syncing Threat Intelligence Feeds	256
Data Sharing Settings	256
Enabling, Disabling, and Configuring a Feed	260
On-Demand Feeds from CB Threat Intel	262
Creating and Adding New Feeds	263
Searching for Threat Reports	265
Threat Report Searches and Results	266
Threat Report Details	268
Ignoring Future Reports	269
15 Creating and Using Investigations	270
Overview of Investigations	271
Viewing Investigations	271
Investigations Menu Bar	272
Event Types	272
Bar Graph	273
Events Table	273
Edit Event Description	273
Child Processes	274
Creating Investigations	274
Adding Events to Investigations	275
Removing Events from Investigations	276
Adding Custom Events to Investigations	276
Deleting Investigations	277
16 Watchlists	279
Overview	280
Viewing Watchlists and their Results	280
The Watchlists Table	281
The Watchlist Details Panel	282
Built-in and Community Watchlists	283
Creating Watchlists	284
Managing Watchlists	289
Watchlist Status	289
Watchlist Expiration	289

Slow or Error-producing Watchlists	290
Editing Watchlists	291
Deleting Watchlists	291
17 Console and Email Alerts	292
Overview of Alerts	293
Enabling Console Alerts	293
Watchlist Alerts	293
Threat Intelligence Feed Alerts	294
Viewing Alert Activity on the HUD Page	295
Managing Alerts on the Triage Alerts Page	296
Reviewing Alerts	298
Alerts Table Data	299
Managing Alert Status	300
Ignoring Future Events for False Positive Alerts	301
Enabling Email Alerts	303
Configuring an Email Server	303
Enabling Specific Email Alerts	305
18 Using the Head-Up Display Page	307
Overview of HUD	308
Viewing the HUD Page	308
Customizing the HUD Page	308
Sortable Columns	308
Endpoint Hygiene Panel	309
Event Monitor Panel	309
Query Duration Panel	310
Resolution Time Panel	311
Saved Searches Panel	311
Sensors Panel	311
Unresolved Alerts Panel	313
A Sensor Parity	314
Sensor Feature Support	315
Sensor Group Feature Support	317
B Sensor Health Score Messages	318
Windows Health Events	319
Priority List	319
Driver and Component Failures	319
Cause	319
Impact	319
Severity Scale	319
Remediation	319
Memory Usage	319
Cause	319
Impact	319
Severity Scale	320
Remediation	320
GDI Handle Count	320

Cause	320
Severity Scale	320
Remediation	320
Handle Count.	320
Cause	320
Severity Scale	321
Remediation	321
Disk Space	321
Cause	321
Impact	321
Severity Scale	321
Remediation	321
Event Loss.	321
Cause	321
Impact	321
Severity Scale	322
Remediation	322
Event Load	322
Cause	322
Impact	322
Severity Scale	322
Remediation	322
macOS Health Events	323
Priority List.	323
Memory Usage	323
Cause	323
Impact	323
Severity Scale	323
Remediation	323
Out of License	323
Cause	323
Impact	323
Severity Scale	324
Remediation	324
Upgrade Issue	324
Cause	324
Impact	324
Severity Scale	324
Remediation	324
Proxy Driver Failure.	324
Cause	324
Impact	324
Severity Scale	324
Remediation	325
Procmon Driver	325
Cause	325
Impact	325
Severity Scale	325
Remediation	325
Netmon Driver	325

Cause	325
Impact	325
Severity Scale	325
Remediation	325
Linux Health Events	326
Priority List	326
Out of License	326
Cause	326
Impact	326
Severity Scale	326
Remediation	326
Apply updated license to the CB Response server.	326
Failed to get Event log Stats	326
Cause	326
Impact	326
Severity Scale	326
Driver Failure	327
Cause	327
Impact	327
Severity Scale	327
Memory Usage	327
Cause	327
Impact	327
Severity Scale	327
Remediation	327

List of Tasks

How to . . .

To access the Sensor Details page:	100
To access the Sensors page:	94
To activate a new account from an invitation:	75
To activate access to a new server from an existing account:	75
To add a new certificate to a server through the console:	128
To add a new threat intelligence feed to the CB Response server:	263
To add events from process or binary searches to investigations:	275
To add Exclusion settings to the sensor group panel on the Sensors page:	116
To add or remove administrator status for a user:	79
To add search terms:	181
To apply one certificate to all sensor groups:	133
To ban a list of process hashes:	165
To ban a process MD5 hash from the Process Analysis page:	164
To block or allow high-impact process searches:	186
To change cloud account details:	78
To change the notification period for an expiring certificate:	131
To change the server certificate for one sensor group:	132
To change the status of all alerts matching a search and/or filter:	300
To change the status of one alert:	301
To change the validation method for server certificates:	130
To change your password:	65
To check for new Threat Intelligence feeds:	255
To clear saved searches:	181
To configure an email server for alerts:	303
To configure two-factor authentication:	38
To configure watchlist expiration:	289
To create an investigation from the Respond menu:	274
To create an isolation exclusion:	153
To create an on-premises user account:	63
To create an OS X/macOS event collection exclusion for a sensor group:	116
To create and attach to a Live Response sensor session:	155
To create custom events:	276
To create or edit a sensor group:	108
To create teams:	60
To create watchlists from Process Search or Binary Search pages:	284
To create watchlists from the Threat Intelligence Feeds page:	287
To create watchlists from the Watchlists page:	286
To delete a CB Response team:	62
To delete a certificate from a server:	131
To delete a user account:	67
To delete a watchlist:	291
To delete investigations:	277
To delete sensor groups:	119
To disable a process hash ban:	173
To disable a threat intelligence feed:	262
To disable console alerts for a threat intelligence feed:	294
To disable process hash bans in a sensor group:	174
To display a table of reports from one threat intelligence feed:	265

To display online documentation from the console: 51

To display only certain search filters on the Process Search page: 178

To do a bulk IOC search on the Binary Search page: 244

To do a bulk IOC search on the Process Search page: 243

To edit watchlists: 291

To enable alerts and syslog recording of blocking events due to hash bans: 172

To enable and configure a threat intelligence feed: 260

To enable console alerts for a threat intelligence feed: 294

To enable console alerts for a watchlist: 293

To enable data sharing with Carbon Black threat intelligence feed partners: 258

To enable email alerts for a threat intelligence feed: 305

To enable email alerts for a watchlist: 305

To enable or disable two-factor authentication: 45

To enable sharing communications: 257

To end a Live Response session with a computer: 159

To end network isolation for one or more endpoints: 152

To execute a saved search: 180

To export sensor data from the Sensors page: 97

To ignore the triggering event for an alert: 302

To install sensors on Linux endpoints: 88

To install sensors on macOS endpoints: 87

To install sensors on Windows endpoints: 85

To invite a user to open a CB Response Cloud account or extend it to a new server: 72

To isolate one or more endpoints from the network: 151

To issue a sensor control request to the sensor: 136

To log into CB Response from an email invitation: 37

To log into the CB Response console after initial login: 43

To log into the CB Response console: 36

To log out of the CB Response console: 45

To manually uninstall Linux sensors: 89

To manually uninstall macOS sensors: 88

To manually uninstall Windows sensors: 86

To mark reports as ignored, use one of the following options: 269

To modify a CB Response team: 61

To move sensors to a new sensor group: 118

To open a Live Response command window without a session: 161

To open the Search Threat Reports page (unfiltered): 265

To open the Triage Alerts page: 297

To output debug logging by restarting the daemon (option 1): 142

To output debug logging by restarting the daemon (option 2): 143

To output debug logging by restarting the daemon (option 3): 143

To perform a binary search: 214

To provide enhanced Analyst permissions to a user: 58

To remove a CB Response Cloud user account: 81

To remove an event from an investigation from the Respond menu: 276

To reposition HUD panels: 308

To reset the API token for a user account: 66

To reset the API token for a user account: 81

To resize HUD panels: 308

To search for a sensor: 96

To search for processes that have block events: 170

To sync all threat intelligence feeds on the page: 256

To uninstall sensors using the console (all platforms):	91
To upload a custom “legacy” certificate during server installation:	128
To use the time filter:	177
To view all block events for a parent process:	170
To view banned hash alerts:	173
To view server and sensor information in the Server Dashboard:	98
To view the available certificates on a server:	127
To view the CB Response Cloud servers to which you have access:	76
To view the HUD page:	308
To view the Threat Intelligence Feeds page:	254
To view user activity:	68

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.

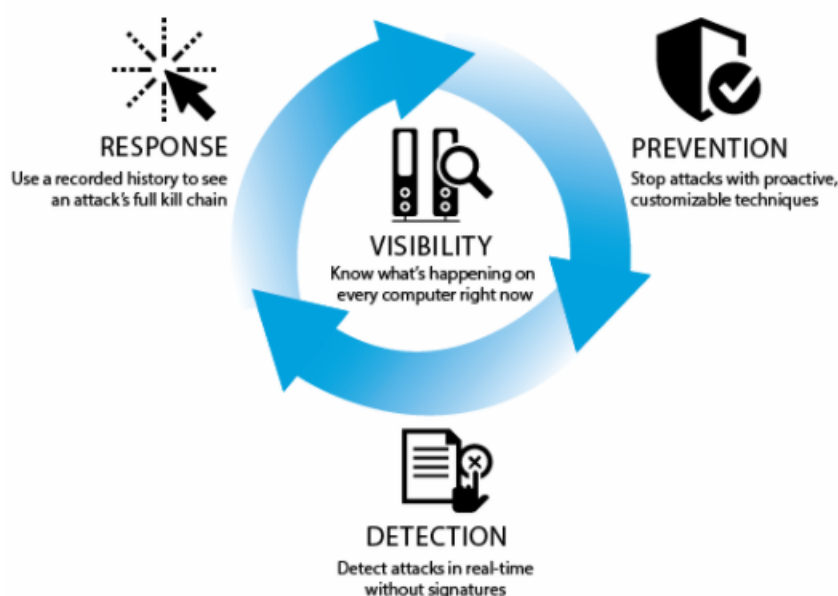
Sections

Topic	Page
What is CB Response?	25
System Architecture	27
Data Flow Diagrams	30
CB Response Workflow Overview	33
CB Response APIs	34

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.



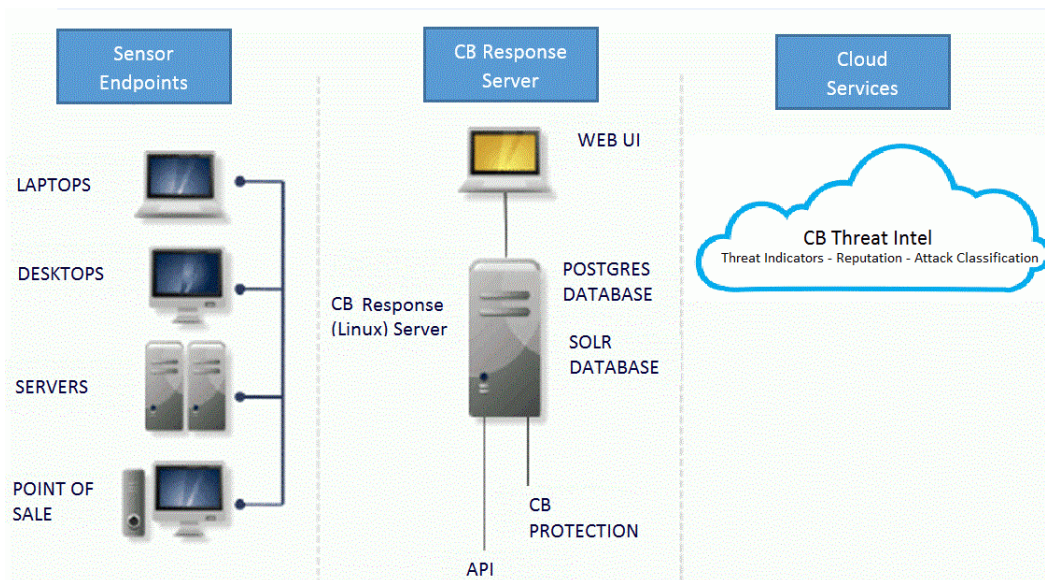
If your company is also using CB Protection, there is integration between it and CB Response. By leveraging CB Protection, you can contain advanced threats by globally blocking or banning them through CB Protection's customizable prevention techniques in the midst of a response.

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	<ul style="list-style-type: none"> • 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 120 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 <code>proxy.pac</code> 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	<ul style="list-style-type: none"> • 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 <code>api.alliance.carbonblack.com</code> and <code>threatintel.bit9.com</code>. • Proxies are supported.
Server to yum Repository	<ul style="list-style-type: none"> • TCP is 443 for HTTPS to <code>yum.distro.carbonblack.io</code>. • 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 them, 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 151.

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 52 and [“Managing User Accounts for Cloud Servers”](#) on page 70 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 available through the navigation bar and top menu, and summarizes the features accessible from those locations.

Sections

Topic	
Logging In (On-Premises)	36
Logging In and Configuring Two-Factor Authentication (Cloud Only)	36
Logging Out (Cloud and On-Premises)	45
CB Response Console Controls	46
Navigation Bar	46
Username Menu	48
EU Data Sharing Banner	49
Notifications	50
Help: User Guide and Customer Support	50

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 with 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 37.
- Configure two-factor authentication. See [“To configure two-factor authentication:”](#) on page 38
- 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)

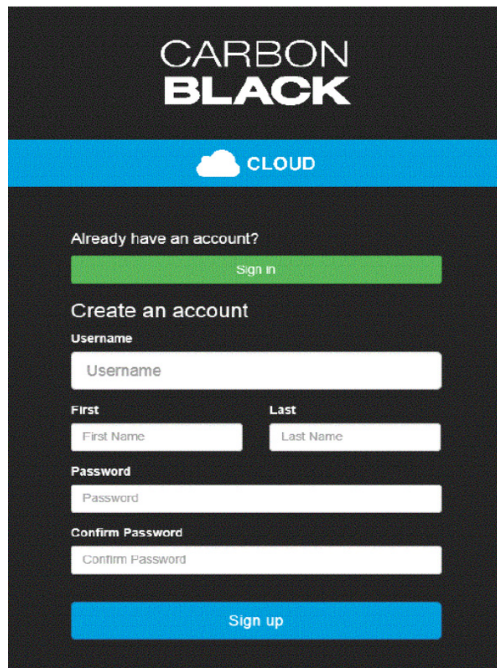
If you have received an email inviting you to access a new CB Response Cloud, use the link in the email to either sign in with an existing account, or create a new account.

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.



The screenshot shows the Carbon Black Cloud login and account creation interface. At the top, the Carbon Black logo is displayed in white on a black background. Below the logo is a blue horizontal bar with a white cloud icon and the word "CLOUD" in white. The main content area is black and contains the following elements:

- A link "Already have an account?" with a green "Sign in" button below it.
- A section titled "Create an account" with a "Username" field.
- Two fields for "First" (First Name) and "Last" (Last Name).
- A "Password" field.
- A "Confirm Password" field.
- A blue "Sign up" button at the bottom.

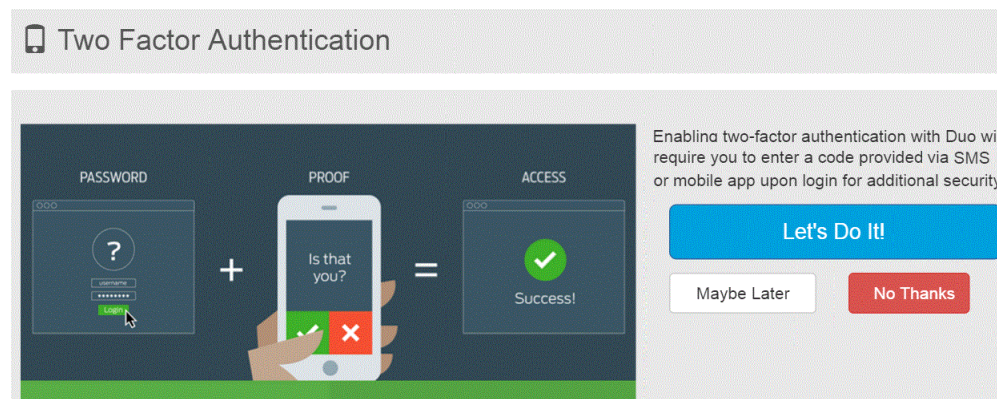
2. Do one of the following:
 - If you already have an account, click **Sign In** and follow the procedures detailed in [“Logging in After Initial Login \(Cloud\)”](#) on page 43
 - To create a new account, enter values in the **Username**, **First**, **Last**, **Password**, and **Confirm Password** fields, 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 45.)

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

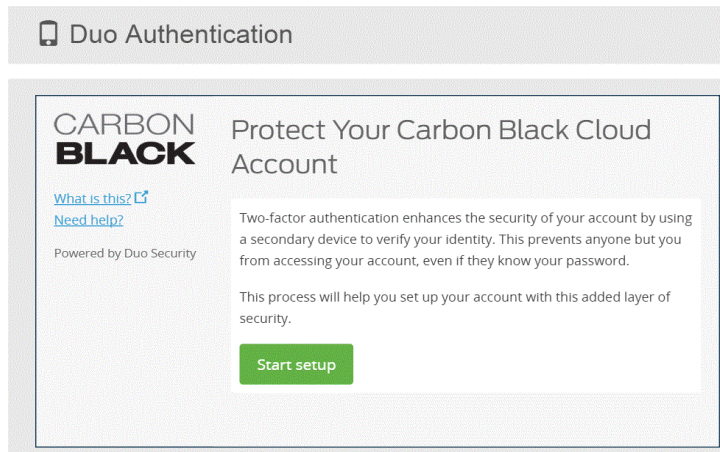
Configuring Two-Factor Authentication (Cloud)

To configure two-factor authentication:

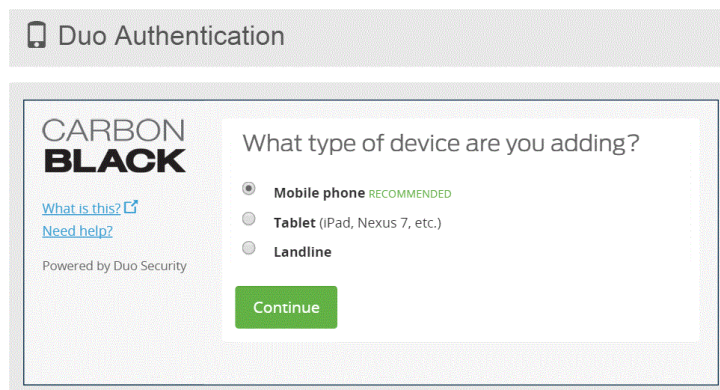
1. Log into the CB Response console.
If you are logging in for the first time, or if you logged in previously and temporarily bypassed enrollment, the Two-Factor Authentication wizard appears.
2. On the first page of the Two Factor Authentication wizard, the following options are available:
 - **Maybe 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.

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

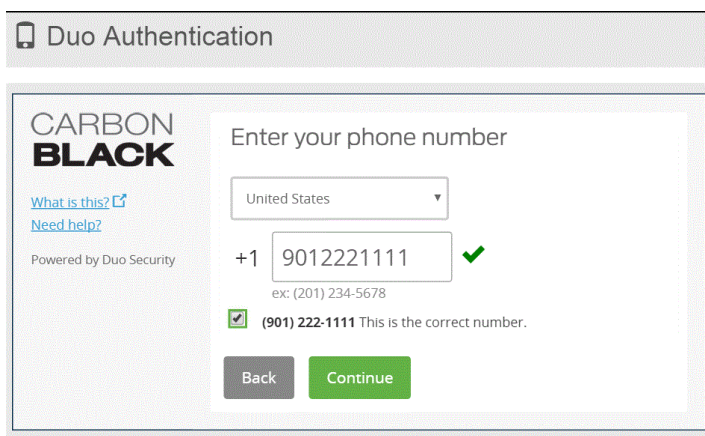


The **What type of device are you adding?** screen appears.

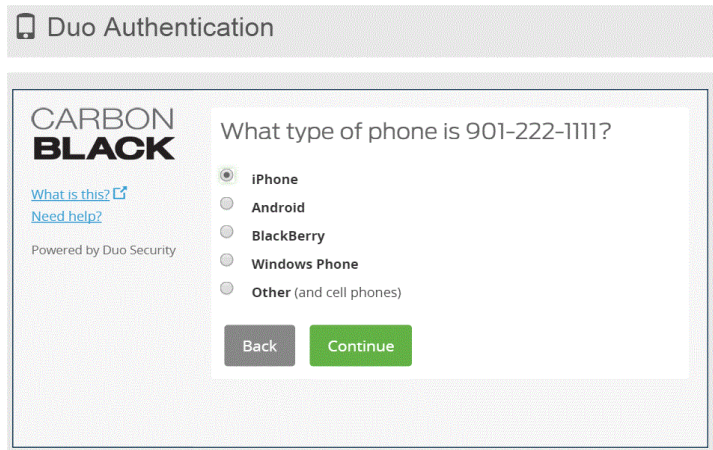


You can add any of several device types for two-factor authentication, including mobile phone devices, tablets, and landlines.

- Select the type of device you are adding and click **Continue**.
- Provide your phone number information and click Continue.

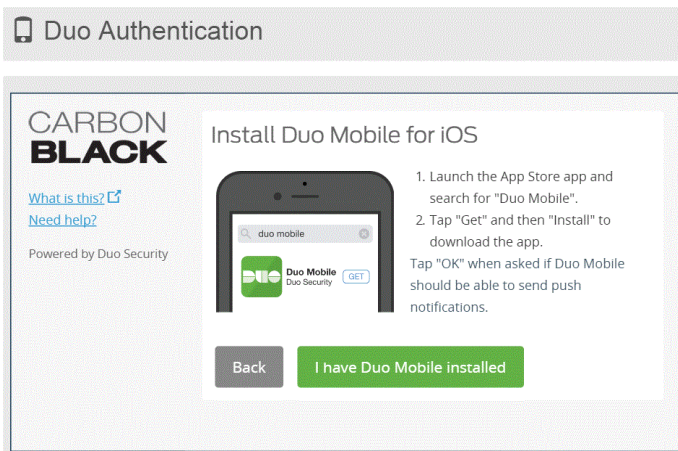


- Select the type of phone device you are using and click Continue.

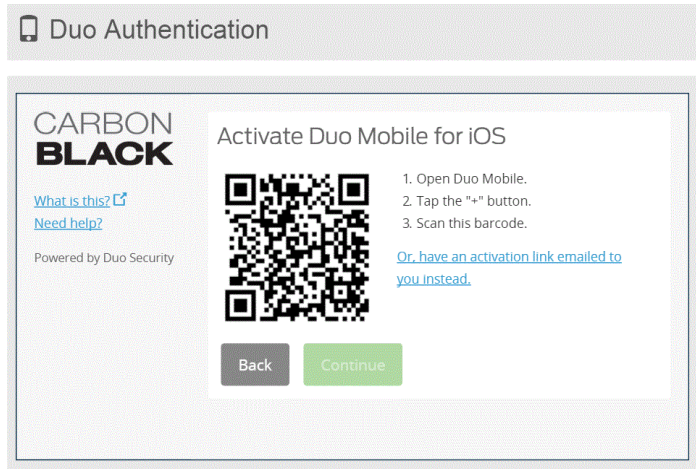


This procedure uses an iPhone as an example, but you can add other types of devices. See Duo product documentation at <https://duo.com> for more information.

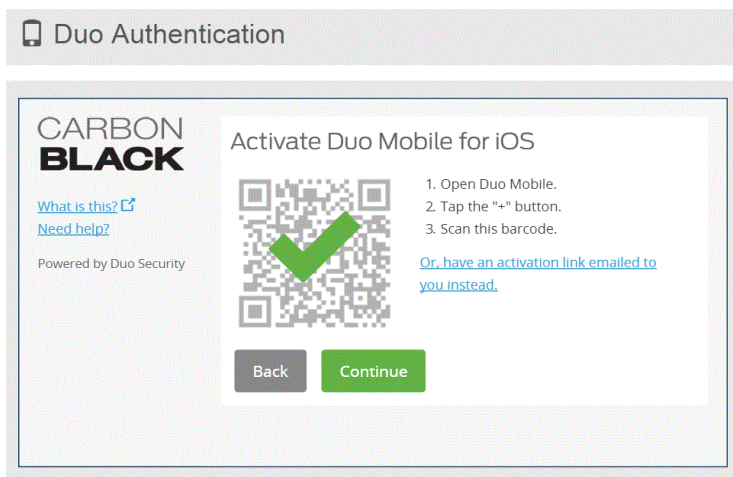
7. Follow instructions to install the Duo Mobile application on your device, and then click **I have Duo Mobile installed**.



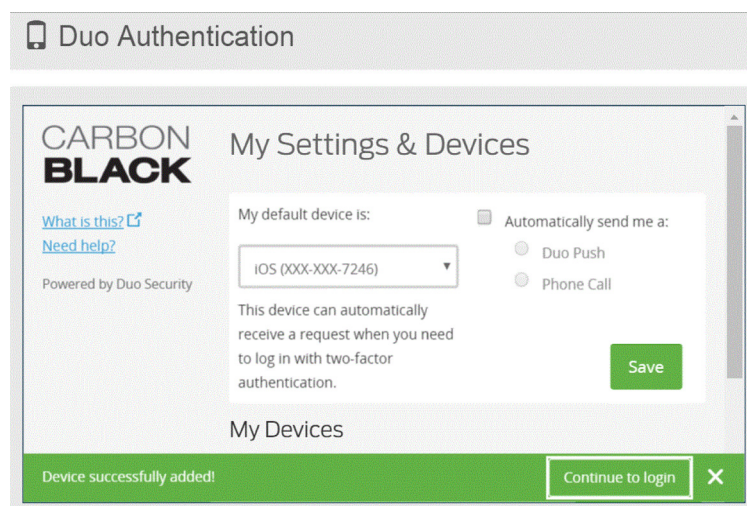
8. On your phone, tap the + button.
9. With your phone, scan the bar code presented in the **Activate Duo Mobile** page:



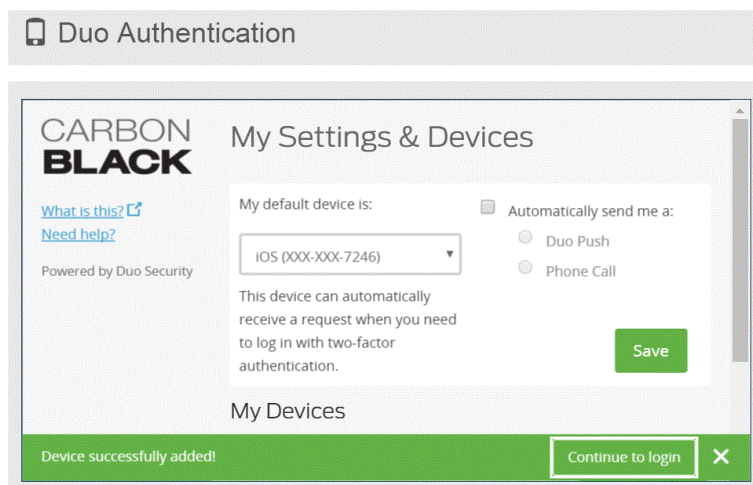
10. When the check mark appears on the bar code indicating success, click Continue.



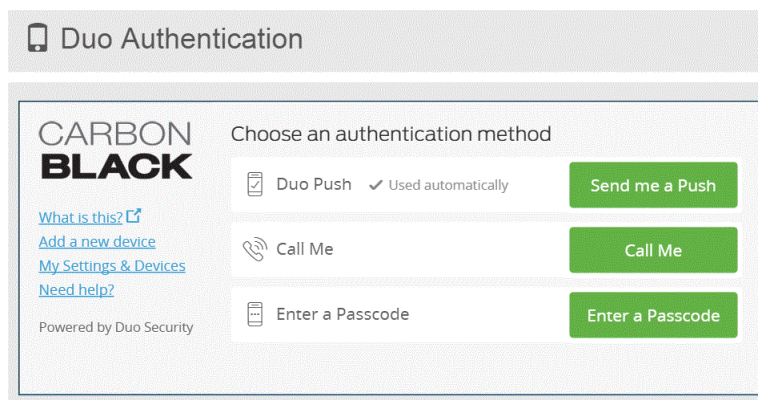
The My Settings & Devices page appears.



11. On the **My Settings & Device** page, make the following selections:
 - a. From the **My default device is** drop-down list, select the device that will be 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 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.



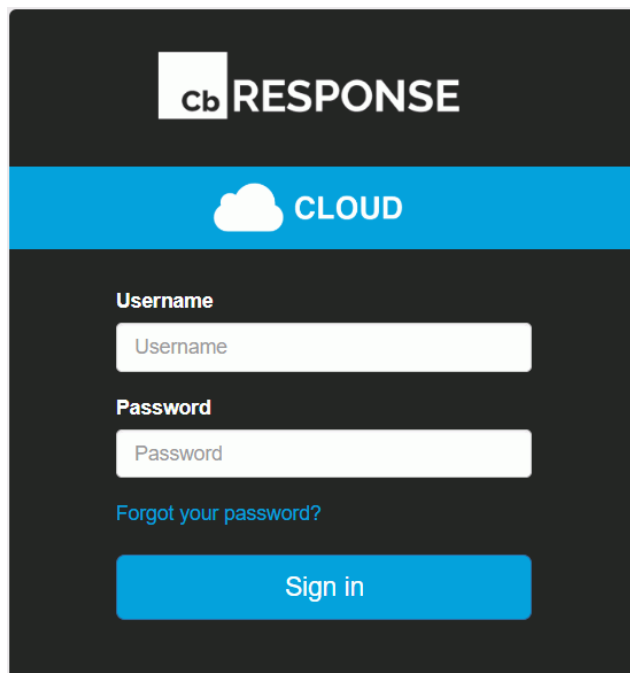
13. When authentication is successful, you are logged into the CB Response Cloud. The CB Response **HUD** (Head-Up Display) page appears.

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.

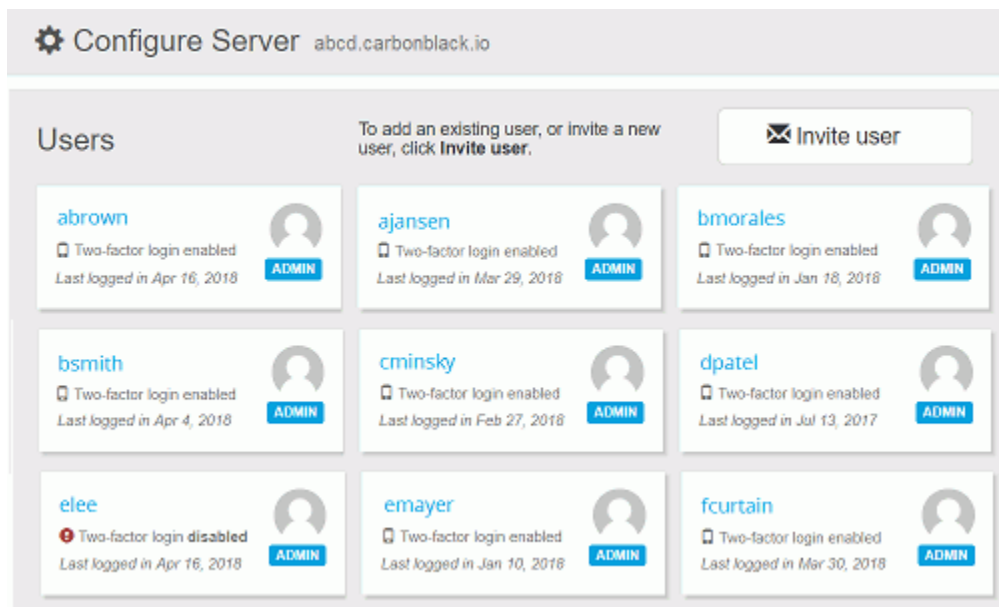


2. Do one of the following:
 - If Username and Password are pre-populated, click **Sign in**.
 - If not pre-populated, enter your **Username** and **Password** and click **Sign in**.

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 38.
- 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 reach other console pages and their features. The examples below show the options available for users with full Administrator/Global Administrator privileges – other users will see options appropriate to their privilege level. Also, the “Teams” option appears only for cloud instances.



Links in the navigation bar are as follows:

Section	Description
Detect	<ul style="list-style-type: none"> • Threat Intelligence provides intelligence feeds. You can set up watchlists (described below), incremental synchronizations and full synchronizations with these feeds. You can also access information about process and binary matches found by each feed. For more information see Chapter 14, “Threat Intelligence Feeds.” • Triage Alerts shows events that match queries defined by watchlists and IOCs (indicators of compromise) defined by feeds. The information provides criteria that is available to search for specific events. For more information see “Managing Alerts on the Triage Alerts Page” on page 296. • Watchlists are saved queries that are performed on process events and binary data stores. The queries contain lists you can use to track specific IOCs. For more information see Chapter 16, “Watchlists.”
Respond	<ul style="list-style-type: none"> • Process Search provides an overview of the sensor process data collection from the sensors that are currently installed. For more information, 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. For more information, see Chapter 12, “Binary Search and Analysis.” • Go Live opens a command line page that provides direct access to sensors. This page is useful when you are performing an investigation (see below), as you can directly access content on endpoints that are running sensors that are providing information. For more information, see Chapter 10, “Responding to Endpoint Incidents.” • Investigations are a collection of tagged process events that are products of search results that come from searching your networks and endpoints for threats. For more information, see Chapter 15, “Creating and Using Investigations.”

Section	Description
<p>Manage</p>	<ul style="list-style-type: none"> • Sensors shows data for sensors and sensor groups. Sensor groups are used to categorize sensors that share the same configuration. You can view, define and update sensors and sensor groups on this page. • Users <ul style="list-style-type: none"> - (on-premises) Displays the User Management page. Tabs on this page allow CB Response administrators to add and configure new users that can log into the CB Response, view user activity, and create and manage teams of users. Teams are groups of users. Users can belong to several teams. Non-administrator users can use this menu item to view their user profile details. - (cloud) Displays CB Response Cloud user accounts that are authorized to access the server. • Teams (cloud only) displays 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. • Server Dashboard shows server statistics such as sensor statistics and server communication status. For more information, see “Monitoring Sensor and Server Information” on page 97. • 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 managed by a CB Response sensor. For more information, 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 from that includes the following options:

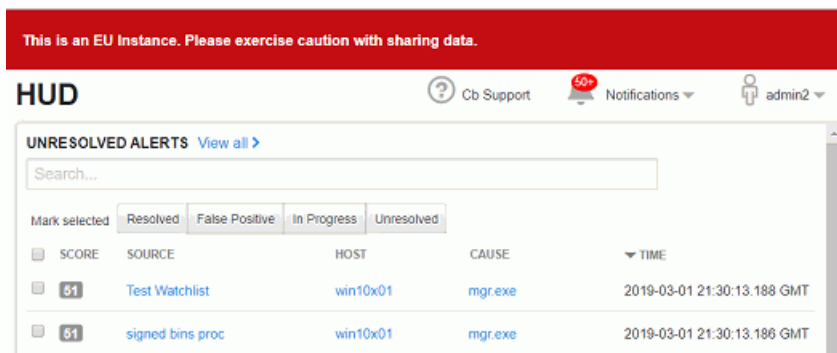
Menu Choice	Description
<p>My Profile</p>	<p>For all users, shows, and allows editing of, the current user’s first and last name, and email address. Shows teams the user is on. Also provides access to dialogs for changing the user’s password and API token.</p> <p>For cloud users, the My profile page 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 77.</p>
<p>Sharing Settings</p>	<p>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 256 for details.</p>

Menu Choice	Description
Settings	<p>For administrators only, this page has several tabs that allow administrators to view and change settings that affect the operation of CB Response:</p> <ul style="list-style-type: none"> • Sites – Provides a menu of sites and the ability to throttle sites by time and day of the week. • E-Mail – Allows you to configure your own alert email server (recommended), use Carbon Black’s email server, or not receive alert email. • 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. • 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 <i>CB Response Integration Guide</i> for more information. • Advanced Settings – <ul style="list-style-type: none"> - 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 185 for details. - EU Data Sharing Banner – Allows administrators to enable and disable the display a red banner at the top of console pages that warns users to be cautious when sharing screenshots or other data from this instance. Note that this can be overridden 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. See “Using Live Response” on page 154 for details. Note that this can be overridden in the cb.conf file by the CbLREnabled setting.
Logout	Logs the current user out of the CB Response console.

EU Data Sharing Banner

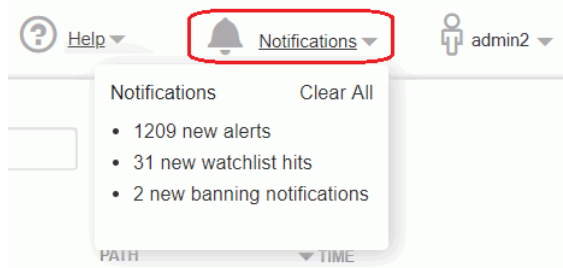
CB Response displays information from all of the endpoints that report to a server through the sensor. Because some of this information may be sensitive, you might need to take extra steps to avoid exposing it in the wrong places. The may include restricting access to features that access this data.

As an extra precaution, CB Response provides administrators have the ability to display a red banner at the top of console pages that warns users to be cautious when sharing screen shots or other data from this instance. This banner is enabled and disabled on the Advanced Settings tab of the Settings page.



Notifications


At the top of the console, the Notifications 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 showing the number of each type of notifications. In addition, clicking on any item on the menu takes you to the page that provides details for the item the notification refers to.



For example, if you clicked on “31 new watchlist hits” in the example above, it would take you 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** on the menu 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.

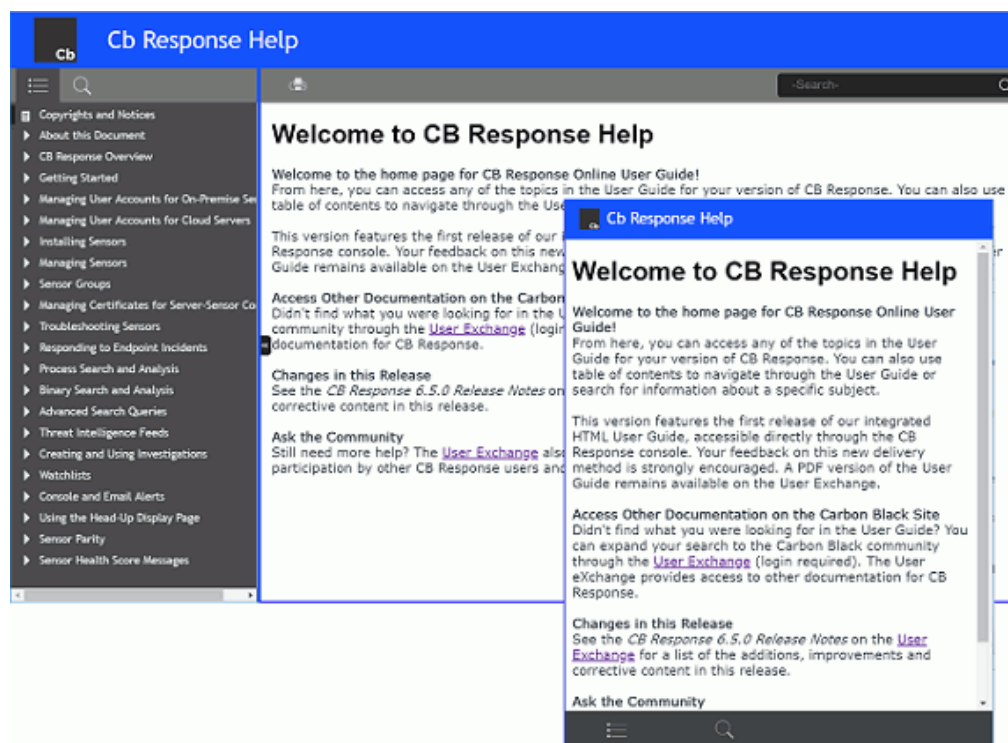


- The **CB Support** choice opens a new browser window or tab showing the login page for the Carbon Black User Exchange, the customer portal for users of all Carbon Black products and services.
- The **User Guide** choice opens a new browser window or tab showing the HTML version of the *CB Response 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.

As with the CB Response console itself, the online version of the User Guide is compatible with Chrome browsers, preferably the latest release. You may be able to display help in other browsers but these are not supported and they might have issues that interfere with display or performance of help content.

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, and their location, vary depending upon 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 to 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 more 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 installed on your premises. This includes managing teams to determine user roles and managing user accounts. For information on managing users for CB Response Cloud, see [“Managing User Accounts for Cloud Servers”](#) on page 70.

Sections

Topic	Page
Overview of User Management (On-Premises)	53
Managing User Access with Teams	54
Role-based Privileges for Teams	54
Adding Enhanced Permissions for Analysts	57
Modifying Teams	61
Creating User Accounts for On-Premise Servers	63
Changing Passwords	65
Resetting API Tokens	66
Deleting User Accounts	67
Viewing User Activity	68

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, which means that this user has full access to all sensors and all features. Once you log in with 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 will have administrator privileges (for all Sensor Groups and features) or not. If a user is a Global Administrator, the other factors below 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 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. For more information, see [“Role-based Privileges for Teams”](#) on page 54.
- **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 that you might want to restrict access to. For example, threat feeds, which are not specific to any Sensor Group, are an important tool for CB Response threat monitoring, and 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 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 57 for details.

A table with more specific details about team privileges appears in the [“Managing User Access with Teams”](#) on page 54.

Important

Creation and management of user accounts and teams is available only to Global Administrators in on-premise installations and 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, and others may 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 also 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, create additional Sensor Groups. See [“Sensor Groups”](#) on page 107 for more information.

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 available from any of the teams.

Although you can assign a user to teams later, it is helpful to have teams set up before you create non-global-administrator users since the capabilities of most users will be determined by the teams they are on.

Role-based Privileges for Teams

The roles you may 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 are allowed to use special features: Live Response, isolation, hash banning, toggling tamper detection, and uninstalling the sensor. See [“Adding Enhanced Permissions for Analysts”](#) on page 57.

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 CB Response gathers about endpoints, 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 visible only 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 will be 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	<p>Viewers: Can view tables and details of sensors in Sensor Groups for which the user has Viewer access.</p> <p>Analyst: In addition to viewing, 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.</p> <p>Analysts can also move sensors between Sensor Groups if they are Analysts for both the source and destination Sensor Groups.</p>
Sensor Groups	<p>Viewers: Can view tables and details of Sensor Groups for which the user has Viewer access.</p> <p>Analyst: In addition to viewing, can perform certain actions involving Sensor Groups for which the user has Analyst access:</p> <ul style="list-style-type: none"> • Can toggle tamper detection if the user also has the 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. <p>An Analyst cannot add or delete a Sensor Group.</p>
Uninstall Sensors	<p>Viewers: No Access</p> <p>Analyst: Can uninstall sensors from the console in Sensor Groups for which the user is an Analyst <i>if the user also has the enhanced permission for uninstalling sensors.</i></p>
Users, Teams and Activity Audit	Only available to Global Administrator (on-premises) or Administrator (cloud)
Tamper detection toggle	<p>Viewer: No Access</p> <p>Analyst: Can turn tamper detection on and off for Sensor Groups for which the user is an Analyst <i>if the user also has the enhanced permission for tamper detection.</i></p>

Feature or Page	Permissions by Role
HUD page	<p>Viewer: Can view the page filtered to show alerts and sensors in Sensor Groups for which the user is a Viewer.</p> <p>Analyst: In addition to viewing data in Sensor Groups for which the user has Analyst access, can take action on alerts.</p>
Threat Intel Feeds	<p>Viewer: No Access</p> <p>Analyst: Can view and modify the page, including enabling and disabling actions on hit (Email Me, Create Alert, or Log to Syslog).</p>
Triage Alerts	<p>Viewer: Can view all binary alerts, and can view other alerts in Sensor Groups for which the user is a Viewer.</p> <p>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.</p>
Watchlists	<p>Viewer: Can view watchlist results for binary searches and other searches involving Sensor Groups for which the user is a Viewer.</p> <p>Analyst: In addition to view access, can also add, modify, and delete watchlists, and can take actions including enabling and disabling email notification, log to Syslog, and alerts.</p>
Process Search	<p>Viewer and Analyst: Can view process search results for Sensor Groups for which the user has at least Viewer access.</p>
Process Analysis	<p>Viewer: Can view process analysis results for Sensor Groups for which the user has at least Viewer access.</p> <p>Analyst: In addition to view access, can take actions for processes in Sensor Groups for which the user is an Analyst <i>if the user also has the enhanced permission for that action</i>. Actions include Isolate host, Go Live, and Ban Hash.</p>
Binary Search (results) & Analysis (details)	<p>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.</p> <p>Analyst: In addition to Viewer access, can ban hashes in the search results <i>if the user also has the enhanced permission to Ban hashes</i>.</p>
CB Live Response	<p>Viewer: No Access.</p> <p>Analyst: Can use Live Response to access and take actions on the endpoints in Sensor Groups for which the user is an Analyst <i>if the user also has the enhanced permission for Live Response</i>.</p>

Feature or Page	Permissions by Role
Investigations	Viewer: Can view the Investigations page. Actions are limited to Export events to CSV and Export timeline to PNG. Analyst: In addition to viewing, 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 <i>if the user also has the enhanced permission for isolating sensors</i> .
Banned Hashes	Viewer: No Access. Analyst: Can ban hashes and remove bans if the user also 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 is defined to allow access to features for monitoring and investigation of suspicious or malicious activity on endpoints. You might also choose to allow certain 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	Enables a user who is an Analyst for any Sensor Group to ban files by hash and remove bans. These bans are applied to all sensors.
Isolate sensor	Enables a user who is an Analyst for a Sensor Group to isolate a sensor <i>in that group</i> from the network and restore the sensor from isolation.
Live Response	Enables a user who is an Analyst for a Sensor Group to connect to and act on a sensor <i>in that group</i> using Live Response.
Tamper detection	Enables a user who is an Analyst for a Sensor Group to disable and enable reporting of tamper events for all sensors <i>in that group</i> .

Enhanced Permission	Description
Uninstall sensors	Enables a user who is an Analyst for a Sensor Group to use the console to uninstall a CB Response sensor in the group.

Note

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 on at least one Sensor Group.

To provide enhanced Analyst permissions to a user:

- From the main navigation bar:
 - For an on-premises server, choose **Users** to display the User Management page.
 - For a cloud instance, choose **Teams** to display the Team Management page and click on the **Users** tab.
- Locate the name of the user you want to give enhanced permissions to and click on the edit user button to the right of the name.

If this is 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 63.
- In the Enhance Analyst permissions panel, check the box next to each permission you want to give this user (the on-premises Edit User pages is shown below).

If the user is not a member of an Analyst team at this point, there is a gray triangle icon 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.

- If necessary, add the user to a team with Analyst permission.

Assign to teams: [Select All / Deselect All](#)

Analysts - East
 Analysts - West
 General Viewers
 Super Analysts

Enhance Analyst permissions:

To give this user enhanced privileges in addition to generic Analyst capabilities, assign them to a team with Analyst privileges on any sensor group.

Ban hashes
 Isolate sensors
 Live Response
 Tamper detection
 Uninstall sensors

5. Click **Save changes**.

The user now has the permissions you chose.

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 of the endpoints in these countries will have sensors that are managed by one CB Response cluster.

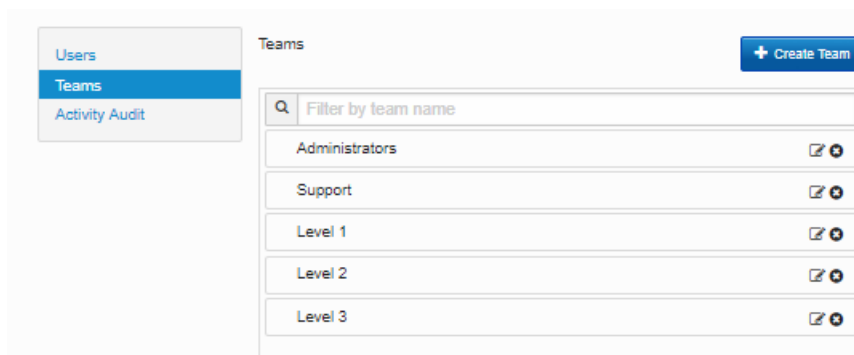
- **Create Administrators:** You might make two users in each country Global Administrators (or for cloud instances, Administrators), so that they can set up their users and user teams. These Global Administrators can also monitor system performance, change Settings that control the behavior of the server, and do anything that is possible with the available features of CB Response. While in this example their primary responsibility is for the sensors in their own country, each Global Administrator is capable of doing any CB Response activity on the any country's endpoints if necessary.
- **Create country-specific Analysts:** You might create four additional users in each country that you assign to teams that make them Analysts for the Sensor Groups corresponding to the endpoints in their own country. This means they will be able to monitor the data from the sensors in these groups.
- **Enhance permissions for some Analysts:** For two of the Analysts in each country, you might add one or more enhanced permissions that allow them to take actions that affect sensors, including isolation, Live Response, uninstalling the sensor, banning hashes, and turning tamper detection on and off.
- **Let Analysts be Viewers for other Sensor Groups:** So that Analysts can be aware of activities or trends that could affect all countries, you could give them the Viewer role for the Sensor Groups in countries where 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

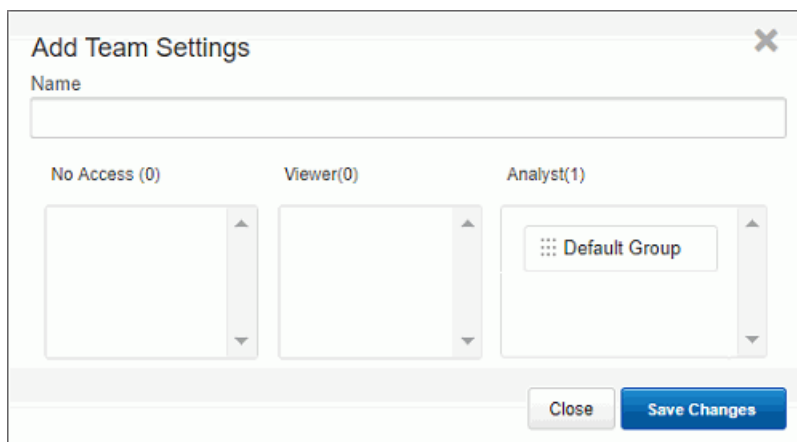
You create teams from the **User Management** page for on-premises installations and the **Team Management** page for cloud installations.

To create teams:

- From the CB Response navigation bar, go to the team management page corresponding to your installation type:
 - For on-premises installations, choose **Users** in the navigation bar and then click **Teams** in the left panel of the User Management page.
 - For cloud installations, from the choose **Teams** in the navigation bar and then click **Teams** in the left panel of the Team Management page.



- Click the **Create Team** button in the top-right corner of the page to display the **Add Team Settings** page:



- In the **Name** field, enter a name for the team.
- Drag and drop the Sensor Groups to the list with the type of permissions that are appropriate for this team.

For example, if you want 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).

- Click **Save Changes**.

Modifying Teams

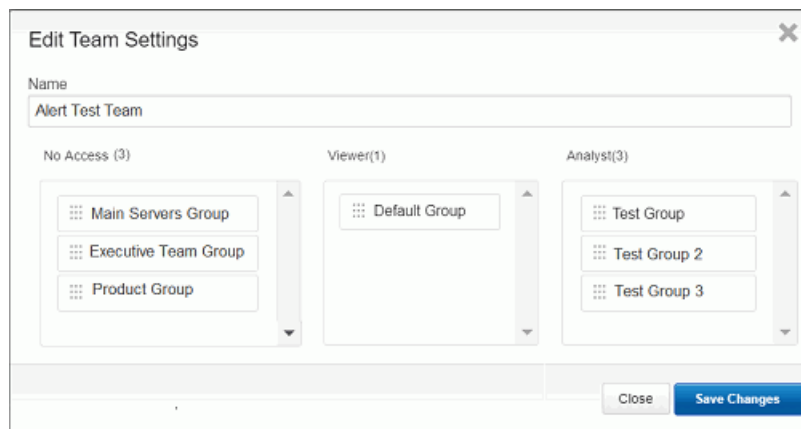
You can modify teams from the **User Management** page for on-premises installations and the **Team Management** page for cloud installations.

To modify a CB Response team:

1. From the CB Response navigation bar, go to the team management page corresponding to your installation type:
 - For on-premises installations, choose **Users** in the navigation bar and then click **Teams** in the left panel of the User Management page.
 - For cloud installations, from the choose **Teams** in the navigation bar and then click **Teams** in the left panel of the Team Management page.
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

You can delete teams from the **User Management** page for on-premises installations and the **Team Management** page for cloud installations. When you delete a team, references to the team in user accounts are deleted as well, but the user accounts remain active.

To delete a CB Response team:

1. From the CB Response navigation bar, go to the team management page corresponding to your installation type:
 - For on-premises installations, choose **Users** in the navigation bar and then click **Teams** in the left panel of the User Management page.
 - For cloud installations, from the choose **Teams** in the navigation bar and then click **Teams** in the left panel of the Team Management page.
2. In the list of teams, click the **Delete (x)** icon to the far right of the team name:



3. In the confirmation window, click **OK** to delete the team.
A popup message appears in the upper right area of the console to report on the success or failure of this 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 since the capabilities of most users will be determined by the teams they are on. See [“Managing User Access with Teams”](#) on page 54 for more information.

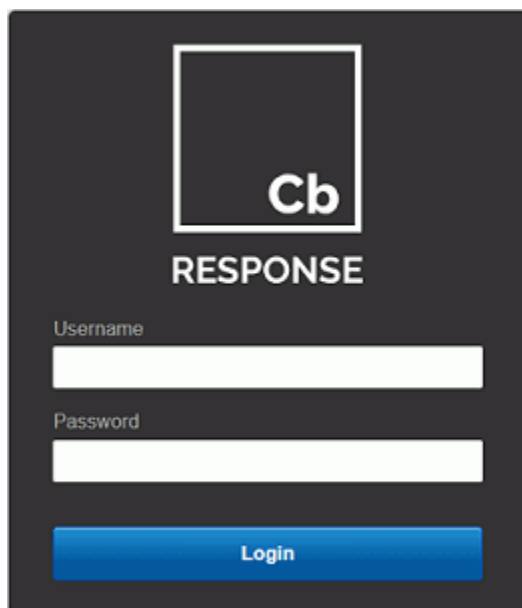
To create user accounts for on-premise servers, log in to the CB Response server console using an account that has Global Administrator status. If no other users have been created yet, use the administrative account and password that were set up in the `cbinit` script during the server installation process.

To create an on-premises user account:

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

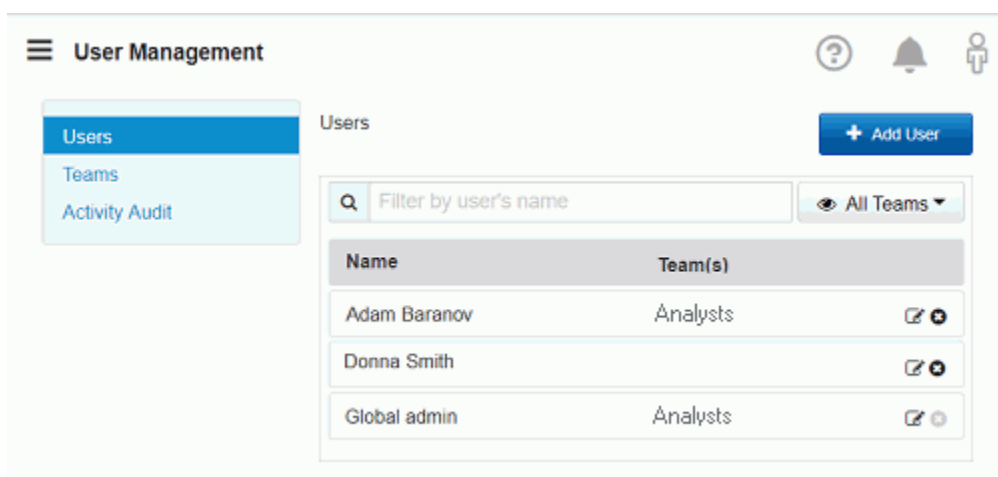
```
https://<your CB Response server address>/
```

The CB Response **Login** window appears:

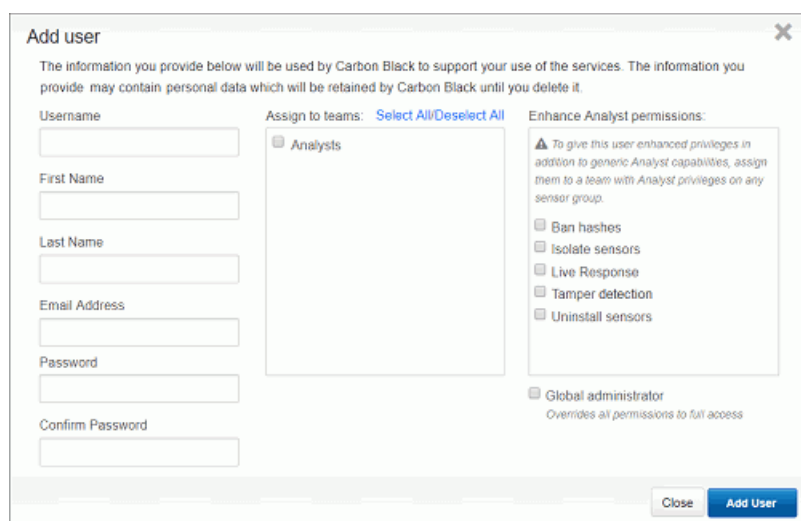


2. Enter the username and password for a Global Administrator account.

- From the navigation bar, click **Users** to display the **User Management** page:



- Click **Add User** in the top-right corner to display the **Add User** page:



- Enter the following information in the **Add User** page:

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 attempt to create a user account with an illegal character, the console will display a warning message.
First Name	First name of the user.
Last Name	Last name of the user.
Email address	Email address for the user.

Field	Description
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 to ensure that it is the one you intended to use.
Assign to teams	Select the teams the user will belong to. The default team is Analysts. Users can belong to more than one team. For more on teams, see “Managing User Access with Teams” on page 54.
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 57.
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

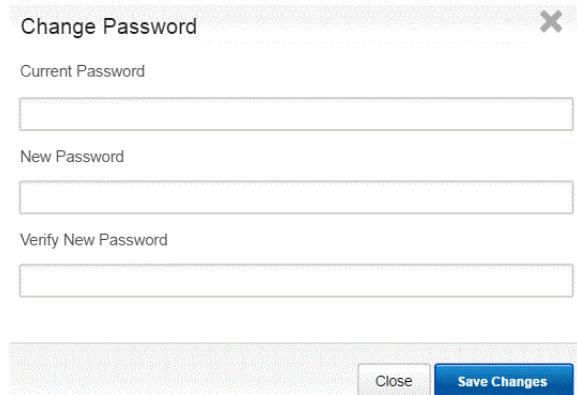
Carbon Black recommends that new users change their passwords after logging in for the first time.

To change your password:

1. Login to the CB Response console.
2. Select **Username > My Profile**. The My Profile dialog opens:

The screenshot shows the 'My Profile' dialog box. On the left, there is a sidebar with 'Profile Info' and 'API Token' options. The main area contains a text block: 'The information you provide below will be used by Carbon Black to support your use of the services. The information you provide may contain personal data which will be retained by Carbon Black until you delete it.' Below this are input fields for 'First Name' (Ann), 'Last Name' (Adams), and 'Email Address' (aadams@mycorp.com). To the right, there is a 'My Teams' section with a dropdown menu showing 'Analysts'. At the bottom, there are links for 'Change Password' and 'Clear Preferences', and a blue 'Save changes' button.

3. In the Profile Info view, click **Change Password** to display the Change Password dialog box:



Change Password ✕

Current Password

New Password

Verify New Password

Close Save Changes

4. Enter the following information:
 - **Current Password**
 - **New Password**
 - **Verify New Password**
5. 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 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 when making calls to the APIs. Users can reset their own API token at any time by following these procedures.

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. Select **Username > My Profile** in the navigation bar.
3. In the **My Profile** window, select **API Token**.
4. Click the **Reset API Token** button to reset the API token.
5. A notification appears briefly in the top-right corner of the console notifying you that the API token has been reset.



Deleting User Accounts

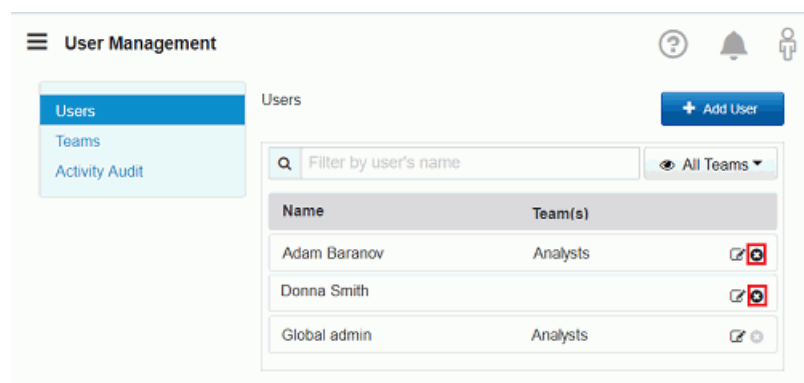
A user account can be removed from the system when that user no longer needs access to the CB Response console or leaves the organization. Users with Global Administrator privileges can delete any account except their own and the built-in administrator account. If the user with the deleted account belongs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

To delete a user account:

1. From the main navigation bar, choose **Users** to display the User Management page:



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 will be grayed out, indicating that the user cannot be deleted while logged in.
3. In the confirmation window, click **OK** to delete the user account.

A popup message appears in the upper right area of the console to report on the success or failure of this deletion. The Activity Audit for this user will also show the command delete action.

Viewing User Activity

CB Response keeps an audit trail of user activity on the console.

To view user activity:

- From the CB Response navigation bar, go to the page corresponding to your installation type:
 - For on-premises installations, choose **Users** in the navigation bar and then click **Activity Audit** in the left panel of the User Management page.
 - For cloud installations, from the choose **Teams** in the navigation bar and then click **Activity Audit** in the left panel of the Team Management page:

Username	Timestamp	Remote Ip	Request Information	Result	Description
admin	2019-02-21 08:49:55.527729-05:00	::ffff:10.42.15.125	POST /api/team	200	OK
admin	2019-02-21 08:47:58.441672-05:00	::ffff:10.42.15.125	GET /api/users	200	OK
Unknown	2019-02-21 08:47:17.914452-05:00	::ffff:10.42.15.125	GET /api/auth	403	Requires Authentication
Unknown	2019-02-21 08:47:06.514204-05:00	::ffff:10.42.15.125	GET /api/auth	403	Requires Authentication

- Here, you can view the following information:

Field	Description
Username	The user name of the user who has accessed the CB Response console.
Timestamp	The full date and time that the user logged into the CB Response console.
Remote IP	The IP address of the computer that the user logged in on.
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 would show an HTTP 200 code response. A request to access a resource the user does not have permission for usually results in redirection to the HUD page, but if that does not occur an HTTP 403 code would display.
Description	The HTTP response description. For example, an HTTP 200 response would show "OK" as a description, while an HTTP 403 response would show a "Requires Authentication" response.

- 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 view the log for user activity directly in the following file:

```
/var/log/cb/coreservices/debug.log
```

User Activity API Audit Logging

When API audit logging is enabled for a server by setting `EnableExtendedApiAuditLogging=True` in the `cb.conf` configuration file, 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:192.168.56.1' status='200' method='GET' path='/api/
v2/sensor'
2017-12-22 11:30:54: username='bill' userid='1'
ip='::ffff:192.168.56.1' status='200' method='GET' path='/api/
v1/alert'
2017-12-22 11:30:55: username='bill' userid='1'
ip='::ffff:192.168.56.1' status='200' method='GET' path='/api/
v1/detect/report/currentmonitoringstatus'
2017-12-22 11:30:55: username='bill' userid='1'
ip='::ffff:192.168.56.1' status='200' method='GET' path='/api/
v3/group'
2017-12-22 11:30:57: username='bill' userid='1'
ip='::ffff:192.168.56.1' status='200' method='GET' path='/api/
v1/feed'
2017-12-22 11:30:57: username='bill' userid='1'
ip='::ffff:192.168.56.1' status='200' method='GET' path='/api/
v1/process'
```

Chapter 4

Managing User Accounts for Cloud Servers

This chapter explains 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 described in [“Managing User Access with Teams”](#) on page 54.

For information on managing user accounts for an on-premises installation, see [“Managing User Accounts for On-Premise Servers”](#) on page 52.

Sections

Topic	Page
Overview of Cloud User Management	71
Creating Cloud User Accounts	72
Accessing Authorized Servers	76
User Account Lockout	76
Viewing and Modifying Cloud User Accounts	77
Changing Security Settings, Email Address or Full Name	77
Changing Administrator / User Status	79
Resetting API Tokens	80
Removing a User Account	81

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 already) 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 or not. 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** they belong to. 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 that you might want to restrict access to. For example, threat feeds, which are not specific to any Sensor Group, are an important tool for CB Response threat monitoring, and 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 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.

See [“Managing User Access with Teams”](#) on page 54 for information about creating and using teams.

Important

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

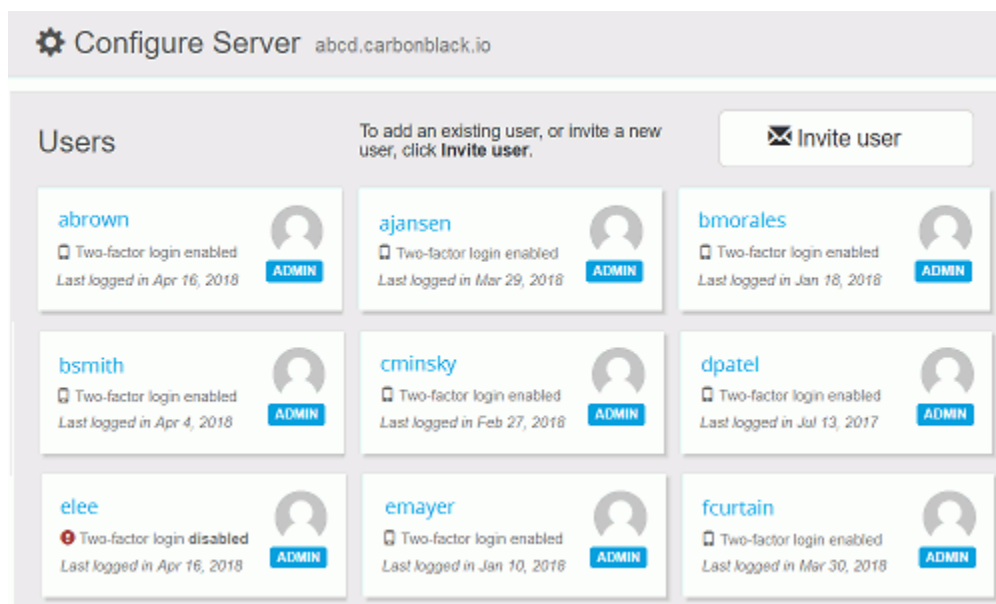
To invite a user to open a CB Response Cloud account or extend it to a new server:

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

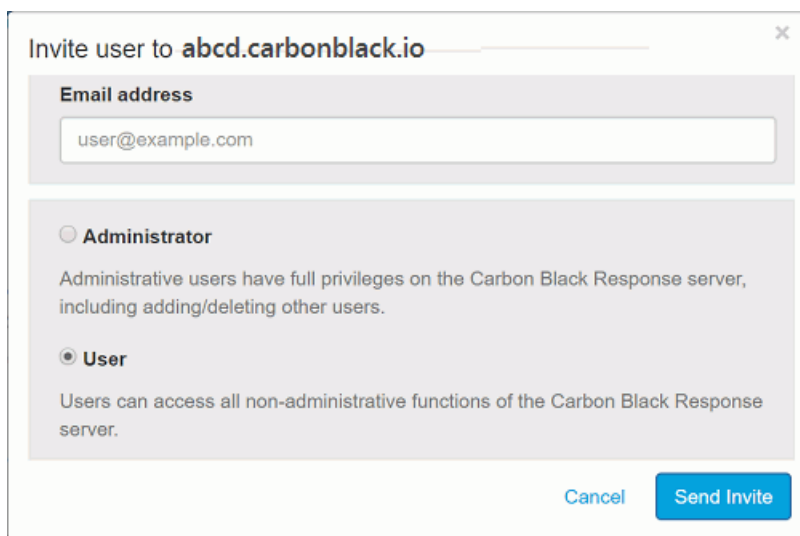
If no other administrator accounts have been created, use an administrator account provided by Carbon Black when you initiated your use of the CB Response cloud.

2. From the console navigation bar, click **Users**.

The **Users** page appears and shows the CB Response Cloud user accounts that are already authorized to access the server. This page also shows users whose invitation has expired without being an 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 in this case:
 - You can re-invite the user by double-clicking the account box and clicking the **Resend Invite** button.
 - You can 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.
 - You can 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 already, 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** as the class of this user on this server.
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**.

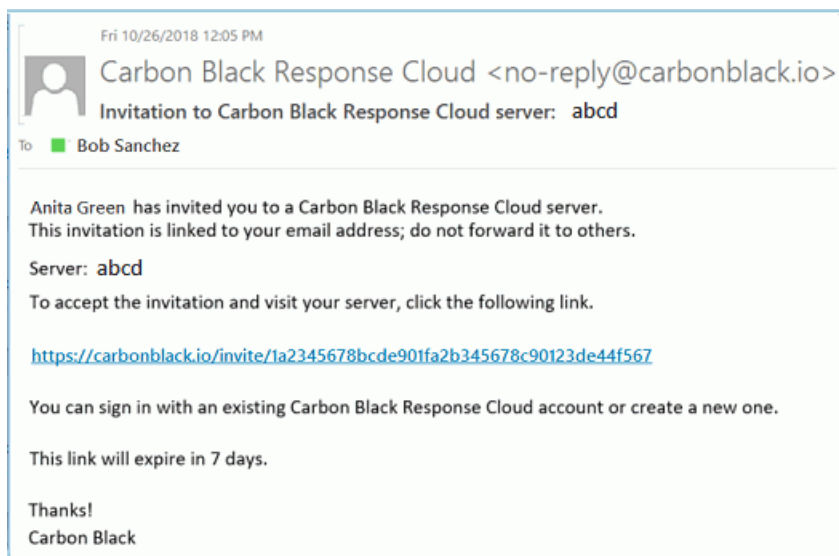
An email is sent to the user containing a link that can be used to create a new account on the server. This link also authorizes an existing user who has access to a different server to log into this one with the same account.

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 sent to the address 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, provide the following details:
 - **Username** – Choose a username for your cloud account. User names are restricted to standard, Latin alphanumeric characters without symbols and punctuation characters.
 - **First** – Enter your first name.
 - **Last** – Enter your last name.
3. Click the **Sign up** button.
You are immediately logged in to the HUD page for the server you were invited to access. In addition, 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 on 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 console 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, listing the servers to which you have access:



4. Click a servername link to access the HUD page for the server.
5. Click the **Configure** button if you want 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 which user information can be viewed or changed for CB Response Cloud users:

- **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 get to this page by clicking Users in the navigation bar on the CB Response console.
 - Clicking on the tile for any user shows 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 79 for more details.
 - An Administrator can also remove a user account via this page. See [“Removing a User Account”](#) on page 81 for more details.
- **Cloud: Account page** – Using the Account page, individual CB Response Cloud users can manage their 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 from two locations: by clicking **View profile on carbonblack.io** from the “My profile” page on the server and by choosing Account from the username menu in the cloud view.

See [“Changing Security Settings, Email Address or Full Name”](#) on page 77 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 the teams each one is on. 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 54 for more details.
- **Server: My profile page** – For the currently logged-in user, the “My profile” page shows the teams a user is on, provides access to the API Token page where the API Token may be changed, and includes a **View profile on carbonblack.io** button that opens a new browser window showing the Account page for this user on the CB Response Cloud.

You get to the My profile page through the username 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. Log into a CB Response Cloud server and then click **Users** in the navigation bar.
2. In the top right of the My Servers page, select **<your username > Account**.

Your cloud user **Account** page appears:

The screenshot shows the 'Account' page with three main sections:

- Basic Info:** Contains input fields for 'First' (jane), 'Last' (doe), and 'Username' (jdoe). A 'Save Changes' button is at the bottom. A placeholder for a profile picture is labeled 'Image from Gravatar'.
- Contact Info:** Contains an 'Email' field with 'jdoe@mycorp.com' and a 'Change Email' button.
- Security:** Contains a 'Password' field with a note 'Your password was last updated June 24, 2017.' and a 'Change Password' button. Below it, 'Two-Factor Authentication' is checked, with a note 'Two-Factor Authentication has been enabled for your account.' and a 'Disable Duo 2FA' button.

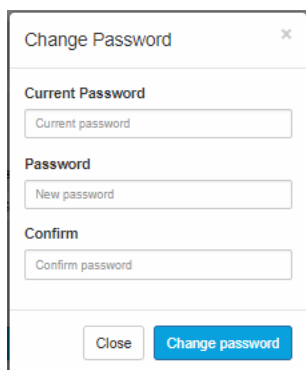
3. In the **Basic info** section, you can modify the following details:
 - **First** – Modify your first name.
 - **Last** – Modify your last name.
 - **Username** – You cannot change your username.
 - **Image** – Upload an image for your account from [Gravatar](#).
4. Click **Save Changes**.
5. In the **Contact info** section, you can change the email address 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.

The 'Change Email Address' dialog box contains the following elements:

- Title: Change Email Address (with a close 'x' button)
- Text: Enter your new email address, and we'll send a message for verification.
- Label: Email
- Input field: A text box containing 'New email'.
- Buttons: 'Close' and 'Change email'.

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**.



The image shows a 'Change Password' dialog box with a close button in the top right corner. It contains three input fields: 'Current Password' with a placeholder 'Current password', 'Password' with a placeholder 'New password', and 'Confirm' with a placeholder 'Confirm password'. At the bottom, there are two buttons: 'Close' and '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 36.

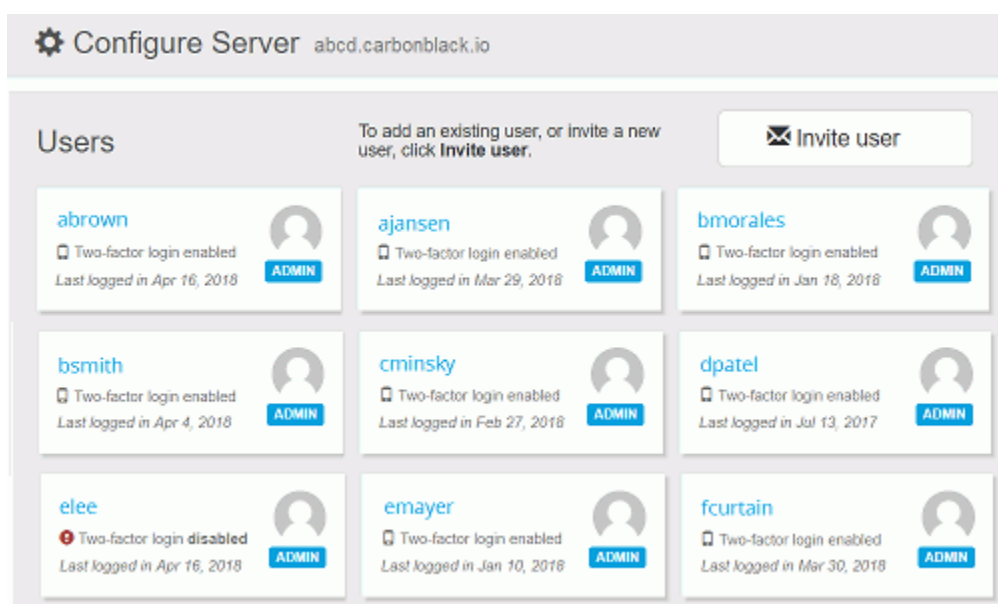
Changing Administrator / User Status

This section explains how to change a user’s status as 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. From the console navigation bar, click **Users**.

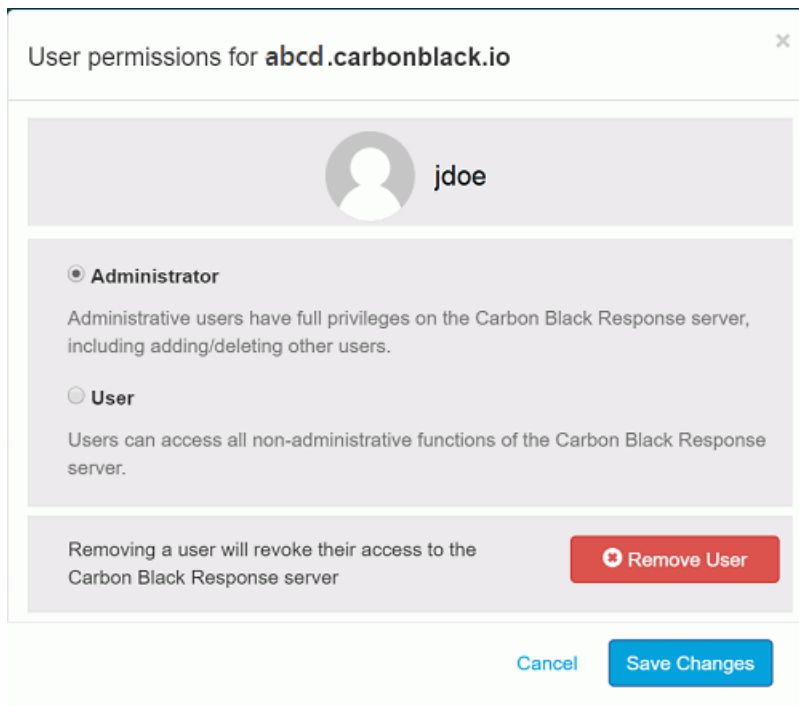
The **Users** page appears and shows the CB Response Cloud user accounts that are already authorized to access the server.




The image shows the 'Users' management interface in the 'Configure Server' console. The header includes a gear icon, the text 'Configure Server', and the URL 'abcd.carbonblack.io'. Below the header, there is a section titled 'Users' with a sub-header: 'To add an existing user, or invite a new user, click **Invite user**.' To the right of this text is an 'Invite user' button. The main area displays a grid of user cards. Each card shows a username, a status indicator (e.g., 'Two-factor login enabled' or 'Two-factor login disabled'), a last logged in date, and an 'ADMIN' badge. The users listed are: abrown, ajansen, bmorales, bsmith, cminsky, dpatel, elee, emayer, and fcurtain.

3. Click the user you want to modify.

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



User permissions for **abcd.carbonblack.io**

 **jdoe**

Administrator
Administrative users have full privileges on the Carbon Black Response server, including adding/deleting other users.

User
Users can access all non-administrative functions of the Carbon Black Response server.

Removing a user will revoke their access to the Carbon Black Response server **Remove User**

Cancel **Save Changes**

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/>

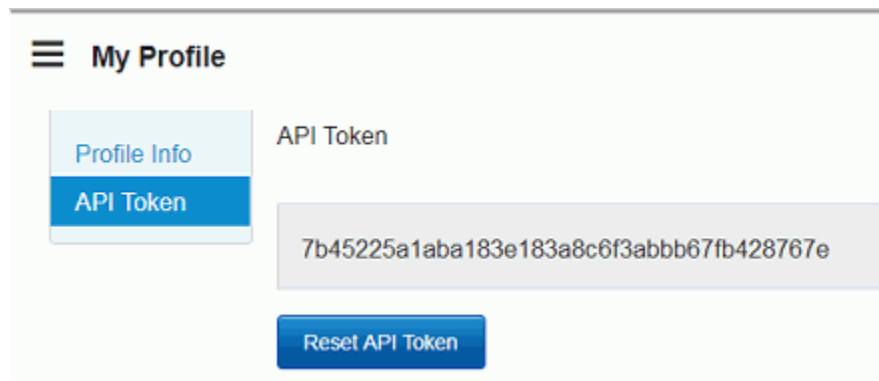
A unique API token is assigned to each CB Response user. It serves as the key authentication mechanism when making calls to the APIs. A user can reset their own API token at any time by following these procedures.

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. Select **Username > My Profile** in the main CB Response menu.
3. In the **My Profile** window, select **API Token**.
4. Click the **Reset API Token** button to reset the API token.
5. 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 on the console. See [“Viewing User Activity”](#) on page 68 for more information.


Removing a User Account

This section explains how to remove a user account from accessing the CB Response Cloud server, terminating access for that account.

To remove a CB Response Cloud user account:

1. From the console navigation bar, click **Users**.
The **Users** page displays those CB Response Cloud user accounts that are authorized to access the server.
2. Click the user you want to remove.
The **User permissions for <server name>** page appears for that user:

User permissions for **abcd.carbonblack.io**

 **jdoe**

Administrator
Administrative users have full privileges on the Carbon Black Response server, including adding/deleting other users.

User
Users can access all non-administrative functions of the Carbon Black Response server.

Removing a user will revoke their access to the Carbon Black Response server **Remove User**

Cancel **Save Changes**

3. In the User permissions dialog box, to completely remove the user account from this CB Response Cloud server, click **Remove User**.
4. When the user account has been removed, the **User permissions for <server name>** page disappears and 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 systems.

- 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 on which CB Response features are supported on sensors running in each sensor operating system.

Sections

Topic	Page
Overview of Sensor Installation	84
Supported Operating Systems and Versions	84
Installing Sensors on Windows	84
Upgrading Sensors	90
Uninstalling Windows Sensors	86
Installing Sensors on macOS Systems	87
Upgrading Sensors on macOS	88
Uninstalling Sensors on macOS	88
Installing Sensors on Linux Systems	88
Upgrading Sensors on Linux	89
Uninstalling Sensors on Linux	89

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 permissions group Local Administrators (or higher).

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](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. 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 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 containing 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 118 for information about sensor group upgrade policies.

4. For the 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 90, 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 90. You must have access to an administrator account on the macOS system to perform the upgrade.

For information on the latest sensors, visit the [Carbon Black User eXchange](#).

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 90, 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 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 in place 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** to display the Sensors page.

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 90.

Note

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

For information on the latest sensors, visit the [Carbon Black User eXchange](#).

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 90, 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 118 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** to display the Sensors page.
2. In the Groups panel, select the Sensor Group with the sensor to uninstall.
3. In the sensors list, select the check box 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 118.

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 54 for information about the roles and permissions 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	96
Monitoring Sensor Status and Activity	94
Monitoring Sensor and Server Information	97
Viewing Sensor Details	100

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 16, “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 host name to get detailed information about a particular sensor.

The Sensors Page

The Sensors page in the CB Response console provides information about sensors and the computers on which they are installed.

To access the Sensors page:

1. From the console navigation bar, select **Sensors**.

The Sensors page is organized as follows:

- Computers (endpoints, or hosts) in your environment with 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 sensors displayed, 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 96.
 3. When the list of sensors cannot fit on a single page, use controls at the bottom of the page to navigate multiple pages as follows:



Showing 1-10 of 181 Items per Page Jump to Page of 19 < 1 2 3 4 5 ... 19 >

- 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 the sensor is installed on.
IP Address	The IP address of the computer the sensor is installed on.
Status	<p>Describes the status of sensor connectivity as follows:</p> <ul style="list-style-type: none"> • Online – Sensor communicated with the CB Response server within the previous expected check-in interval. • Offline – Sensor was unable to 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. <p>If known, offline status might include one of the following reasons:</p> <ul style="list-style-type: none"> - 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. <p>If a sensor is being uninstalled, one of the following status descriptions might appear:</p> <ul style="list-style-type: none"> • 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	<p>In a clustered environment, the server ID to which a sensor sends data.</p> <p>For a standalone instance, the value is 0 (zero).</p>
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:

- On the Sensors page, do one of the following:
 - In the search box, enter characters in the name of the computers you want to find. Searching works incrementally as you type and is case-insensitive. Search results include computers with sensors installed that have a name matching the search string; if **Show uninstalled sensors** is selected, matching computers with 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: <ul style="list-style-type: none"> 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.

- 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 above the list of sensors, select 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 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 contains data only for the 20 sensors 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 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 detail page for the sensor, the message “This host has been isolated from the rest of the network” appears at the top, and **Remove isolation** is the only option for this feature on the **Actions** list.

For more information about the isolation feature, see [“Isolating an Endpoint”](#) on page 151.

- **Remove isolation** – From an isolated state, rejoins a computer to the network so 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

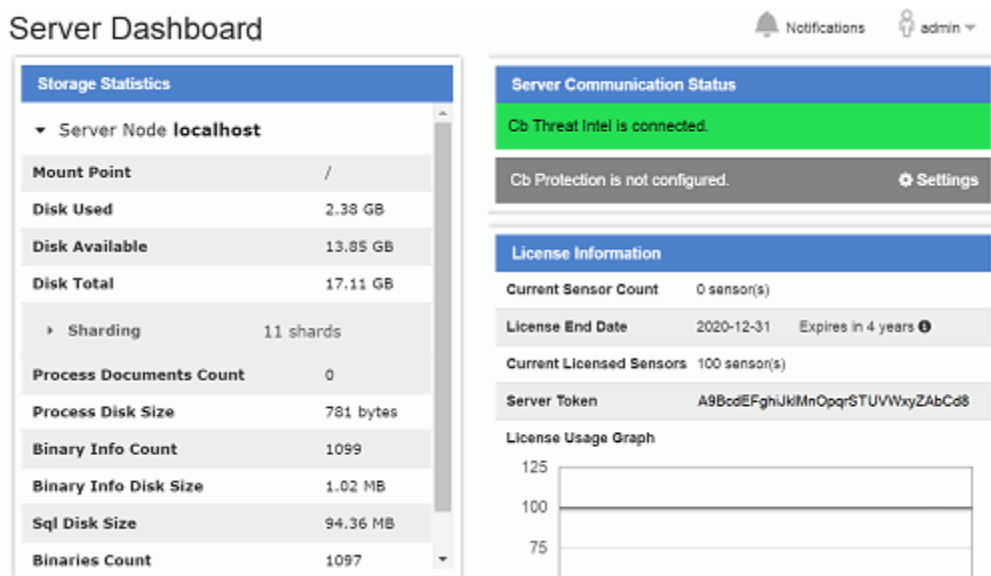
This section describes how to view this information in the Server Dashboard, and descriptions of the details that display there.

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. From the navigation bar, select **Server Dashboard** to open the Server Dashboard page:



2. Review the following **Storage Statistic** information:

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 the 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 .
Process Disk Size	The amount of disk space taken up by the process documents.
Binary Info Count	The number of binaries seen by the sensor. This is the same number as the total number of binaries on the Binary Search page .

Field	Description
Binary Info Disk Size	The total number of bytes of binary information 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, as follows:

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 the following **Server Communication Status** information:

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 following **License Information** information:

Field	Description
Current Sensor Count	Total number of unique sensors online 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.
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:

Sensors Cb Support Notifications dan2

LAPTOP-6 Related processes: 311626 | Related binaries: 6343 Go Live Actions

Sensor Vitals	
Status:	Online
Next Expected Checkin:	2018-04-30 15:04:09.415 GMT (26 seconds)
Health Score:	100
Sensor Id	78
Node Id	0
Node Address	https://sensors.abcd.carbonblack.io:443
Node Hostname	abcd.carbonblack.io
Shard Id	0
Registration Time:	2017-12-11 21:03:06.069 GMT (about 5 months)
Last Checkin:	2018-04-30 15:03:40.452 GMT (3 seconds)
Sync-mode:	No abcd
Restart Pending:	No
Uninstall Pending:	No
Sensor Version:	6.1.6.80405
Sensor Uptime:	4 days
Queued EventLog Size:	221.29 KB
Queued BinaryLog Size:	0 bytes
Network Isolation:	Not isolating

EP Agent	
EP Agent Installed	Yes
EP Agent Host Id	15

Computer Vitals	
Hostname	LAPTOP-6
OS Version	Windows 10 Enterprise, 64-bit
IP Address/MAC Info:	10.331.3.67 — 00:20:58:9f:22:f7
Computer Domain Name:	LAPTOP-6
Computer SID:	S-3-6-54-1234567890-2345678901-1234567890
Amount of RAM:	4 GB
Free Disk Space:	7.71 GB
Total Disk Space:	30.68 GB
Host Uptime:	18 days
Power State:	Running
Clock Delta:	0 seconds

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.

Process	Endpoint	Start Time	Filemods	Modloads	Netconns	Children	
mDNSResponder /usr/sbin/mDNSResponder	macosx1	Jan 23, 2018 10:17 PM GMT		5	163	2	>
launchd /sbin/launchd	macosx2	Jan 23, 2018 10:16 PM GMT	5		1	591	>
mDNSResponder /usr/sbin/mDNSResponder	macosx6	Jan 23, 2018 8:58 PM GMT			974		>

- From the Process Search (or Sensors) page, click the name of the endpoint:

Process	Endpoint	Start Time	Filemods	Modloads	Netconns	Children
mDNSResponder /usr/sbin/mDNSResponder	macosx1	Jan 23, 2018 10:17 PM GMT		5	163	2
launchd /sbin/launchd	macosx2	Jan 23, 2018 10:16 PM GMT	5		1	591
mDNSResponder /usr/sbin/mDNSResponder	macosx6	Jan 23, 2018 8:58 PM GMT			974	

Sensor Details Heading and Options

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



- The name of the host computer on which the sensor is installed.
- The number of processes related to the sensor activity and a link to the Process Search page with embedded search criteria for processes related to this sensor.
- The number of binaries that are related to the sensor activity and a link to the Binary Search page with embedded search criteria for binaries that are related to this sensor.
- Clicking **Go Live** opens an interactive live session on the sensor’s host computer so that you can execute commands in real time to help isolate or eradicate a threat. For more information about this feature, see [“Using Live Response”](#) on page 154.
- The **Actions** menu provides the following options:
 - **Move to group** – Moves the sensor to another sensor group.
 - **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.
 - **Uninstall** – Uninstalls the sensor from the host computer.
 - **Isolate** – Isolates a computer from the rest of the network, leaving only the connections that are needed for access to its sensor by the CB Response server. For more information about this feature, see [“Isolating an Endpoint”](#) on page 151.

Sensor Vitals

The Sensor Vitals panel displays the following information about sensor status:

Field	Description
Status	The connectivity status of the sensor (Online or Offline). If CB Response is not running or if there is a communication problem, the status displays as Offline .
Next Expected Checkin	The date and time (GMT) of the next time the sensor is expected to check into the CB Response server. Also gives a time estimate of either when the next checkin is expected (for example, 21 seconds) or when the last checkin was expected (for example, for an offline sensor, about 4 months ago).

Field	Description
Health Score	A numeric score (1-100) that indicates the overall health of the sensor. For example, 25 would be a poor score and 100 would be a perfect score. Also, shows a message describing sensor issues, such as "Elevated handle count" or "Elevated memory usage."
Sensor Id	The internal CB Response sensor guid of the sensor's 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.
Last Checkin	The date and time (GMT) of the last time the sensor checked into the CB Response server. Also, provides a time estimate of when the next checkin was expected (for example, 4 months ago for an offline sensor).
Sync-mode	Shows if the sensor is in the process of synchronizing with the server.
Restart Pending	Shows if the sensor host is in the process of restarting.
Uninstall Pending	Shows if the sensor is in the process of being uninstalled from the sensor host.
Sensor Version	The current version of the sensor.
Sensor Uptime	The duration of time that the sensor has been actively running on the host computer.
Queued EventLog Size	The size of the queued log (in bytes) of unprocessed events on the sensor host.
Queued BinaryLog Size	The size of the queued log (in bytes) of unprocessed binaries on the sensor host.
Network Isolation	Shows if the sensor host is being isolated from the rest of your network and the Internet. For information about isolating hosts, see "Isolating an Endpoint" on page 151.

EP Agent

The **EP Agent** panel shows if the CB Protection agent is installed, and if so, displays the ID of the CB Protection agent host.

Computer Vitals

The **Computer Vitals** panel shows the following details about the computer on which the sensor is installed:

Field	Description
Hostname	Host name of the computer on which the sensor is installed.
OS Version	Version of the operating system on the computer on which the sensor is installed.
IP Address/MAC Info	IP and MAC address of the computer on which the sensor is installed.
Computer Name	Name of the computer on which the sensor is installed. This can be the same name as the host name.
Computer SID	Unique security identifier for the computer.
Amount of RAM	Amount of RAM that is available on the computer.
Free Disk Space	Amount of free disk space that is available on the computer.
Total Disk Space	Amount of total disk space that is available on the computer.
Host Uptime	Length of time that the host computer has been running since the last boot time of the sensor. If the sensor has not rebooted since its installation, this reflects the install time. Otherwise, it is the last boot time.
Power State	Indicates whether the host computer 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.

Teams

The **Teams** panel shows the teams that have access to this sensor and the permissions that team members have. This is the same information that was defined in **Sensors > Edit Settings > Permissions**.

For information about teams and permissions, see [“Permissions Settings”](#) on page 115.

Configuration

The **Configuration** panel shows the following configuration information for the sensor:

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 110.
Server Name	URL of the CB Response server.
Sensor Upgrade Policy	Upgrade policy for this sensor. See “Upgrade Policy Settings” on page 118 for details of sensor group upgrade policies.

Sensor Activity

The **Sensor Activity** panel contains:

- The activity types that the sensor has engaged in
- The date and time of each activity (for the duration of the time that the sensor has been up and running)

Sensor Data Queued - Historical View

The **Sensor Data Queued - Historical View** panel contains:

- The number of event and binary files that are queued for processing by the sensor.
- The date and time (GMT) for each set of files.

Sensor Comm Failures

The **Sensor Comm Failures** panel shows the timestamp and failure code of communication failures between the sensor and the server.

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, then locate “13” on this page.

Sensor Driver Diagnostics

The **Sensor Driver Diagnostics** panel shows information about the sensor driver.

CB Response OS X sensors have these 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. Unloading the core kernel driver immediately can cause system instability if other (typically security) products 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 these 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 send Tamper events to the CB Response server directly. 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 **cb сенсор.ko.<kernel version>** where the **<kernel version>** is one of the currently supported kernels.
 - 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, queuing the activity up 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.

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

The Sensor Driver Diagnostics panel provides 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 with large number of network connections (e.g. 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 allow you to disable the DNS name resolution in data collection for network connections, reducing the amount of netconn traffic on these systems. This is done by configuring the following Windows registry key as shown:

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

Sensor Event Diagnostics

The **Sensor Event Diagnostics** panel shows the date and timestamp (in GMT) of sensor events, along 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

Sensor Component Status

The **Sensor Component Status** panel shows information about threads and the sensor component.

Sensor Resource Status

The **Sensor Resource Status** 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 of the performance problem. The panel shows these 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 amount of memory in use by all applications on the computer, including memory that has been temporarily paged to disk at the date and time displayed in the Timestamp field.
Handles	The number of handles in use at the date and time displayed in the Timestamp field.

Sensor Upgrade Status

The Sensor Upgrade Status panel shows details about the sensor upgrade process.

Chapter 7

Sensor Groups

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

Sections

Topic	Page
Overview of Sensor Groups	108
Create or Edit a Sensor Group	108
Moving Sensors to Another Group	118
Deleting Sensor Groups	119

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 installing sensors. When you edit a sensor group, settings are the same as when you create a sensor group. (See ["Deleting Sensor Groups"](#) on page 119.)

To create or edit a sensor group:

1. In the navigation bar of the CB Response console, click **Sensors**.
The Sensors page appears.
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 or Edit Group panel, sensor group settings are organized in sections you can expand or collapse as needed.

Note

To quickly open or close all sections of sensor group settings at once, click **Display All Sections** or **Collapse All Sections** at the top right of the Create Group or Edit Group page.

3. Navigate the sensor group sections to complete settings in the following categories:
 - **General** – See “[General Settings](#)” on page 110.
 - **Sharing** – See “[Sharing Settings](#)” on page 110.
 - **Advanced** – See “[Advanced Settings](#)” on page 113.
 - **Permissions** – See “[Permissions Settings](#)” on page 115. 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 115.
 - **Isolation Exclusions** – See “[Isolation Exclusions](#)” on page 153.
 - **Exclusions** – See “[Exclusion Settings \(OS X/macOS only\)](#)” on page 115.
 - **Upgrade Policy** – See “[Upgrade Policy Settings](#)” on page 118.

Note

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 page refreshes to show 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 your changes.

Sensor group changes take effect after 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 of the Create or Edit Group panel of the Sensors page includes the following settings:

Setting	Description
Name	The name of the sensor group; alphanumeric characters only.
Sensor Process Name	<p>(Optional, Windows-only) An alternate name for the sensor group process. The default name of the process is <code>cb.exe</code>.</p> <p>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.</p> <p>If you change the name of the sensor process, the process will run with this name instead of the default <code>cb.exe</code>. This will not change the Windows service display name, but it will change the name of the actual executable that is run.</p>
Server URL	The URL that the sensor group uses to communicate with the CB Response server. This URL is the same one used to log into the CB Response server, prefixed with <i>sensors</i> . Use HTTPS and specify the secure port in the URL.
Site Assignment	<p>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 loginname > Settings > Sites, and then assign the site to this sensor group by selecting the site in this field.</p> <p>Note: Modifying bandwidth settings for sites requires Global Administrator status for on-premise installations and Administrator status for the cloud.</p> <p>Additional information about site throttling is available in the <i>Carbon Black Enterprise Response - Operating Environment Requirements</i> on the User eXchange.</p>
Assign Server Certificate	Assign a server certificate to all sensors in the group. Only sensors that check in will receive this update. This field also includes a <i>Manage certificates</i> link that goes to the Server Certificates tab of the Settings page. See Chapter 8, 'Managing Certificates for Server-Sensor Communication' for details.

Sharing Settings

The Sharing section of the Create or Edit Group panel of the Sensors page includes the following settings.

Setting	Description
Share Binary hashes with Carbon Black	<p>Select this option to be notified of any binary flagged by Carbon Black Collective Defense Cloud.</p> <p>For more information about this choice, do the following from the Sharing section of the Create or Edit Group page:</p> <ol style="list-style-type: none"> 1. Click Share Settings to open the Sharing page for the CB Response server. 2. On the Sharing page, scroll down 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	<p>Select this option to:</p> <ul style="list-style-type: none"> • 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 participating in the community program. <p>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.</p> <p>For more information on this page, see “Process Search and Analysis” on page 175.</p>

Setting	Description
Allow Carbon Black to analyze unknown binaries	<p>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.</p> <p>For more information about this choice, do the following:</p> <ol style="list-style-type: none"> 1. Click Share Settings to open the Sharing page for the CB Response server. 2. On the Sharing page, scroll down to Endpoint Activity Sharing. 3. In the Carbon Black column next to Complete Binaries, click the current setting (Enabled, Disabled, or Partial) for a description.
Allow CB Inspection to analyze unknown binaries	<p>Select this option to detect new variants of known malware by sharing the full binary content of unknown executable files. Binaries will be uploaded and shared with Carbon Black.</p> <p>For more information about this choice, do the following:</p> <ol style="list-style-type: none"> 1. Click Share Settings to open the Sharing page for the CB Response server. 2. On the Sharing page, scroll down to Endpoint Activity Sharing. 3. In the CB Inspection column next to Complete Binaries, click the current setting (Enabled, Disabled, or Partial) for a description.

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

- Click the **Share Settings** link in the sensor group Sharing section.
- Select **Sharing Settings** from the user menu in the top right corner of the console (*loginname* > **Sharing Settings**).

For more information, see ["Data Sharing Settings"](#) on page 256.

Advanced Settings

The Advanced section of the Edit Group panel on the Sensors page includes the following settings:

Setting	Description
Sensor-side Max Disk Usage	<p>Contains two options to limit sensor disk consumption on clients by raw available space (in megabytes) or percentage of the total space available. The sensor(s) will limit the amount of space they use on clients based on the smaller of these two values:</p> <ul style="list-style-type: none"> • In the MB field, enter the maximum available space on the client, in megabytes (between 2 and 10240), that sensors can use. • In the % field, enter the maximum percentage of total disk space (between 2 and 25) on the client that sensors can use.
Filter known modloads (Windows and OS X only)	<p>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 provides a method for reducing the amount of known good events reported into the server.</p> <p>For more information on Windows modloads, see https://technet.microsoft.com/en-us/magazine/2007.09.windowsconfidential.aspx.</p>
Process Banning	<p>When selected, enables process hash bans in this group. By default, this setting is disabled and process hash bans prevent banned processes from running.</p> <p>For more information, see “Banning Process Hashes” on page 163.</p>
Tamper Detection (Windows only)	<p>When selected, the sensor identifies when attempts are made to modify the sensor's binaries, disk artifacts, or configuration.</p> <p>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 also has enhanced permission for isolating sensors.</p>
VDI Behavior Enabled	<p>When selected, 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.</p> <p>Note: VDI support must be globally enabled in order to use this feature. For more information, see “Server VDI Support” in the <i>CB Response Integration Guide</i>.</p>

Setting	Description
Retention Maximization	<p>These settings change how sensor process data that contains only modload processes or only modload and cross processes is recorded on the server.</p> <p>Minimum Retention makes this data more easily searchable but leaves a bigger footprint and can lead to a reduction in data retention time.</p> <p>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.</p> <ul style="list-style-type: none"> • 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. <p>Note: Recommended and Maximum Retention may result in false positives in the results of cmdline searches. See “Retention Maximization and cmdline Searches” on page 242 for more detail.</p> <p>Note: This setting was called “Data Suppression Level” in pre-6.5 versions of CB Response.</p>
Alerts Critical Severity Level	<p>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.</p> <p>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:</p> <ul style="list-style-type: none"> • 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 with 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. <p>For more information about alerts, see Chapter 17, “Console and Email Alerts”.</p> <p>For more information about threat intelligence feed scores, see “Threat Intelligence Feeds” on page 251.</p>

Permissions Settings

In the Permissions section of the Edit Group panel on the Sensors page, you define user team permissions for sensors groups.

Available permission levels are as follows:

- **No Access** – When users in a team try to access or view details on a host in this sensor group, the system generates an HTTP 405 response that says “The method you are using to access the file is not allowed.”
- **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.
- **Analyst** – Users can configure the sensor host and group details.

For information about user teams and access levels, see [“Managing User Access with Teams”](#) on page 54.

Event Collection Settings

In the Event Collection section of the Create or Edit Group panel on the Sensors page, you can define which types of events for CB Response 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 CB Response sensor to record the user name associated with each running process. This associates endpoint activity with the operating system user account.
- **Cross process events** – Enables the CB Response sensor to record instances when a process crosses the security boundary of another process. While 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 or Edit Group panel on the Sensors page. This Exclusions section allows you to define paths on OS X/macOS systems and customize event collection at those paths to improve performance or just to eliminate unnecessary data. With this feature, you can specify that, for example, 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 (for standalones) or cluster to make the new setting effective:

- For standalone server:

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

```
sudo sudo cbcluster stop
```

(...wait for all the nodes to shut down, and then...)

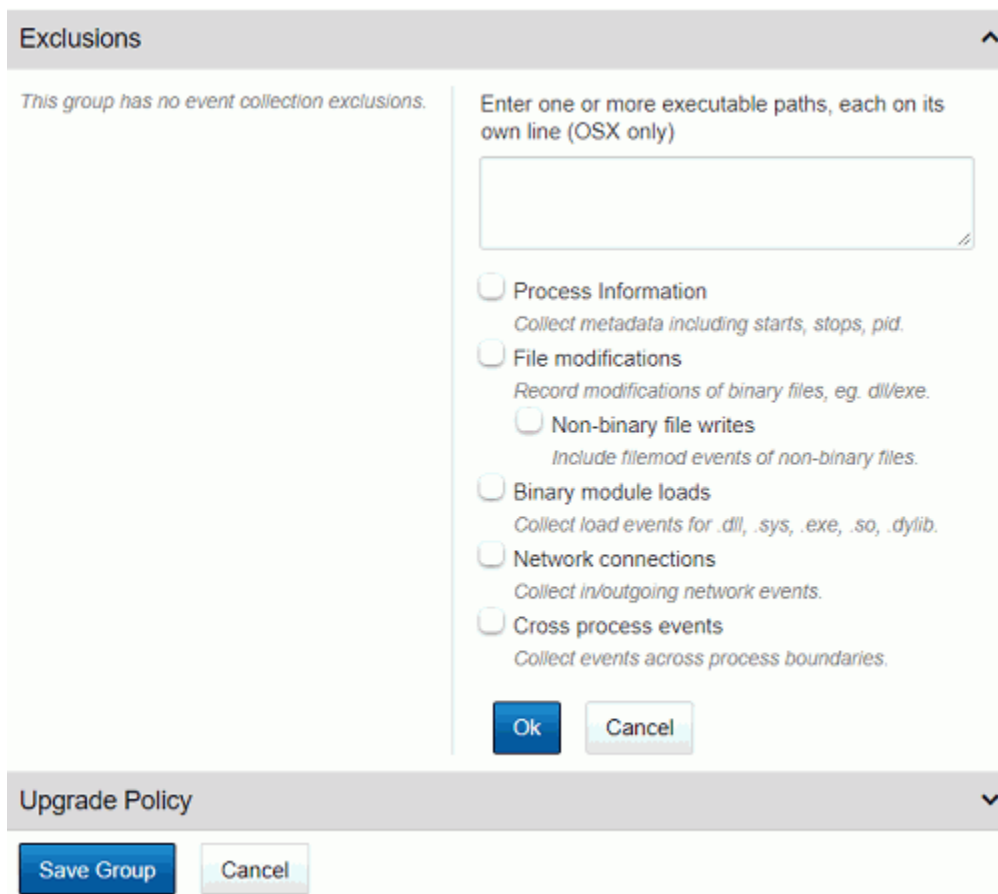
```
sudo sudo cbcluster start
```

Creating Exclusions

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

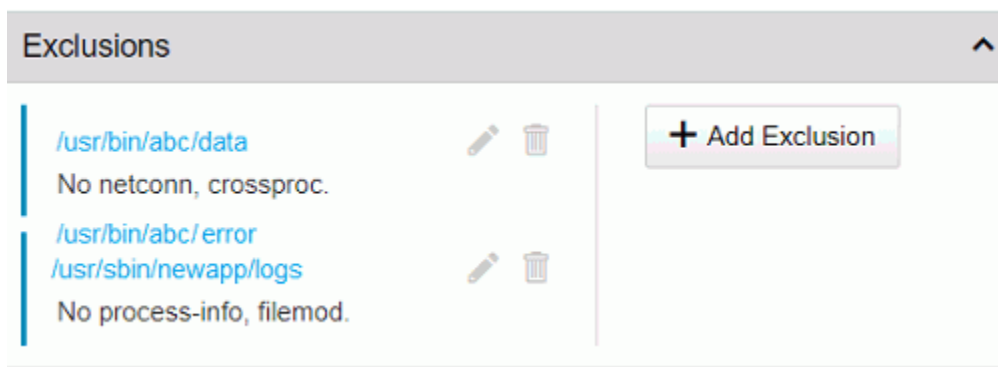
To create an OS X/macOS event collection exclusion for a sensor group:

1. In the navigation bar of the CB Response console, click **Sensors**.
The Sensors page appears.
2. In the Groups panel of the Sensors page, click the gear icon(⚙) next to the Sensor Group for which you want to create exclusions.
The Edit Group panel appears.
3. Click the **Exclusions** bar to expand that panel.
4. Click the **Add Exclusion** button.
The Exclusion configuration fields are exposed.



5. Enter the path(s) you want to affect with this exclusion in the box in the upper right of the panel. Put each path on a new line. You must use a complete path, without wildcards.
6. Check the box next to each type of information you do not want to collect for the specified paths.
7. When you have finished specifying the types of data that will not be collected, click the **Ok** button.

The Exclusion is saved and displayed in the panel.



8. If you want to add an additional exclusion, click the Add Exclusion button again, provide the path(s) and data exclusions for those paths, and click **Ok** when finished.
9. When you have finished creating exclusions, click the **Save Group** button.

You can edit or delete any exclusion listed on the Exclusions panel using the icons to the left of its name.

Upgrade Policy Settings

The Upgrade Policy section of the Create or Edit Group panel on the Sensors page contains options to set the policy for upgrading installed sensors in the group, for the Windows, OS X, and Linux platforms.

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** – Install a specific version for all sensors in a group. This keeps 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 or vetted.

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. In the navigation bar, select **Sensors** to display the Sensors page.
2. In the Groups panel, click to select the sensor group containing the sensors you want to move.
3. In the Sensors panel, select the check box(es) next to the sensor(s) that you want to move.
4. Click **Actions > Move to group**.
The Move Sensors Confirmation dialog box appears.
5. From the drop-down list, select the sensor group to which you want to move the selected sensor(s), and then click **Okay**.

The selected sensor(s) 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'](#) for more on this topic.

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 will no longer have access to sensors that belong to the sensor group.

To delete sensor groups:

1. In the navigation bar, select **Sensors** to display the Sensors page.
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** to remove the sensor group from the list.

Chapter 8

Managing Certificates for Server-Sensor Communication

This chapter describes how CB Response uses HTTPS and TLS to secure communication between endpoints and the server, as well as for two-way authorization between the two. 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	121
Server-Sensor Certificate Requirements	122
How CB Response Supports Multiple Certificates	123
Managing Certificates on the Server	126
Viewing Certificate Information in the Console	126
Substituting a Legacy Certificate during Server Installation	127
Adding Certificates through the Console	128
Choosing a Validation Option	129
Changing the Expiration Notification Period	130
Deleting Certificates	131
Upgrades from Previous Server Releases	131
Assigning Certificates to Sensor Groups	132
Sensor Support for Certificate Management	133

TLS Server Certificate Management Overview

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

Prior to server version 6.4.0, CB Response established the trust between endpoints and the server by using “certificate pinning,” 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 (i.e. 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 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, reducing the exposure to a compromised server certificate. Also, you can 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 will be disabled.

See [“Sensor Support for Certificate Management”](#) on page 133 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 used seamlessly in a clustered environment.

In addition to allowing you to provide your own certificates, CB Response provides new certificate visibility features that can be useful for both 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 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 add stricter validation methods. A certificate using standard validation continues to allow sensor and server to communicate even after it expires while 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 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 swap out certificates, for example if one is about to expire.
- **View the certificate for a sensor** – The Sensors page shows the server certificate 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 used for sensor communications must meet the following requirements:

- The files you provide must be valid certificate and key files (i.e. 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 may fail silently.

- 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 example below shows how you could 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 for the minions.

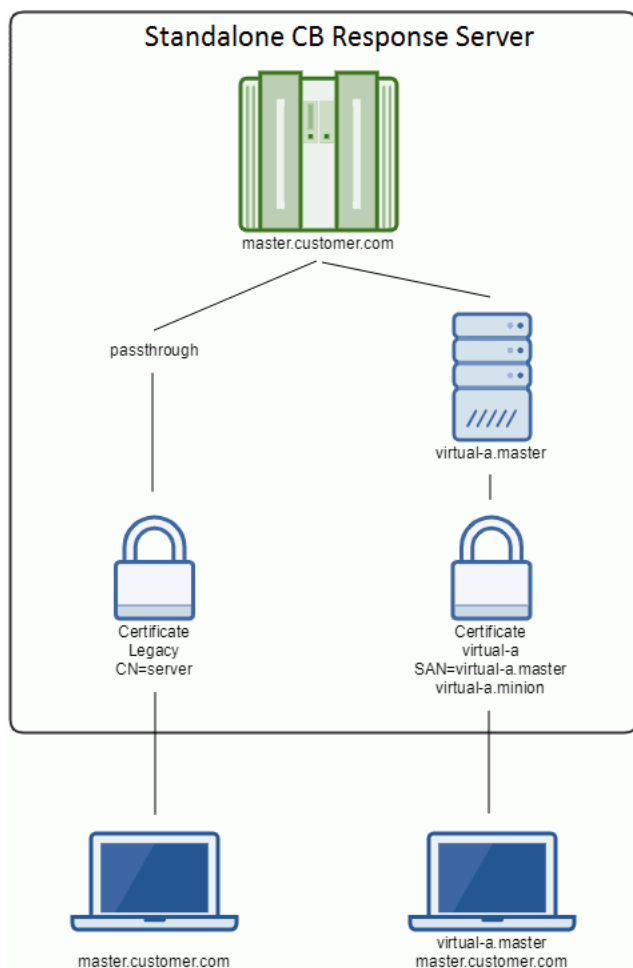
```
Certificate A
  CN=<something>
  SAN.DNS.1=virtual-a.master
  SAN.DNS.2=virtual-a.minion

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

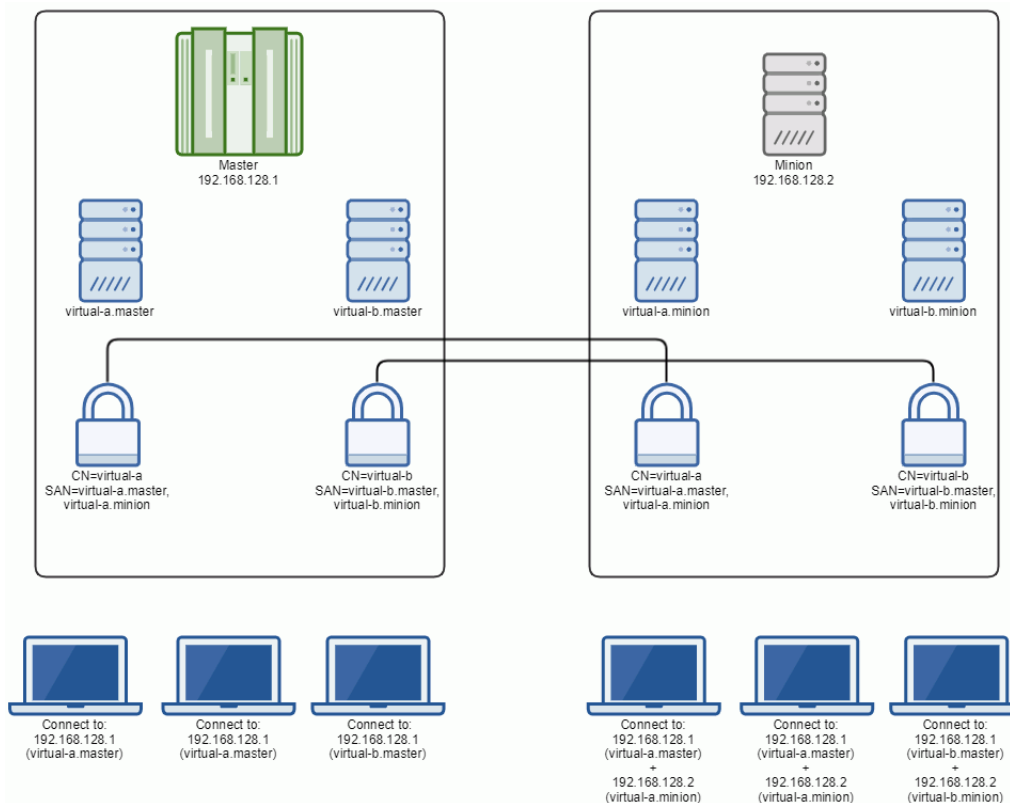
On an upgrade to CB Response 6.4.0, the previously used self-signed certificate is named "Legacy" and would be served via the default server access route (using no virtual servers), which is supported for all sensor versions and configurations. New server installations also include a "Legacy" certificate.

Note

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

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 work 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 copied to minions automatically.
- 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 have to have two distinct SAN DNS entries.
- Sensors need to distinguish only **two** different virtual addresses to be able to communicate to the servers (master and minion). Because resolution of those virtual names happens internally and that 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.

If you want to use server certificates other than the default legacy certificate, you have two opportunities to do that:

- 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 127 for details.
- Once 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 128 for details.

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

- Choose the validation methods that sensors use for certificates. See [“Choosing a Validation Option”](#) on page 129 for details.
- 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 132 for details.

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 used 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 available on the current server. It also shows the validation method being used for these certificates. See [“Choosing a Validation Option”](#) on page 129 for more about certificate validation methods.

Settings

Server certificate validation mode

- Standard validation Certificate pinning only. Requires matching sensor and server certificate.
- Strict certificate validation Standard validation as well as additionally requiring validation against a trusted Certificate Authority on Sensors. This also includes checking if the certificate has expired.

Save Changes

Server certificates Notify me 30 days before a certificate expires + Add certificate

NAME	SENSORS	THUMBPRINT	SAN	EXPIRY DATE	ADDED BY	DATE ADDED	ACTIONS
Legacy	2	12:AB:03:98...		2029-05-06	System	a month ago	Actions
cert-a	51	3D:05:BB:13...	virtual-a.master, virtual-a.minion	2024-04-29	admin	20 days ago	Actions
cert-b	14	5A:B2:02:76...	virtual-b.master, virtual-b.minion	2029-04-30	admin	4 days ago	Actions
cert-c	8	10:E1:83:42...	virtual-c.master, virtual-c.minion	2029-04-30	admin	4 days ago	Actions

To view the available certificates on a server:

1. In the upper right of the console, select **Settings** on the username menu.
2. On the Settings page, click **Server Certificates** in the left panel.

The table of certificates is displayed.

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 suitable for use with the standard pinning validation method. By default, this is a certificate produce by the server itself. 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 122.

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 not, the default legacy certificate will be used instead, 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, just with 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 128 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 where 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 one or more certificates to the CB Response Server using the console for use 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 122.

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

1. In the upper right of the console, select **username > Settings**.
2. On the Settings page, click **Server Certificates** in the left panel.
3. Click the **Add certificate** button.
4. In the Add certificate dialog, provide a unique name for the certificate that will help you identify its purpose when listed in CB Response (use 50 or fewer alphanumeric characters without spaces).

Add certificate X

Name

Both X509 certificate and private key must be in ASCII PEM encoded format without encryption. Two different SAN.DNS entries are required. For further information please consult User Guide or contact Carbon Black support.

Upload certificate

Choose File No file chosen

Upload private key

Choose File No file chosen

Add Cancel

5. Under *Upload certificate*, click **Choose File** and provide the path to a certificate file meeting the requirements described in [“Server-Sensor Certificate Requirements”](#) on page 122.
6. Under *Upload private key*, click **Choose File** and provide the path to the ASCII PEM-encoded, unencrypted key file for this certificate.
7. 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

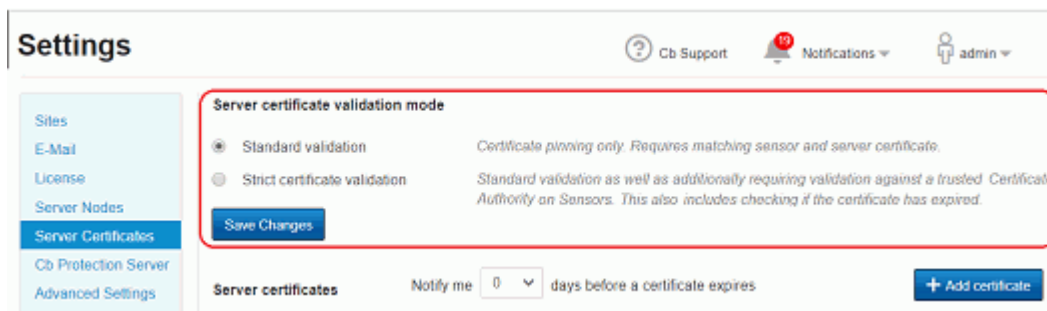
You can choose one of two validation methods that sensors use for the server certificates used to secure server-sensor communication. The validation method can be set through the console method described below or by providing a value in the `cb.conf` file for `CbServerSSLCertStrictCheck`, in which case 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 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 133 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. In the upper right of the console, select username > Settings.
2. On the Settings page, click Server Certificates in the left panel.
3. Under Server certificate validation mode, there are two radio buttons:
 - **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 immediately under the radio buttons, and if you are certain you want to make this change, click **Confirm** on the confirmation dialog.

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

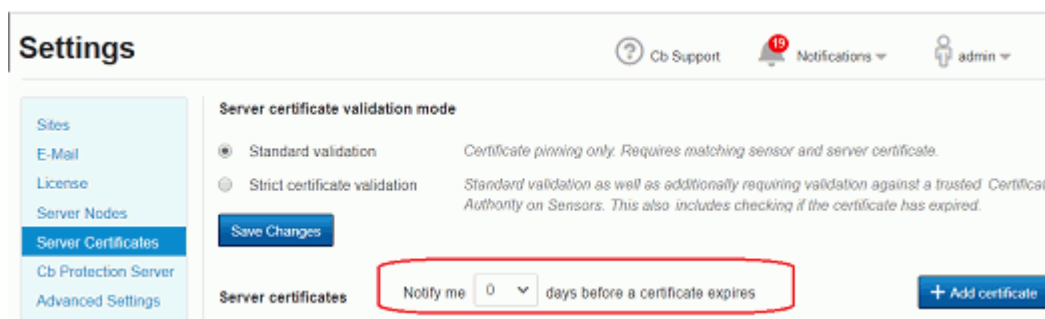
Caution: As the confirmation dialog states, changing validation method has the potential to disable communication between sensor and server. Make sure you have configured certificates properly before changing this setting, especially if you are changing to strict validation.

Changing the Expiration Notification Period

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

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. In the upper right of the console, select **username** > **Settings**.
2. On the Settings page, click **Server Certificates** in the left panel.
3. In the *Notify me* drop-down above the table, choose the number of days in advance you would like to be warned about expiration of any of your 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 created during server installation.

To delete a certificate from a server:

1. In the upper right of the console, select **username** > **Settings**.
2. On the Settings page, click **Server Certificates** in the left panel.
3. On the Server Certificates page, check that the certificate you want to delete does not have any sensors using it. If the certificate is not in use, choose **Delete** on the Actions menu for that certificate.
4. In the confirmation dialog, if you are sure you want delete this certificate, click **Delete**.

Caution: Once 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 expecting 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 immediately after the upgrade.

Once the upgrade is complete, you can begin implementing 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. Once 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 is not able to connect with the new certificate, it reverts to previous sensor certificate as a fall-back scenario.

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 133.

In clustered environments, certificate changes are propagated automatically 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 you have mapped into different sensor groups. You can use different server certificates to authenticate the connections between the CB Response Server and different sensor groups, reducing the exposure to a compromised server certificate. This also allows you to manage certificate expiration on a per sensor group basis.

You might 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. Once 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.

The screenshot shows the 'Edit Group' configuration page. The 'General' section is expanded, showing several fields: 'Name' (Admin Group), 'Sensor Process Name' (Legacy), 'Server URL' (cert-a, cert-b, cert-c, cert-d), and 'Site Assignment' (cert-d). The 'Assign Server Certificate' dropdown menu is highlighted with a red box, showing 'cert-d' selected. The page also includes a 'Display All Sections' link in the top right corner.

To change the server certificate for one sensor group:

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

5. 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. In the upper right of the console, select **username > Settings**.
2. On the Settings page, click **Server Certificates** in the left panel.
3. On the Actions menu for the certificate you want to delete, choose **Assign to all sensor groups**.



4. In the confirmation box, if you are certain you want to apply this to all sensor groups, 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 choose 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 one on the server. If you choose Strict validation, the exact hash match is still required, plus additional validation criteria that vary by platform. The table below shows the different validation criteria that are available for the sensor versions each platform.

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

- **OS X (macOS) sensor 6.2.5** – This **supports** the new certificate management features and handles strict validation as shown in the table below.
- **Windows sensor 6.2.2** – This **does not support** certificate management. However, version 6.2.3, which does supporting certificate management, is expect to be released shortly after server 6.4.0. See the table for strict validation details.

Windows XP and Windows Server 2003 will not support TLS certificate swap, regardless of the CB Response Sensor version.

- **Linux sensor 6.1.10** – This **does not support** certificate management but continues to use the default “Legacy” certificate. Monitor the Carbon Black User Exchange for 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+ (planned release June 2019)
Exact certificate match (certificate pinning)	Yes	Yes
Expiration date	Yes	Yes
Certificate validation chain	-	Yes
Hostname matches (SAN=)	-	Yes
Revocation check	-	-
Key Usage is Server Auth (1.3.6.1.5.5.7.3.1)	-	Yes

Upgrading to Sensors that Allow Certificate Management

If you want to use the certificate management features of CB Response and need to upgrade your sensors to a version 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. Once the sensors are updated, you can then 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 on different OS platforms. Other information that could be useful for troubleshooting appears in [Appendix B, “Sensor Health Score Messages”](#).

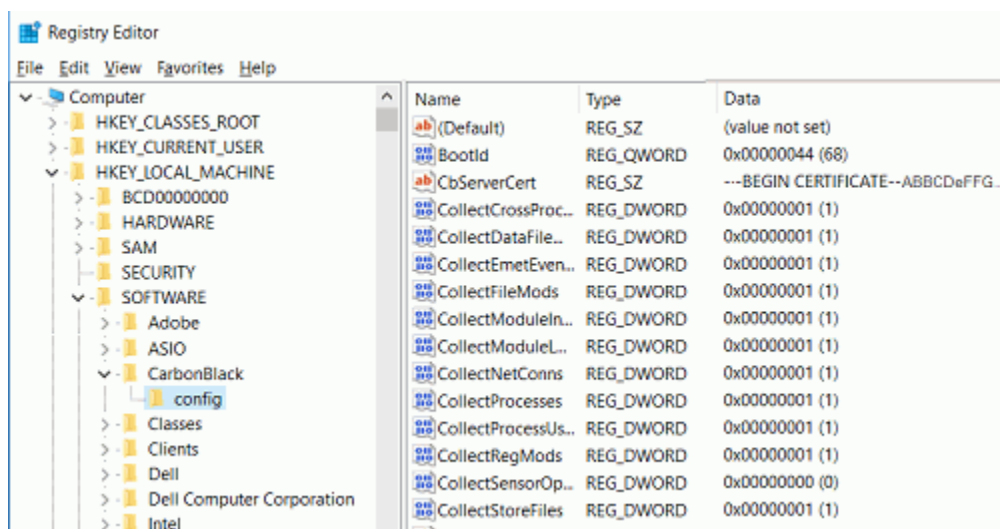
Sections

Topic	Page
Troubleshooting Windows Sensor Installations	136
Troubleshooting Linux Sensor Installations	139
Troubleshooting OSX Sensor Installations	143
Diagnostic Uploads Utility	145

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 and 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 obtain 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 – Triggers 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 attempt to the server is 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 triggered attempt does not occur until after the current operation is complete.
 - b. 201 – Triggers 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	Shows top-level event logging statistics.
MachineStatistics.log	Shows general system, process, and kernel statistics.
ModuleInfo.log	Shows internal module statistics.
NetConnEvents.log	Shows network event logging statistics. Note: To reduce netconn traffic for systems with a large number of network connections, see “Reducing the Impact of Netconn Data Collection (Windows)” on page 106.
RawEventStats.log	Shows 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`

```

Administrator: Command Prompt
C:\Windows\CarbonBlack>dir
Volume in drive C is OSDisk
Volume Serial Number is 3C9C-F25D

Directory of C:\Windows\CarbonBlack

05/20/2013  04:28 PM  <DIR>          .
05/20/2013  04:28 PM  <DIR>          ..
05/15/2013  01:44 PM           3,052,536  cb.exe
05/20/2013  04:28 PM  <DIR>          DebugLogs
05/20/2013  04:28 PM  <DIR>          eventlogs
05/17/2013  04:32 PM  <DIR>          InstallLogs
05/20/2013  04:28 PM           0 Sensor.LOG
05/15/2013  04:16 PM           47 SensorUpgrade.LOG
05/17/2013  04:32 PM       167,872  uninst.exe
05/15/2013  04:15 PM  <DIR>          upgrade
                4 File(s)      3,220,455 bytes
                6 Dir(s)  620,578,656,256 bytes free

C:\Windows\CarbonBlack>sc control carbonblack 201

SERVICE_NAME: carbonblack
        TYPE               : 10  WIN32_OWN_PROCESS
        STATE                : 4   RUNNING
                        (STOPPABLE, NOT_PAUSABLE, ACCEPTS_SHUTDOWN)
        WIN32_EXIT_CODE       : 0   (0x0)
        SERVICE_EXIT_CODE   : 0   (0x0)
        CHECKPOINT           : 0x0
        WAIT_HINT            : 0x0

C:\Windows\CarbonBlack>dir Diagnostics\SensorComms.log
Volume in drive C is OSDisk
Volume Serial Number is 3C9C-F25D

Directory of C:\Windows\CarbonBlack\Diagnostics

05/20/2013  04:28 PM           822 SensorComms.log
                1 File(s)      822 bytes
                0 Dir(s)  620,578,500,608 bytes free

C:\Windows\CarbonBlack>_

```

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 troubleshoot errors that could have been caused during Linux sensor installation.

General Logging

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

- For a version 4.x to 5.0 Linux sensor:

```
/var/log/cb/sensor/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 will result in the generation of log files for `WARNING` and `ERROR` levels:

- `/var/log/cb/sensor/cbdaemon.WARNING`
- `/var/log/cb/sensor/cbdaemon.ERROR`

- For a version 5.1 or higher 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 will result 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
```

Installation Verification

The following is a manifest of installed files:

Path	Additional Notes
<code>/etc/init.d/cbdaemon</code>	Sensor daemon script
<code>/usr/sbin/cbdaemon</code>	Sensor daemon executable
<code>/lib/modules/\$(uname -r)/kernel/lib/cbsensor.ko</code>	Sensor kernel module
<code>/etc/sysconfig/modules/cbsensor.modules</code>	Kernel autostart file
<code>/var/lib/cb/config</code>	Settings file
<code>/var/lib/cb/sensorsettings.ini</code>	Settings file

- 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-v5.2.0.60603-1.x86_64
```

Note

The version number will vary depending 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	CURL error codes can be found at 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`

Determine Server URL

To determine the server URL used by the sensor, follow the instructions in “[Sensor Communication History](#)” on page 141 to create a communication log and dump the contents of the generated log file. The server URL appears at the top.

Trigger 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.
- `g_eventFilter` – Controls which event types are generated.

These arguments can be passed as follows:

- In the `/etc/sysconfig/modules/cbsensor.modules` file
- By issuing the 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
#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

```

^^ are the available levels.

Simply OR them together to create the log level mask that you want.

Daemon Debug Options

You can output debug logging a daemon by restarting it with debug logging enabled. There are several ways to do this, all of which require you to stop and start the daemon.

To output debug logging by restarting the daemon (option 1):

1. Stop the daemon by issuing this command:

```
sudo service cbdaemon stop
```

2. Edit the `/etc/rc.d/rc3.d/S25cbdaemon` file to change `daemon $daemon` to `daemon $daemon info`.

```

start () {
    check
    [ -f $config ] || exit 6

    check_kernel

    echo -n $"Starting $prog: "

    # start daemon
    daemon $DAEMON
    RETVAL=$?
    echo
    [ $RETVAL = 0 ] && touch $lockfile

    return 0
}

```

3. Start the daemon by issuing this command:

```
sudo service cbdaemon start
```

To output debug logging by restarting the daemon (option 2):

1. Stop the cbdaemon by issuing this command:

```
sudo service cbdaemon stop
```

2. Run the daemon directly from the prompt by issuing this command:

```
sudo /usr/sbin/cbdaemon info
```

To output debug logging by restarting the daemon (option 3):

1. Stop the daemon by issuing this command:

```
sudo service cbdaemon stop
```

2. Run the daemon in the foreground by issuing this command:

```
sudo /usr/sbin/cbdaemon debug info
```

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/cbsensor/sensordiag.sh
```

Troubleshooting OSX Sensor Installations

This section describes places to check to troubleshoot errors that could have been caused during OSX sensor installation.

Installation Verification

The following is a manifest of files that should be installed:

Path	Additional Notes
/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. Trigger 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 <code>errno</code>
25	HTTP errors	HTTP error code
200	Curl errors	CURL error codes can be found at 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 140 to create a communication log and dump the contents of the generated log file. The server URL appears at the top.

Trigger an Immediate Checkin to the Server

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

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

Trigger 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 is 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 **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 enddate 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 in the CB Response console. There are three options for Crash data uploads and two for Diagnostics and Environment data:

- **Disabled** – The default setting for each diagnostics type. If this is set, neither automatic crash file uploads nor manual triggering of any uploads is available.
- **Manual** – If this is set, you manually trigger the sensordiag utility tool via the command line – no data is automatically uploaded. Running the tool collects and uploads the data type specified by a command-line switch (Crash, Diagnostic, Environment or any combination).
- **Automatic** – If this is set, 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, choosing 'Automatic' also enables the sensordiag utility so that you can manually collect and upload these data types.

Allow Upload of Sensor Diagnostics Data

Enable sensors to collect diagnostics data and upload to Carbon Black for troubleshooting. Collected data includes application logs, system hardware configuration and application configuration information from deployed sensors. Data collected is limited to technical information about the system software and hardware. Application data, binaries and user data from the systems is never included.

Disabled
Do not upload sensor diagnostics data to Carbon Black.

Manual
Upload diagnostics data manually by using a utility installed on the sensor.

Automatic
Upload diagnostics data automatically when fault conditions are detected on the sensor.

[More Detail >](#)

File Transfer and Security

Each collection of sensor diagnostic files is packaged as a zip file on the sensor before upload 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 accessed by authorized Carbon Black representatives.

To avoid sending oversized files over a http request, the uploaded file size is limited to 5MB. Files larger than that will be split into multiple files during transit. For example, a 7MB file will be uploaded in chunks of 5MB and 2MB.

Files are uploaded to the CB Response server on first come first server basis. Only one file can be uploaded to a server at a time. If more sensors check in for an upload while an upload is in progress, then 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.

Once 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

(1)- File names and file paths (which include the file name) as well as IP Addresses of machines connected to will appear in CbOxSensorService log files.

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

(3)- Device names may appear in the log along with IP Addresses (1).

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

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

(6)- Host name is shown in system profiler information.

Chapter 10

Responding to Endpoint Incidents

This chapter describes how to respond to endpoint incidents using CB Response features such as isolating endpoints, using CB Live Response, and banning process hashes.

Sections

Topic	Page
Overview of Incident Response	150
Isolating an Endpoint	151
Using Live Response	154
Extending Live Response	162
Live Response Activity Logging and Downloads	162
Banning Process Hashes	163

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. For example, you can continually monitor the issue or re-image the affected systems.

CB Response provides the following methods for responding to threats directly from the console:

- **Endpoint Isolation** – You can isolate a computer from the rest of the network, leaving only connections needed for access to its sensor by the CB Response server.
- **CB Response Live Response (Live Response)** – You can isolate a computer from the rest of the network. After being isolated, the computer will only have network access to support communicating with the CB Response server.
- **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, if you find a malicious process currently running on a sensor-managed computer, you can isolate that computer immediately to prevent the spread of the problem. Then, Live Response to end the process and perform any other file removal or needed repairs.

On the other hand, if the incident CB Response identified is not ongoing, isolation may not be necessary. In that case, you could use Live Response to remediate or further investigate it on affected machines, 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 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 52 or “[Managing User Accounts for Cloud Servers](#)” on page 70 for more information about user roles and privileges.

Isolating an Endpoint

You can isolate one or more Windows, OS X, or Linux 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 153):

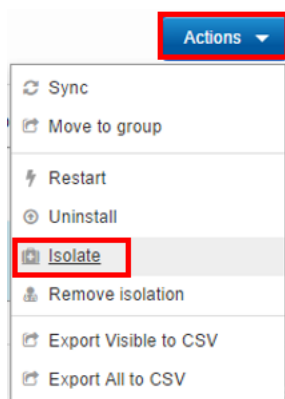
- The CB Response server can communicate with an isolated computer.
- To allow the sensor to communicate with the CB Response server, ARP, DNS, and DHCP services remain operational on the sensor’s host. (For Windows operating systems prior to Vista, ICMP (for example, ping) will remain operational.)
- 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 one or more endpoints from the network:

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 you want to isolate.
2. In the navigation bar, select **Sensors**.
3. On the Sensors page, check the box next to each endpoint you want to isolate.
4. From the **Actions** drop-down list, select **Isolate**:



5. In the confirmation box, click **OK** to confirm that you want to isolate these computers.

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

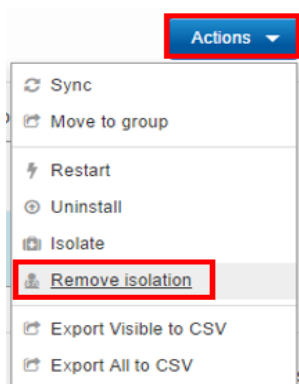
When you designate an endpoint for isolation, its status on the server first moves into in “isolation configured” state waiting for its next check-in. Because of this, there could be a period of several minutes before the endpoint is actually isolated. When it checks in, the server tells the sensor to 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 any remediation steps you plan to take on the systems that show malicious activity. For example, you might use Live Response to investigate or modify a computer. When you are finished, restore connectivity to the endpoints you isolated.

To end network isolation for one or more 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 you want to restore.
2. In the navigation bar, select **Sensors**.
3. On the Sensors page, check the box next the endpoints for which you want to restore network connectivity.
4. From the **Actions** drop-down list, select **Remove isolation**.



5. In the confirmation dialog, if you are certain you want to restore network connectivity, click **OK**.

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 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. In 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).

Isolation Exclusions

This group has no isolation exclusions.

Brief Description *

Test server

IP address or URL to exclude *

192.168.2.1

Enable this exclusion

Ok Cancel

6. Select **Enable this exclusion** and then 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 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. 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. Beginning with version 6.3.0, there are two ways to enable and disable it:

- 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.*
- For on-premises servers, you can edit the `CbLREnabled` setting in the `cb.conf` file to specify that Live Response is enabled (True) or disabled (False). This fixes the state of Live Response so that it cannot be modified through the console user interface.

Once 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 154.
- **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 161.

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. Navigate to the **Sensor Details** page for the computer you want to access by double-clicking on the computer name wherever it appears as a link. If you are not already on a page that shows the computer name:
 - a. In the navigation bar, select **Sensors**.
 - b. On the Sensors page, double-click the name of the computer.
For more information on how to access the Sensor Details page, see [“Viewing Sensor Details”](#).
3. On the Sensor Details page, click **Go Live**.

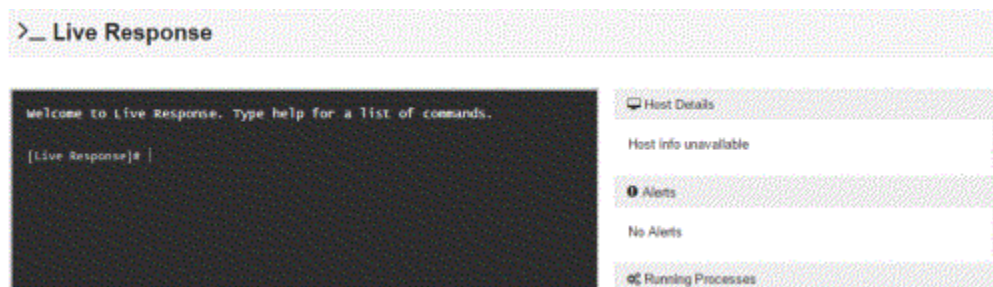


The **Live Response** page appears with a command window on the left and an information panel on the right. The command window prompt shows the name of the host and the current directory in which Live Response is active. The information panel includes:

- **Host Details**
- **Alerts** related to the host
- **Running Processes** on the host

There is a status indicator (dot) and message immediately above the command window. The dots have 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.
- **Grey** – A session cannot be established with the sensor—for example, 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 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 <code>cls</code> 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).

Command	Description
dir	Return a list of files in the current directory or the specified directory if it is added to the command, (for example, <code>dir c:\temp</code> or <code>dir /tmp</code>)
drives	List the drives on the remote host. This is for Windows only.
exec[processpath]	<p>Execute a background process specified in the <code>processpath</code> argument on the current remote host. By default, process execution returns immediately and output is to <code>stdout</code> and <code>stderr</code>.</p> <p>Options may be combined:</p> <ul style="list-style-type: none"> • <code>exec -o outputfile processpath</code> – Redirect the process output to the specified remote file, which you can download. • <code>exec -w processpath</code> – Wait for the process to exit before returning. <p>You could combine the options as shown in the example below to execute and capture the output from a script:</p> <pre>exec -o c:\output.txt -w c:\scripts\some_script.cmd</pre> <p>You must provide the full path to the process for the <code>processpath</code> argument. For example:</p> <pre>c:\windows\system32\notepad.exe</pre>
execfg <i>[processpath]</i>	<p>Execute a process on the remote host and return <code>stdout/stderr</code>.</p> <p>For example, this command prints the output of <code>ipconfig</code> to the screen:</p> <pre>execfg c:\windows\system32\ipconfig /all</pre>
files <i>[-s session]</i> <i>[action] [option]</i>	<p>Perform actions over cache-stored session files.</p> <p>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.</p> <p>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 <code>-s</code>.</p> <p>A list of sessions is available through the <code>sessions</code> command. If attached to a sensor, the current session is assumed unless otherwise specified.</p> <p>There are three available actions:</p> <ul style="list-style-type: none"> • <code>list</code> – List all the cached files that are available in the specified session by file ID. • <code>get [id]</code> – Get the file <code>[id]</code> from the cache. • <code>delete [id]</code> – Remove the file <code>[id]</code> from the cache.

Command	Description
get <i>[path]</i>	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.
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: <code>help dir</code>
hexdump	Output the first 50 bytes of the file in a hexdump format.
kill	Terminate the specified process.
memdump <i>[filepath]</i>	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 uploaded to the server and can be downloaded through the CB Response console. Memory dumps can take several minutes, and an (*) icon in the Live Response window indicates that it is still in progress. This is for Windows 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 <i>[remotepath]</i>	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.
pwd	Print the current working directory.
reg	View or modify Windows registry settings. The syntax of this command is: <code>reg [action] [key] [options]</code> See “Registry Access in Live Response” on page 159 or use <code>help reg</code> 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 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.

To end a Live Response session with a computer:

- 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	<p>Format: reg query <i>[key or value] [options]</i></p> <p>Options:</p> <p><i>(none)</i> – If no option switch is specified, query for the specified key</p> <p>-v – Query for the specified value</p> <p>For example:</p> <pre>reg query HKLM\Software\Microsoft\Windows\CurrentVersion\Run</pre>
add	<p>Format: reg add <i>[key] [options]</i></p> <p>Options:</p> <p>-v – Value for the key to be added</p> <p>-d – Data for the key to be added</p> <p>-t – Type of the key to be added; accepted types are:</p> <ul style="list-style-type: none"> • REG_NONE • REG_BINARY • REG_SZ • REG_EXPAND_SZ • REG_MULTI_SZ • REG_DWORD • REG_DWORD_BIG_ENDIAN • REG_QWORD <p>For example:</p> <pre>reg add HKLM\Software\Microsoft\Windows\CurrentVersion\Run -v calc -t REG_SZ -d c:\windows\system32\calc.exe</pre>
delete	<p>Format: reg delete <i>[key or value] [options]</i></p> <p>Options:</p> <p><i>(none)</i> – If no option switch is specified, delete the specified key</p> <p>-v – Delete the specified value</p> <p>For example:</p> <pre>reg delete HKLM\Software\Microsoft\Windows\CurrentVersion\Run -v calc</pre>

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 also 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:

- In the navigation bar, select **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 <i>[id]</i>	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 <i>[id]</i>	Attach to the session whose ID is provided. The <code>session</code> 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 <code>cls</code> command for this purpose.
files -s <i>[id]</i>	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 <i>command</i>	Show the description of the specified command with additional details (such as options) if available. For example: <code>help dir</code>
sensor <i>[options]</i>	List sensors managed by this CB Response server. Options: <code>-i [1.2.3.4]</code> – Return all sensors with specified IP address <code>-n [host_str]</code> – Return all sensors with matching host name <code>-a</code> – 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 <code>-n</code> and <code>-i</code> are specified, only <code>-i</code> is used.

Command	Description
session	<p>Manage Live Response sessions. With no argument, lists all open sessions and their ID numbers, which can be used with the <code>attach</code> command.</p> <p>Options:</p> <ul style="list-style-type: none"> • session new <i>[id]</i> – Create a new session for the sensor whose ID number is provided. You must provide a <i>sensor</i> ID, not a <i>session</i> ID. • session list <i>[-v]</i> – List existing sessions. If the <code>-v</code> option is included, closed sessions are included. This option (without <code>-v</code>) is the default when no additional arguments are used. • close <i>[id]</i> – 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 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` – This is where you would go to begin troubleshooting a Live Response issue. It contains debug information related to the functional operation of the Live Response components and communication between sensor and server.

- `/etc/cb/liveresponse-logger.conf` – This is where you can change the level of information in the `debug.log`.
- `/var/log/cb/audit/live-response.log` – This file is for auditing Live Response activity. It keeps a log of all commands executed on an endpoint, the sensor ID, IP address, and hostname of the endpoint, and the username and account information of the user who executed each one.
- `/var/cb/data/liveresponse` – This directory is where files that are “get” and “put” using Live Response are stored. It also contains the output of all commands executed. For example, if you do a process listing, the list goes into this directory in JSON format. If you download a file (for example, using the archive command), that also appears in this directory (under `/tmp`) and on the host 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. 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 your 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 that you can use 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. DLLs cannot be banned, and so the button does not appear for them.
- 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 and use the dialog that appears 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 could have serious consequences. While CB Response sensors prevent you from banning most critical processes, it is still 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.

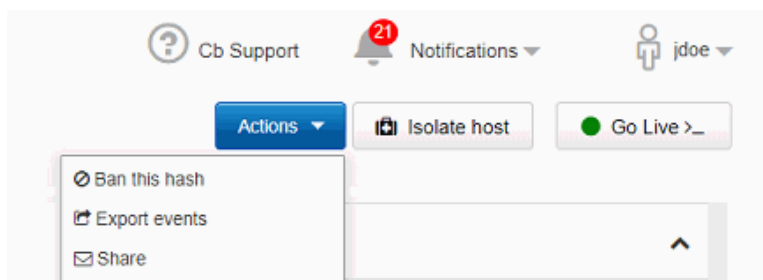
The following procedure describes how to ban an MD5 hash that is listed in any search or other tables in CB Response.

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 with Ban hashes permission and membership on a team with the Analyst role for any Sensor Group.
2. Navigate to the Process Analysis page for the process you want to ban. (See [“Process Search and Analysis”](#) on page 175 for details on accessing these pages.)
3. From the **Process Analysis** page, click **Ban this hash**.

Note

This button only appears if the binary is an EXE. DLLs cannot be banned.



- The **Confirm Banned Hashes** page appears and lists this hash, whether it is known, and the number of computers at this site on which the hash for this process has been seen.

⊘ **Confirm Banned Hashes**

Please carefully review the lists of hashes below. Remove items if you do not want to ban them.

Known Hashes

The following hashes are recognized by Carbon Black. After you press the Ban button below, Carbon Black will prevent them from executing.

<input checked="" type="checkbox"/>	A7BEBF17943A42A7ECFC02B4F40CAE7A		1 computers in 6 processes	
-------------------------------------	----------------------------------	--	----------------------------	--

Notes

Banning the hashes listed above will prevent them from executing.

- Clicking this **Trashcan** icon deletes the hash from the list of those to be banned. For single-hash-ban operations, click **Cancel** at the bottom of the page.
- Add information in the **Notes** box, so that you know why you banned the hash. This can include a file name, threat report identification, or anything else that would be helpful if you or another CB Response administrator or investigator went back to examine the ban.
- 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:

- 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.
- In the navigation bar, select **Banned Hashes**.
The Managed Banned Hashes page appears.

Manage Banned Hashes

25 hashes are banned + Ban More Hashes

View All Banned Previously Banned Sort By MD5 Asc Desc

Hash	Notes	Latest Block	Total Blocks	Hosts w/ Blocks	Banned
f76a4e881845718da5092dd6125f5b12	badstuff 12/12/17	N/A	0		<input type="checkbox"/> ▼
dbc188467bbdeb15c2b79460ce878998	suspicious exe 3	about 7 days	1	DESKTOP-12	<input type="checkbox"/> ▼

- Click the **Ban More Hashes** button in the top-right corner to open the Add Hashes to Ban List dialog box.

⊕ Add Hashes to Ban List ✕

MD5 hashes to ban

Enter one hash per line

Notes

Mention why you are banning these hashes

Close
Ban Hashes

- In the **MD5 hashes to ban** field, enter the MD5 hashes for the processes you want to ban. Each hash must be on its own line.
- 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.
- When you have entered the hashes and notes, click **Ban Hashes** to display the Confirm Banned Hashes page.

⊘ **Confirm Banned Hashes**

Please carefully review the lists of hashes below. Remove items if you do not want to ban them.

Known Hashes

The following hashes are recognized by Carbon Black. After you press the Ban button below, Carbon Black will prevent them from executing.

<input checked="" type="checkbox"/>	A7BEBF17943A42A7ECFC02B4F40CAE7A		1 computers in 6 processes	
-------------------------------------	----------------------------------	--	----------------------------	--

Notes

Cancel
Ban

Banning the hashes listed above will prevent them from executing.

Note

Each hash in the list you provided has a Trashcan icon next to it, so that you can remove it from the list of hashes to be banned if needed. The page also 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** 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 168.
- View block events on the Process Analysis page. See [“To view all block events for a parent process:”](#) on page 170.
- 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 172.

- Monitor banning alerts on the Triage Alerts page. See [“To view banned hash alerts:”](#) on page 173.
- Enable and disable bans on the Manage Banned Hashes page. See [“Disabling a Hash Ban”](#) on page 173.

The Manage Banned Hashes Page

The Manage Banned Hashes page allows you to 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, and any ban may be enabled or disabled by checking or unchecking its box.
- **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.

Hash	Notes	Latest Block	Total Blocks	Hosts w/ Blocks	Banned
f76a4e881845718da5092dd6125f5b12	badstuff 12/12/17	N/A	0		<input type="checkbox"/>
dbc188467bbdeb15c2b79460ce878998	suspicious exe 3	about 7 days	1	DESKTOP-12	<input type="checkbox"/>

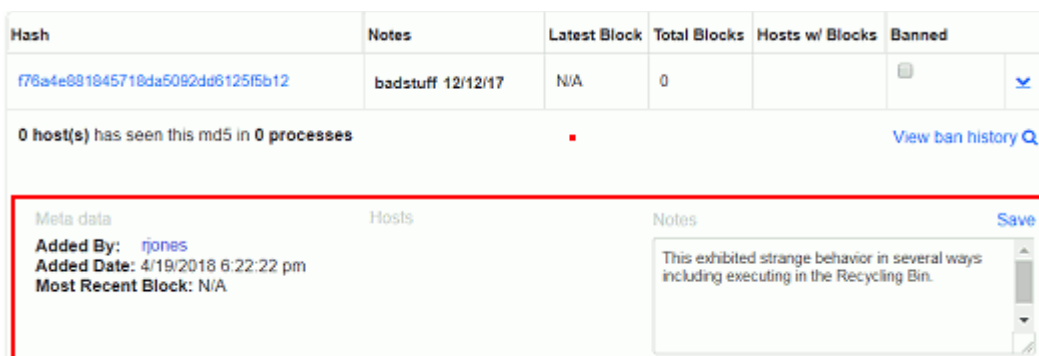
The following table describes fields on this page. Note the data in the table that reports on blocks caused by bans requires that the Banning Events feed on the Threat Intelligence Feeds page is enabled. (For more information on this page, see [“Threat Intelligence Feeds”](#) on page 251.)

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.

Column	Description
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.

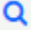
When you expand the row for a ban using the blue down arrow, information about the ban and the process it bans appears in the panel itself. You also can use navigation links there to go to other locations for more information.

Note the data in the table that reports on blocks caused by bans requires that the Banning Events feed on the Threat Intelligence Feeds page is enabled.



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 on the user name navigates to the table of users on the User Management page.
Hosts	The computers on which the process controlled by this ban has been blocked.
Notes	Any user-created notes about the ban or the hash and allows them to be edited and saved.

Column	Description
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.
 (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

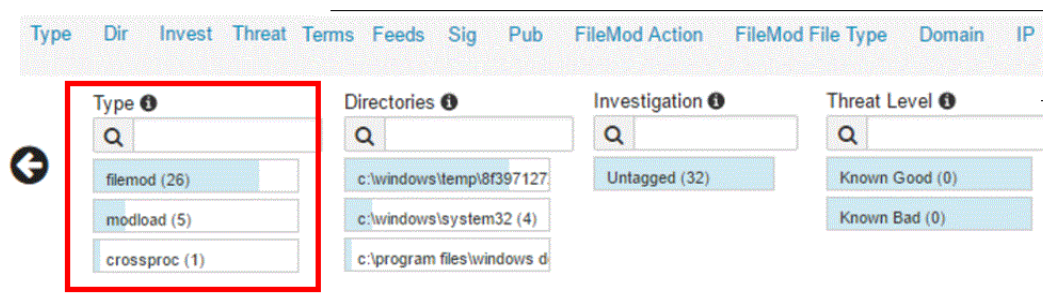
When a process is blocked because of a CB Response hash ban, that is an indication that some user or process attempted an activity you did not want to happen. Even though the activity was blocked, you might want to investigate the attempt.

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 may also be used to trigger alerts and be included in the syslog output from CB Response. See [“Enabling Alerts and Syslog Output for Banning Events”](#) on page 172.

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. (For more information, see [“Process Search and Analysis”](#) on page 175.)

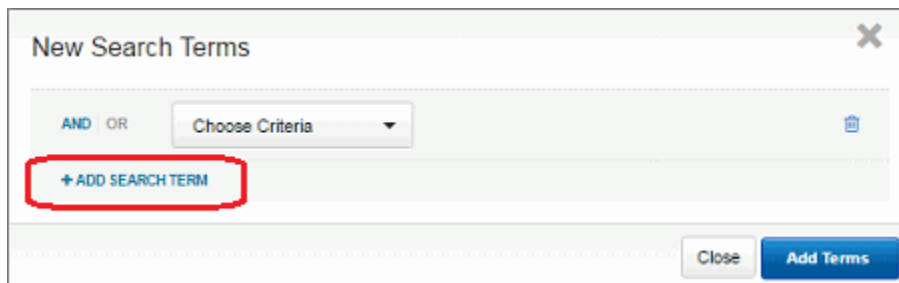


Searching for Blocked Processes

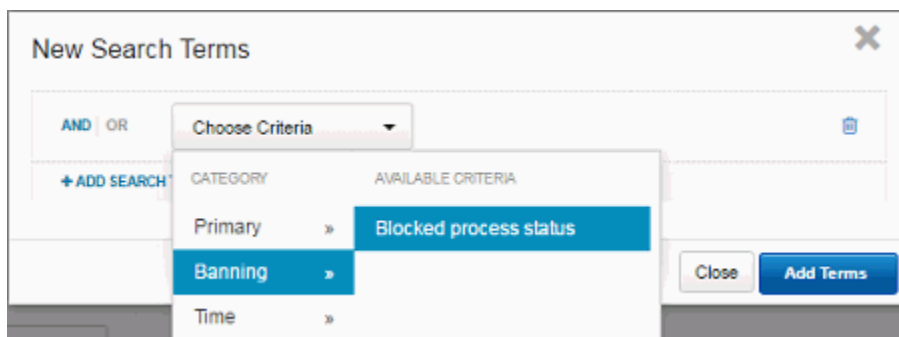
The Process Search page allows you to search for processes that have been affected by a process hash ban. This includes processes successfully blocked as well as those that could not be blocked for various reasons.

To search for processes that have block events:

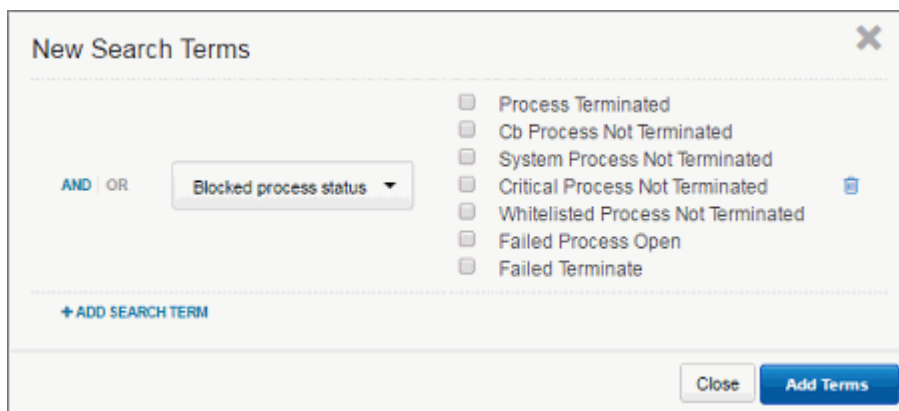
- In the navigation bar, click **Process Search**.
- Click **Add Search Terms**.



3. In the New Search Terms dialog box, click **Choose Criteria** to open the criteria menu and select **Banning > Blocked process status**.



4. A series of **Blocked Process Status** choices is added to the New Search Terms dialog box. You can select one or more block conditions to search for.



5. If you want to search only for successful blocks, select **Process Terminated**. If you want to search for other block events, check the boxes for those events. When you have chosen all of the relevant boxes, click **Add Terms** and run the search. The search results are updated to show block events matching your criteria .

Note

You also can enter one of the following queries for blocked processes manually:

- `block_status:processterminated`
- `processblock_count:[1 TO *]`

Enabling Alerts and Syslog Output for Banning Events

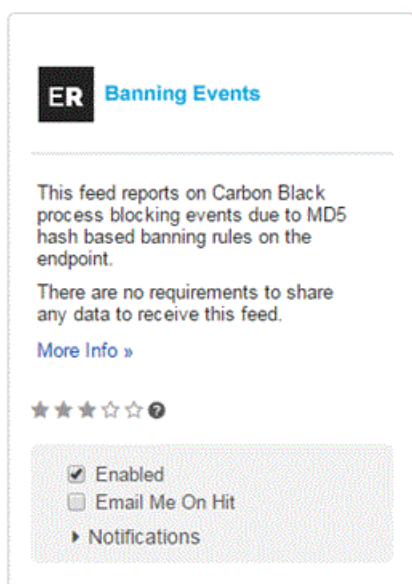
Unless banning is disabled entirely, process block events are sent to the CB Response server and viewable on the Process Analysis page. (For more information, see [“Process Search and Analysis”](#) on page 175.)

To configure alerts and syslog output for process blocks, a special **Banning Events** panel is available on the Threat Intelligence Feeds page. (For more information on this page, see [“Threat Intelligence Feeds”](#) on page 251.) This is not a feed in the normal sense, since the events for blocks are sent to the server whether or not 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 CB Threat Intel.

To enable alerts and syslog recording of blocking events due to hash bans:

1. In the navigation bar, click **Threat Intelligence**:
2. On the Threat Intelligence Feeds page, locate the **Banning Events** feed.



3. In the Banning Events feed panel, click **Notifications**, and then select the notification types you want to create: **Create Alert** and/or **Log to Syslog**.



4. If you also want 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, as well as the time of the most recent attempt and the system it occurred.

To view banned hash alerts:

1. In the navigation bar of the CB Response console, click **Triage Alerts** to display the Triage Alerts page. For more information, see [“Managing Alerts on the Triage Alerts Page”](#) on page 296
2. In the search box for the **Feed** filter, enter `cbbanning` and press Enter, or if it already appears on the list, click on `cbbanning`.

Any **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 that was blocked, the Process Analysis page for the parent process appears, and will include a **blocked** event.

Time ^	Type	Description
2015-06-05 16:14:05.738 GMT	modload	Loaded c:\windows\system32\wdi.dll Signed (bf1fc3f79b863c914687a737c2f3d681)
2015-06-05 16:14:05.737 GMT	modload	Loaded c:\windows\system32\ndfapi.dll Signed (18d4729031314f8c217cdfcc599ef4e)
2015-06-05 16:14:05.00 GMT	blocked	Process c:\program files (x86)\pad\pad.exe (18365b3d9c3ade5ee8ecd36791ee57c8)

Disabling a Hash Ban

After a hash ban is created, it will always appear on the Manage Banned Hashes page. However, you can turn bans on and off if you choose.

To disable a process hash ban:

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. In the navigation bar, click **Banned Hashes**.
The Manage Banned Hashes page appears.

The screenshot shows the 'Manage Banned Hashes' interface. At the top, there is a notification '4 hashes are banned'. Below this is a table with the following columns: Hash, Notes, Latest Block, Total Blocks, Hosts w/ Blocks, and Banned. The 'Banned' column contains checkboxes and dropdown arrows for each row.

Hash	Notes	Latest Block	Total Blocks	Hosts w/ Blocks	Banned
eb37b51c37e857765ef6a36caeee1925	afdsasdf	N/A	0		<input checked="" type="checkbox"/> ↓
cda9f1373805af88f6a4f2064bba24d		N/A	0		<input type="checkbox"/> ↓
b4255003ca47513688035a6a234cd778		N/A	0		<input checked="" type="checkbox"/> ↓
4925fb11e01acde190baf038d8e0e9e8	wordpad on w2k8sp2x86	N/A	0		<input checked="" type="checkbox"/> ↓

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

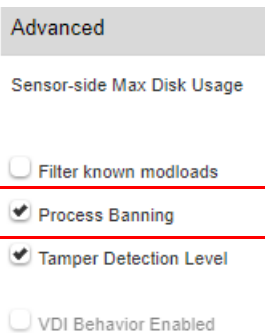
The ability to ban process hashes is enabled by default. However, you can disable or restrict it in certain ways.

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 configured on the server will be ignored by sensors in that group, and no processes will be blocked by CB Response on those sensors. For more information on sensor groups, see [“Sensor Groups”](#) on page 107.

To disable process hash bans in a sensor group:

1. From the navigation bar, select **Sensors**.
2. In the Groups panel, click the name of the sensor group to exempt from hash bans.
3. At the top of the *<name>* Group panel to the right, click **Edit**, and then click **Advanced** to expand the Advanced sensor group settings.
4. Deselect the **Process Banning** option.



5. Click **Save Group**.

Chapter 11

Process Search and Analysis

This chapter describes how to perform detailed process searches and in-depth analysis of the processes in search results.

Sections

Topic	Page
Overview of Process Search	176
Results Table	187
Process Analysis Page	190
Analysis Preview Page	209

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 setting `SensorLookupInactiveFilterDays` determines whether sensors that have not checked in for the 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 will not be included in search results.

To access the Process Search page:

- Select **Process Search** in the navigation bar.

Process Search Cb Support Notifications User

Filters Settings

Username (0) Info
[Search Input]

Process Name (0) Info
[Search Input]

Group (0) Info
[Search Input]

Hostname (0) Info
[Search Input]

Parent Process (0) Info
[Search Input]

Process Path (0) Info
[Search Input]

Process MD5 (0) Info
[Search Input]

Query

[Search Input] Contains text... Last 3 days


[+ ADD SEARCH TERMS](#) [RESET SEARCH](#) [GROUP BY PROCESS](#)

Time Filters

In the **Process Search** page, you can specify that the results show only processes appearing in events that occurred within a specified time period. For example, if you select the **Last 3 days** option, the search results will show processes that appear in events within the last three days.

To use the time filter:

1. Select a time filter from the **Time of Last Update** drop-down list.
2. Click **Search**.



Search Filters

Search filters provide another way 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.

Process	Endpoint	Updated	Start Time	PID	Username	Regmods	Filemods
polkitd /usr/lib/polkit-1/polkitd		May 14, 2019 3:11 PM GMT	Apr 25, 2019 5:18 PM GMT	681	root		
(unknown) (unknown)		May 14, 2019 3:06 PM GMT	Apr 25, 2019 5:18 PM GMT	681	polkitd		
polkitd /usr/lib/polkit-1/polkitd		May 14, 2019 3:02 PM GMT	Apr 25, 2019 5:18 PM GMT	681	root		
polkitd /usr/lib/polkit-1/polkitd		May 14, 2019 2:56 PM GMT	Apr 25, 2019 5:18 PM GMT	681	root		

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**.
The **Choose Filters to Display** window appears.
2. Use the check boxes to enable/disable those filters to display.
Disabling a filter removes it from the view, and if it is part of the search query, those pieces of the query are removed. Enabling a filter places it back in view.
3. Click **Save**.

The screenshot shows the 'Process Search' interface with a 'Choose Filters to Display' dialog box open. The dialog box contains a list of filters with checkboxes and descriptions:

- Username**
This filter shows most common user context seen executing a given process. Use this filter with the Process Name filter to find processes with unexpected usernames.
- Process Name**
This filter indicates which processes reported the largest number of events. The most common processes appear at the top of the list. Less common processes appear at the bottom.
- Group**
Use this filter to identify which sensor groups reported the largest number of process events. This filter is only useful if you organize your sensors into multiple groups.
- Hostname**
This filter shows which endpoints reported the largest number of process events. Use it to identify the endpoints that are running the highest number of processes, or scroll down to find the endpoints that are running the fewest.
- Parent Process**
This filter names the processes that most frequently spawn other processes. Spawned processes include those created by childproc, fork, and crossproc.
- Process Path**
This filter shows the most commonly occurring executable paths for spawned processes. Use this in conjunction with the parent process filter to find unexpected behaviors on your endpoints.
- Process MD5**
This filter shows the MD5 hashes most frequently reported by your endpoints. Use this with the Process Name filter to find processes with unexpected hashes.

A blue 'Save' button is located at the bottom of the dialog box. The background interface shows various filter categories like Username (14), Process Name (50+), Group (2), Hostname (5), and Parent Process (33) with search bars and lists of results.

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 yellow. Click to deselect a filter row.

Process Search

Filters ⚙️

Time of Last Update

Last 3 days ▾

Process Name (29) ℹ️

Q

mDNSResponder (19.1%)	▮
Google Chrome (18.9%)	▮
mds (12.4%)	▮
launchd (12.4%)	▮

Group (1) ℹ️

Q

default group (100.0%)	▮
------------------------	---

Hostname (1) ℹ️

Q

hays (100.0%)	▮
---------------	---

Parent Process (4) ℹ️

Q

launchd (91.0%)	▮
Google Chrome (4.1%)	▮
GoogleSoftwareUpdateAgent (4.1%)	▮
systemstats (0.8%)	▮

Filter Row Percentages

The top (or only) row within a filter is the one that has occurred more than any other process within that filter. Filters can contain one or more filter rows. Filter row percentages indicate the percentage of processes that have occurred within that particular filter. This is always equivalent to 100% when you add up all filter rows within a filter.

Information Icon

You can also hover over the information (“i”) icons for specific information about each filter.

Filter Search Fields

Each filter contains a **Search** field in which you can enter search parameters to further refine search results based on that search filter.

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 the 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, ranked according to the process start time, with the most frequently executed processes at the top. You can also sort the results according to the count of events or last update time.

The **Search** field and criteria fields can be used independently from one another, or they can be used in combination. When used in combination, the system combines them using an AND operator.

Clicking the blue **Search** button executes a search with any parameters you have selected. By default, searches are constrained to the past three days. You can select a different time range for your search by changing the **Process events** drop-down list located under the **Search** button. When viewing events on the Process Search page, the time displayed next to the process is the time when the event was recorded by the sensor, not the time the server received it.

Saved 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 drop-down list.

The selected saved search is loaded and executed:

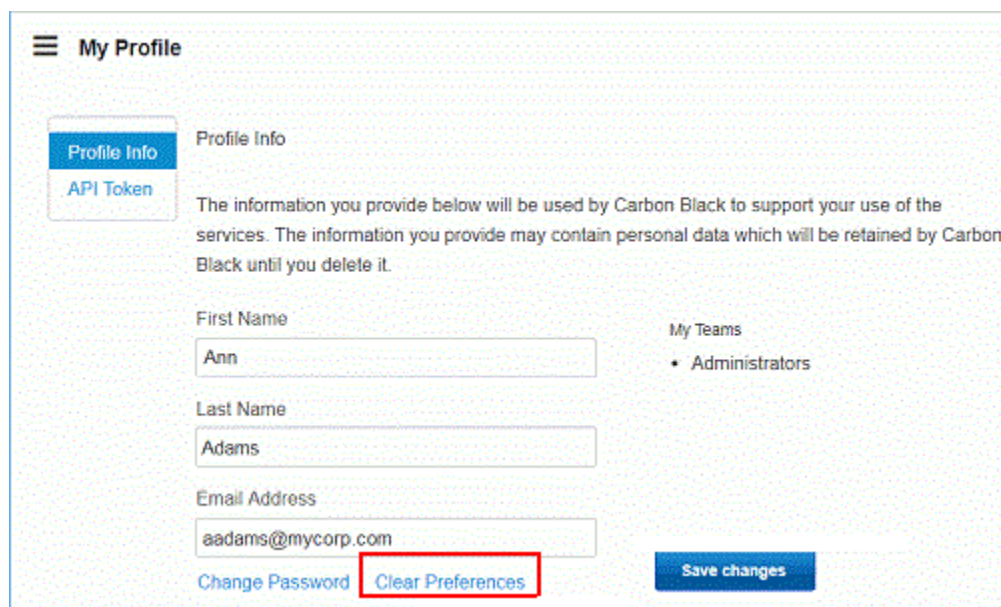


Clear Preferences

You can clear saved searches by accessing your **My Profile** page and clicking **Clear Preferences**.

To clear saved searches:

1. In the top-right corner of the CB Response console, select **username > My Profile**.
The **My Profile** window opens.
2. Click **Clear Preferences**:



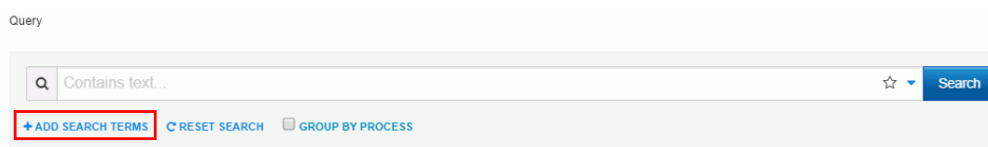
The screenshot shows the 'My Profile' page. On the left, there is a navigation menu with 'Profile Info' and 'API Token'. The main content area is titled 'Profile Info' and contains a warning: 'The information you provide below will be used by Carbon Black to support your use of the services. The information you provide may contain personal data which will be retained by Carbon Black until you delete it.' Below this are three input fields: 'First Name' (Ann), 'Last Name' (Adams), and 'Email Address' (aadams@mycorp.com). To the right, there is a 'My Teams' section with a dropdown arrow and 'Administrators'. At the bottom, there are three buttons: 'Change Password', 'Clear Preferences' (highlighted with a red box), and 'Save changes'.

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 as many search terms as needed (in the form of AND/OR operators) by clicking **Add Search Terms** on the **Process Search** page beneath the **Search** field:

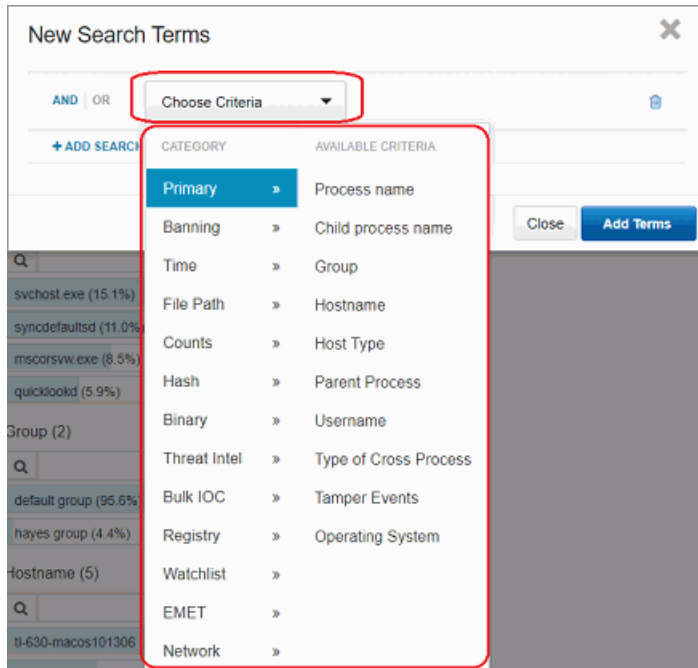


The screenshot shows the 'Query' section of the Process Search page. It features a search bar with the placeholder text 'Contains text...'. To the right of the search bar is a star icon and a 'Search' button. Below the search bar, there are three buttons: '+ ADD SEARCH TERMS' (highlighted with a red box), 'RESET SEARCH', and 'GROUP BY PROCESS'.

The **New Search Terms** dialog box appears:

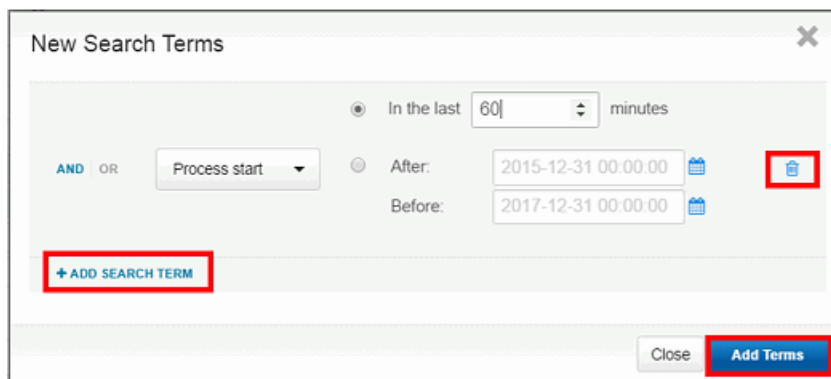


- In the New Search Terms window, select a type of search term from the **Choose Criteria** drop-down list.



- Click **Add Search Term** to add more terms.

In the example below, a **Process start** search term will display processes that have started **in the last 60 minutes**.




- To remove a search term, click the **Delete** (trashcan) icon to the right of the search term.

5. When you have added all search terms, click **Add Terms** in the bottom-right corner of the New Search Terms window to add the search terms to your search:


Reset Search

To reset and remove all search terms, click **Reset Search** on the Process Search page:



Group By Process

To de-duplicate multiple search results associated with the same process (combine results with the same process into groups), select **Group By Process** on the Process Search page:



Search Result Warnings

If no search results match your search criteria, the Process Search page displays the following warning:

Results

Process	Endpoint	Start Time
No Results		
If you applied filters, try deselecting them to widen your search.		
Create a watchlist to get alerts on processes that match this search in the future.		

If search results are too large, you can select one of the options below to pare down the search results:

Results

Process	Endpoint	Start Time
---------	----------	------------

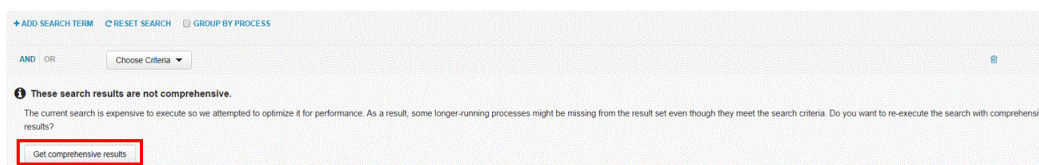
More than 10,439 Results

You can use filters to narrow your search, or [view all 10,439 results](#)

Get Comprehensive Results Button

A **Get Comprehensive Results** button appears on the Process Search page if a search query spans both current and older data collected prior to version 6.1, and the query has complex search terms requiring 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 collected prior to 6.1.
- If you request comprehensive results, the server returns full search results for current data, but excludes data 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. From the navigation bar, click **Process Search**.
The Process Search page appears.
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. For information on performing advanced search queries, see "[Advanced Search Queries](#)" on page 224.
 - b. Click **Choose Criteria** to display a drill-down 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 finish entering your search string or selecting search criteria options, click **Search**.

The search results appears in the **Results** table. For more information, see [“Results Table”](#) on page 187

Note

For information on performing advanced queries in CB Response, see [“Advanced Search Queries”](#) on page 224.

Managing High-Impact Queries

Certain process searches could 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 244 for more on this type of search

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 will not be 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 or another user determine that one of these settings 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.

The second alternative is to reconfigure the search to avoid the blocked condition. See [Chapter 13, “Advanced Search Queries,”](#) for more information about creating more 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, 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 uneditable. See [“Process Search Settings in cb.conf”](#) on page 186.

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. In the main console menu, choose **<username> > Settings**.
3. In the left menu on the Settings page, choose **Advanced Settings**.
4. Check (or uncheck) the box for the search type you want to block (or unblock).
5. Click the **Save changes** button in the lower right corner of the page.

**Process Search Settings in cb.conf**

There are two settings in the `cb.conf` file that 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 interface)
- `ForceBlockCoreJoinsInSearchTo` (interacts with “Block Searches that included Binary Metadata with Large Binary Stores” in the console interface)

By default, neither of these are present in the `cb.conf` file, which allows console users with Global Admin privileges to block or allow the related search through the Process Search Settings on the Advanced Settings page.

If a setting is `True`, process searches in the relevant category are **blocked**, and no user, including a Global Admin, can change this 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 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 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 Showing 10 of 111 Sort by None Edit Columns Create Watchlist Export CSV

Process	Endpoint	Updated	Start Time	PID	Username	Regreads	Filemods	Modloads	Netcons	Children	Tags	Hits
cmd	server12.mycorp.local	Nov 12, 2016 10:01 PM	Nov 12, 2016 10:01 PM	9998	root			10		4		>
cmd	server12.mycorp.local	Nov 12, 2016 11:01 PM	Nov 12, 2016 11:01 PM	10044	root			10		4		>
cmd	server12.mycorp.local	Nov 13, 2016 12:01 AM	Nov 13, 2016 12:01 AM	10056	root			10		4		>
cmd	server12.mycorp.local	Nov 13, 2016 1:01 AM	Nov 13, 2016 1:01 AM	10069	root			10		4		>
cmd	server12.mycorp.local	Nov 13, 2016 2:01 AM	Nov 13, 2016 2:01 AM	10081	root			10		4		>
cmd	server12.mycorp.local	Nov 13, 2016 3:01 AM	Nov 13, 2016 3:01 AM	10094	root			10		4		>
cmd	server12.mycorp.local	Nov 13, 2016 4:01 AM	Nov 13, 2016 4:01 AM	10107	root			10		4		>
cmd	server12.mycorp.local	Nov 13, 2016 5:01 AM	Nov 13, 2016 5:01 AM	10119	root			10		4		>


Results Table Features

A few features appear above the **Results** table:



- **Showing** – Use this option to adjust the maximum number of search results that display on a given page. The default is 10 results per page.
- **Sort by** – Use this menu to sort search criteria by one of 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.
- **Export CSV** – Use this option to export the first 1000 process search results to a .csv file in a comma-separated value format for reporting, retention, or compliance. Each row will contain 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.

Results Table Row Details

On each row within the **Results** table, the following information appears:

Title	Description
Icon	The icon of the process or program that was executed.
Process	The name of the executable file that was run, for example, <code>notepad.exe</code> . Underneath the process name, the path on the file system from which the process was executed appears.
Endpoint	The elapsed time since the most recent execution of the process as well as the endpoint associated with the result.
Updated	The timestamp for when the process was last updated.
Start Time	The timestamp for when the process started.
PID	The Process ID.
Username	The username associated with this process.

Title	Description
Regmods	<p>The number of Windows registry modifications that were made by the execution of this process.</p> <p>Regmods are color-coded as follows:</p> 
Filemods	<p>Contains a color-coded dot if the execution of the process resulted in file modifications. Filemods are color-coded as follows:</p> 
Modloads	<p>Contains a color-coded dot if the execution of the process resulted in loaded modules. Modloads are color-coded as follows:</p> 
Netconns	<p>Contains a color-coded dot if the execution of the process resulted in attempted or established network connections. Netconns are color-coded as follows:</p> 
Children	<p>Contains a color-coded dot if the execution of the process resulted in generated child processes. Children are color-coded as follows:</p> 
Tags	<p>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 as follows:</p> 
Hits	<p>Contains a color-coded dot if the execution of the process resulted in watchlist or feed hits. Hits color-coded as follows:</p> 
>	<p>The Process Analysis page with details about the process executable file. For more information about process analysis, see "Process Analysis Page" on page 190.</p>

Process Analysis Page

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 176.
2. In the **Results** table, locate the process that you want to further analyze.
3. Click the arrow (>) to the right of that process:

Results Showing 10 of 1,491 Sort by None Create Watchlist Export CSV

Process	Endpoint	Regmods	Filemods	Modloads	Netconns	Children	Tags	Hits
svchost.exe c:\windows\system32\svchost.exe	21 days ago laptop-7	●	●	●	●	●		>
svchost.exe c:\windows\system32\svchost.exe	21 days ago laptop-7	●	●	●	●	●		>
rundll32.exe c:\windows\system32\rundll32.exe	21 days ago laptop-7	●	●	●				>
svchost.exe c:\windows\system32\svchost.exe	21 days ago laptop-7	●	●	●	●	●		>
svchost.exe c:\windows\system32\svchost.exe	21 days ago laptop-7	●	●	●	●	●		>

4. The **Process Analysis** page displays activity that CB Response collects for a specific process:

Process Analysis

Google Chrome server-2 rjones Running an hour ago a month

Process Host User State Last Activity Duration

Actions Isolate host

/Applications/Google Chrome.app/Contents/MacOS/Google Chrome

Command Line - Copy

Process: Google Chrome

PID: 320
OS Type: osx
Path: /Applications/Google Chrome.app/Contents/...
Username: rjones
MDS: 5d3c79eb4fd7edfeb1aa6fdb7d92c9
Start Time: 2016-12-09T16:12:15.613Z
Interface IP: 192.168.214.128
Server Comms IP: 10.36.4.22

Google Chrome: Signed by

Alliance Feeds 0 hit(s) in 0 report(s)

Type Dir Invest Threat Terms Feeds Sig Pub FileMod Action FileMod File Type Domain IP Reg Action Reg Hive Child Path Child MDS

Type Directories Investigation Threat Level Search Terms Feeds

Event Timeline

Number of Events

Time	Type	Description
2016-12-09 18:58:02.382 GMT	filemod	Deleted /Users/rjones/Library/Application Support/Google/Chrome/ /com.google.Chrome.EeWRRH
2016-12-09 18:58:02.381 GMT	filemod	Created /Users/rjones/Library/Application Support/Google/Chrome/ /com.google.Chrome.EeWRRH
2016-12-09 18:58:02.381 GMT	filemod	First wrote to /Users/rjones/Library/Application Support/Google/Chrome/ /com.google.Chrome.EeWRRH
2016-12-09 18:53:32.373 GMT	filemod	Deleted /Users/rjones/Library/Application Support/Google/Chrome/Default/ /com.google.Chrome.iPDZ3q

Process Analysis Features

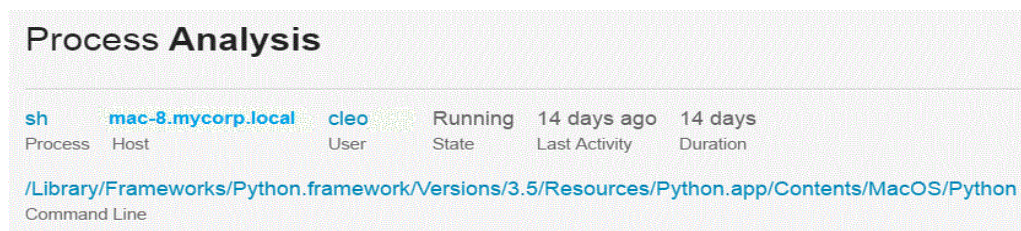
The **Process Analysis** page contains various features to help you more deeply investigate process details, such as:

- [“Process Summary”](#) on page 192
- [“Isolate Host”](#) on page 192
- [“Go Live”](#) on page 193
- [“Actions Menu”](#) on page 193
- [“Interactive Process Tree”](#) on page 195
- [“Process Execution Details”](#) on page 196

- [“Binary Metadata”](#) on page 197
- [“Feeds”](#) on page 198
- [“On Demand Feeds”](#) on page 198
- [“EMET Protections Enabled \(Windows Only\)”](#) on page 199
- [“Process Event Filters”](#) on page 199
- [“Event Timeline”](#) on page 202
- [“Process Event Details”](#) on page 202

Process Summary

The process summary information is located in the top-left corner of the **Process Analysis** page and displays general process execution details. For example:




The screenshot shows a table with the following data:

Process	Host	User	State	Last Activity	Duration
sh	mac-8.mycorp.local	cleo	Running	14 days ago	14 days

Below the table, the command line is displayed: `/Library/Frameworks/Python.framework/Versions/3.5/Resources/Python.app/Contents/MacOS/Python`

Isolate Host

The **Isolate host** button is located at the top-right corner of the Process Analysis page:



Use this option to isolate a computer. For example, you might discover that suspicious files have been executing from a particular computer and you want to prevent them from spreading to other computers in your network.

When a computer (host) 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, even though the computer will not work as expected.


Note

In order 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 with Analyst privileges for the endpoint you want to isolate.

The computer remains isolated until this option is disabled or the computer reboots. See [“Isolating an Endpoint”](#) on page 151 for more information.

Go Live

The **Go Live** button is located in the top-right corner of the Process Analysis page:



This option is useful when you are investigating an IOC. After you have identified a computer with suspicious activity, you can directly access the content on that system. You can open an interactive live session to the end-point host and execute commands in real time to help isolate or eradicate the threat. See ["Using Live Response"](#) on page 154 for more information.

Actions Menu

The **Actions** drop-down list includes the following options:

- **Ban this hash** – Creates a ban of the process displayed on the Process Analysis page. 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 when the ban is created are terminated. See ["Banning Process Hashes"](#) on page 163 for more information.
- **Export events to CSV** – Downloads a `Report.zip` archive to your local computer. The files in the archive 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 199 for more information.

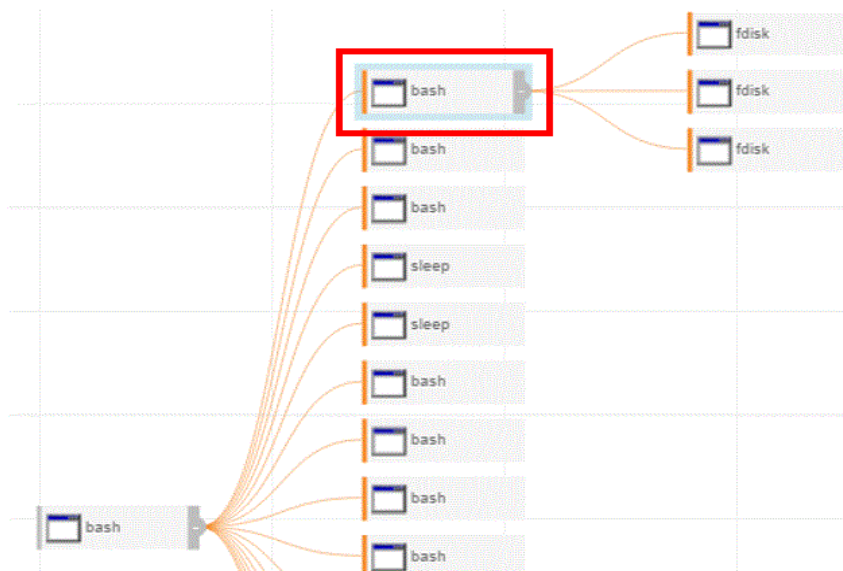
The screenshot displays the CB Response search and analysis interface. At the top, there are various filter tabs: Type, Dir, Invest, Threat, Terms, Feeds, Sig, Pub, FileMod Action, FileMod File Type, Domain, IP, Reg Action, Reg Hive, Child Path, and Child MD5. Below these are search filters for Type, Directories, Investigation, Threat Level, Search Terms, and Feeds. The 'Type' filter is highlighted with a red box and shows a list of process types: modload (21), childproc (2), and crossproc (2). Below the filters is an 'Event Timeline' chart showing the number of events over time, with a red box around the chart area. At the bottom, there is a table of events with columns for Time, Type, and Description. The table is also highlighted with a red box.

Time	Type	Description
2016-08-31 02:00:52.358 GMT	childproc	PID 3676 ended c:\windows\system32\conhost.exe Signed (d5669294f78a7d48c318e22d5085ba7)
2016-08-31 02:00:52.358 GMT	crossproc	A handle to a thread in this process was opened with change rights by c:\windows\system32\lsass.exe (382100e75b0f4088aeeef228c5ceffad)
2016-08-31 02:00:52.358 GMT	modload	Loaded c:\windows\system32\rport4.dll Signed (2522d170c08f370a208d1305dd580909)
2016-08-31 02:00:52.358 GMT	modload	Loaded c:\windows\system32\taskschd.dll Signed (8e2833eef1c1f525587b1acc52054od5)
2016-08-31 02:00:52.358 GMT	modload	Loaded c:\windows\system32\vmvort.dll Signed (d80e5d199b770cb89dbdc52ab58c7519)
2016-08-31 02:00:52.358 GMT	modload	Loaded c:\windows\system32\user32.dll Signed (33094e2182c451bcfd690f734b1c4ef)
2016-08-31 02:00:52.358 GMT	modload	Loaded c:\windows\system32\ole32.dll Signed (5a085dcb8be5f4e58dc7be033a411e3e)
2016-08-31 02:00:52.358 GMT	modload	Loaded c:\windows\system32\oleaut32.dll Signed (6835d94fdaab39e008e8490bd3e88ca3)

- **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), as well as 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, with the relevant child process highlighted, as shown in this example:



You can interact with the process tree by clicking other 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 on 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 only display up to 15 child processes; either 15 unsuppressed, 15 suppressed, or 15 of both types.

For processes with 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.

Process Execution Details

Details about the process execution appear in the panel on the top-right side of the Process Analysis page:

Process: coreduetd	
PID	84
OS Type	osx
Path	/usr/libexec/coreduetd
Username	root
MD5	e1fb6c97138ad89d6d7b37e5973eff2
SHA-256	6167fb0a27da20206f01713f9669509abc4db625276 5d42a80b7134cd5501fbc
Start Time	2019-02-27T22:16:56.371Z
Interface IP	10.260.142.64
Server	10.260.142.64
Comms IP	
coreduetd: Signed by	

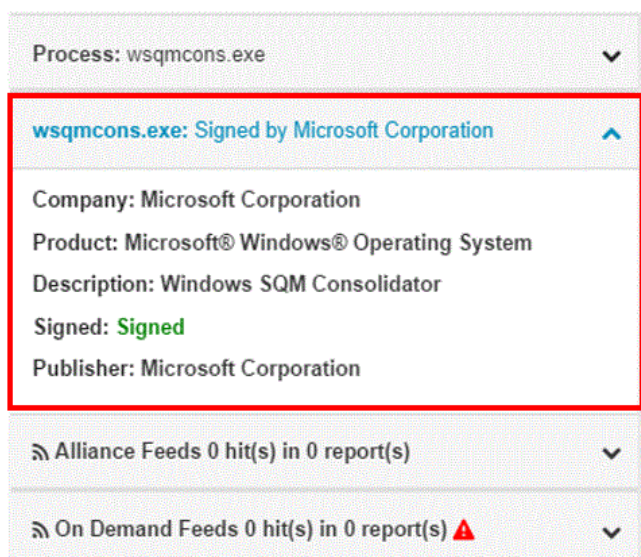
If the process is an executable, the following information is displayed:

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 on the host computer who executed the process.
MD5	The MD5 hash value of the process.
SHA-256	<p>The SHA-256 hash value of the process.</p> <p>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 capable of generating SHA-256 hashes.</p> <p>For files 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 another sensor that supports SHA-256.</p>
Start Time	The date and time of the process execution.
Interface IP	<p>The IP address of the network <i>adapter</i> on the sensor on which the network connection to the CB Response server was made.</p> <p>Note: Pre-5.1 sensors do not report an Interface IP.</p>

Field	Description
Server Comms IP	The IP address from which the server recognizes the sensor reporting data. If the sensor is communicating through a Proxy or NAT device, the address will be for that device, not the sensor itself. Otherwise it should be the same as the Interface IP address.

Binary Metadata

Details about the binary metadata (specifically, digital signature information) appear in the panel on the top-right side of the Process Analysis page:



The following information appears:

Field	Description
<i><process executable file name></i>	The name of the process executable file.
Company	The name of the company that created the process executable file.
Product	The product for which the process executable file was created.
Description	A text description of the product .
Signed	Shows if the process or module that was executed or loaded has been signed by the publisher.
Publisher	The official publisher of the process executable file.

Feeds

The **Feeds** panel, located on the top-right side of the Process Analysis page, shows if the process event details had any hits from CB Threat Intel partner feeds:

The screenshot shows the Feeds panel for the process explorer.exe. It is expanded to show one hit from the Alliance Feeds. The hit is from VirusTotal, dated 5-19-2016, with a score of 1. The IOC hash value is A5675939CF0F99B20B5A3CFCC3C1B46A. Below this, the On Demand Feeds section is collapsed and shows 0 hits.

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

The **On Demand Feeds** panel, located on the top-right side of the Process Analysis page, shows if the process event details had any hits from on-demand feeds:

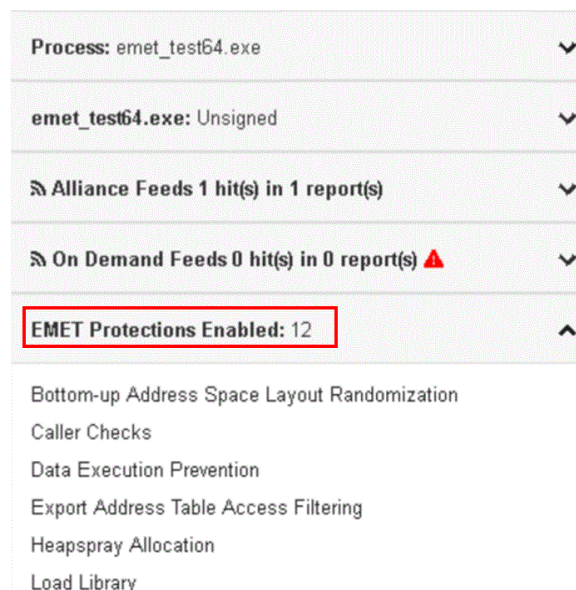
The screenshot shows the On Demand Feeds panel for the process wsqmcons.exe. It is expanded to show a message: "Sharing of data to the Carbon Black Threat Intelligence Cloud is not enabled for the group that contains this process. Go to Sharing Settings to enable sharing." The message is highlighted with a red border.

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 where you can set this up. For more information, see [“On-Demand Feeds from CB Threat Intel”](#) on page 262.


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.



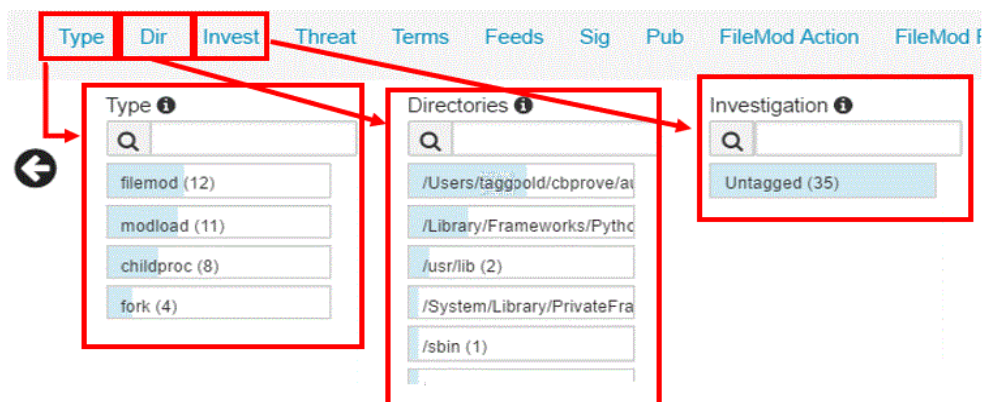
Process Event Filters

Several filters provide high-level details on events that occurred in the process executable file. You can access the filters by clicking the blue filter name in the filter menu.




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 allow you to scroll through the list of available filters that are not visible on your page.



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** 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. For more details on process events, see “Process Event Details” on page 202. <ul style="list-style-type: none"> • 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


Filter	Menu Bar Name	Description
Directories	Dir	The directories used in this process.
Investigation	Invest	The tagged status for events in this process for any investigations.
Threat Level	Threat	Shows report scores for events associated with CB Threat Intel hits in this process. For more information, see “Threat Intelligence Feed Scores” on page 253
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 associated with network connections that were made by this process.
IP Address	IP	The IP addresses 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	Shows paths to child processes that were created by this process.
Childproc md5s	Child MD5	Shows MD5 files of child processes that were created by this process.

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. For example, if you have been filtering or searching in any of the filters, you can reset them to their original state.

Event Timeline

The **Event Timeline** is located beneath the process event filters discussed in “[Process Event Filters](#)” on page 199. This is useful for investigating IOCs if you want to view 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/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 back and forth across the timeline, the Process Event Details below are updated. For more information on Process Event Details, see “[Process Event Details](#)” on page 202.



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 essentially zoom in on event segments in the top graph to view event counts for particular time segments.



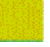
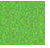



Process Event Details

The **Process Events Details** view for a selected process appears as a table with several rows at the bottom of the Process Analysis page:

<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Time ^	Type	Description	Q	Search
<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	2016-05-09 15:23:41.737 GMT	fork	Process id 5160		
<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	2016-05-09 15:23:41.737 GMT	childproc (suppressed)	PID 5160 started /usr/bin/uname Signed (78f4785c0c51531f1c01d4161c96563f)		
<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	2016-05-09 15:23:41.737 GMT	childproc (suppressed)	PID 5160 ended /usr/bin/uname Signed (78f4785c0c51531f1c01d4161c96563f)		
<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	2016-05-09 15:23:41.736 GMT	modload	Loaded /usr/bin/uname Signed (78f4785c0c51531f1c01d4161c96563f)		
<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	2016-05-09 15:23:41.734 GMT	modload	Loaded /usr/lib/libDiagnosticMessagesClient.dylib Signed (34657d5c5dce60bc438065e4e86add4b)		
<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	2016-05-09 15:23:41.734 GMT	modload	Loaded /usr/lib/libc++.1.dylib Signed (36e61c9c4b6046fb3ea60030b3af2340)		
<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	2016-05-09 15:23:41.734 GMT	modload	Loaded /usr/lib/libc++abi.dylib Signed (63d163cf8989e92eb2dccffac785d771)		
<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	2016-05-09 15:23:41.734 GMT	modload	Loaded /usr/lib/libauto.dylib Signed (41a110ef2f47d396453feef5775a3a0)		
<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	2016-05-09 15:23:41.734 GMT	modload	Loaded /usr/lib/libobjc.A.dylib Signed (27bc356550009fb1a372fbe28b69fceb)		
<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	2016-05-09 15:23:41.734 GMT	modload	Loaded /usr/lib/system/libxpc.dylib Signed (0b3ba8db1261786a06d91d60fae8bcf2)		
<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	2016-05-09 15:23:41.734 GMT	modload	Loaded /usr/lib/system/libunwind.dylib Signed (9b46f4cd12e68047f68c368ca1a931d3)		

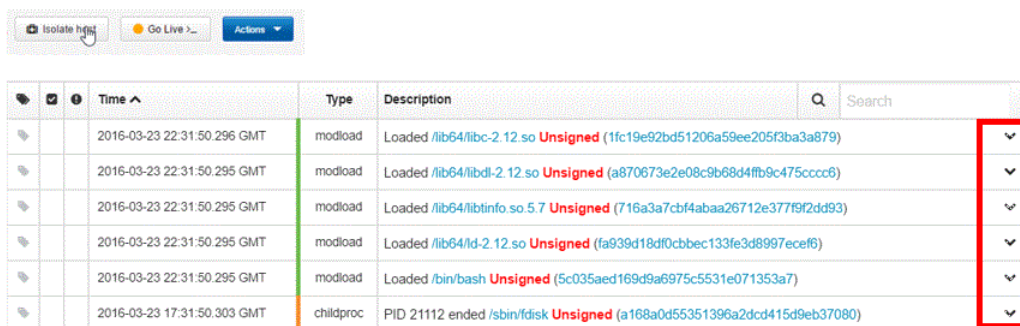
The process events rows show the following details:

Heading	Description
Tag 	Shows if an event is tagged for an investigation. You can click the tag icon to select this event for future investigation. After you select the tag icon, it turns blue to indicate that it is now included in an investigation.
Trusted Events <input checked="" type="checkbox"/>	Shows if the event is trusted. When you click on the row, the trust information appears with a link to the source.
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).


Heading	Description
Type	<p>The process event type. (For more information on event types, see "Process Event Types" on page 206.)</p> <ul style="list-style-type: none"> • 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	<p>The operation that the Type event performed. (For more information on event types, see “Process Event Types” on page 206.)</p> <p>The Description column might contain:</p> <ul style="list-style-type: none"> • filemod – Might display “Deleted” or “Created” and then provide the path to the file that was modified. • modload – Might display the module that was loaded by the process. Modload descriptions could 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 – Might display the Windows registry key that was created or modified. • netconn – Might display the connection made, including the IP address (including hostname unless DNS resolution is excluded for the host), port, and protocol. • childproc – Might display the child process start time, end time, and PID of the selected parent process. • fork – (OS X and Linux only) Might display the instance’s parent process, forked with a different Process ID (PID). • posix_exec – (OS X and Linux only) Might display the instance’s process that is loaded and the new binary image. • crossproc – (Windows only - not supported on Windows XP/2003) Might display the action it performed, for example, opening a handle or starting or ending processes. • blocked – Displays 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 used in the attempt to run the process
Search	<p>Allows you to 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.</p>

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 computers that have seen this file modification and the number of processes in which the file modification occurred on those computers.
modload	This includes following information: <ul style="list-style-type: none"> The number of computers that have seen the MD5 hash for the module that was loaded and the number of processes the MD5 appears in on those computers. 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	The number of computers that have seen a modification of a registry key, and the number of processes in which the registry modification occurred on those computers. Windows sensors only.
netconn	The number of network connections that the execution of this process either attempted or established.

Event Type	Details																								
<p>childproc</p>	<p>The following information:</p> <ul style="list-style-type: none"> • The number of computers that have seen the MD5 in the description and the number of processes in which this 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 (for example, “about one month ago”), username of the user 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. For more information, see “Advanced Settings” on page 113. This image shows suppressed vs. unsuppressed child processes. Note that suppressed child processes are labeled Suppressed in the process tree (see “Interactive Process Tree” on page 195). <table border="1" data-bbox="594 1024 1328 1234"> <tbody> <tr> <td>2016-04-21 17:04:14.861 GMT</td> <td>childproc</td> <td>PID 4322 ended /Users/prove/tools/forkexec/forwriteexec: Unsigned (99dd58b0a0a5c47271e876da79f)</td> </tr> <tr> <td>2016-04-21 17:04:14.861 GMT</td> <td>childproc</td> <td>PID 4322 started /Users/prove/tools/forkexec/forwriteexec: Unsigned (99dd58b0a0a5c47271e876da79f)</td> </tr> <tr> <td>2016-04-21 17:04:09.705 GMT</td> <td>childproc (suppressed)</td> <td>PID 4320 ended /Users/prove/tools/forkexec/fork: Unsigned (c3a277905556a29f26f8197d203b966)</td> </tr> <tr> <td>2016-04-21 17:04:09.705 GMT</td> <td>childproc (suppressed)</td> <td>PID 4320 started /Users/prove/tools/forkexec/fork: Unsigned (c3a277905556a29f26f8197d203b966)</td> </tr> <tr> <td>2016-04-21 17:04:03.408 GMT</td> <td>childproc</td> <td>PID 4318 started /Users/prove/tools/forkexec/forkexec: Unsigned (615486fc748222db2c13bd7770132a)</td> </tr> <tr> <td>2016-04-21 17:04:03.407 GMT</td> <td>childproc</td> <td>PID 4318 ended /Users/prove/tools/forkexec/forkexec: Unsigned (615486fc748222db2c13bd7770132a)</td> </tr> <tr> <td>2016-04-21 17:03:58.394 GMT</td> <td>childproc (suppressed)</td> <td>PID 4317 ended /bin/ls Signed (a5b67fd3f99a224f31db0eeaf25a6668a)</td> </tr> <tr> <td>2016-04-21 17:03:58.394 GMT</td> <td>childproc (suppressed)</td> <td>PID 4317 started /bin/ls Signed (a5b67fd3f99a224f31db0eeaf25a6668a)</td> </tr> </tbody> </table> <p>Note</p> <p>The process tree shows a maximum of 15 child processes; either 15 unsuppressed, 15 suppressed, or 15 of both types.</p> <p>For processes with 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.</p>	2016-04-21 17:04:14.861 GMT	childproc	PID 4322 ended /Users/prove/tools/forkexec/forwriteexec: Unsigned (99dd58b0a0a5c47271e876da79f)	2016-04-21 17:04:14.861 GMT	childproc	PID 4322 started /Users/prove/tools/forkexec/forwriteexec: Unsigned (99dd58b0a0a5c47271e876da79f)	2016-04-21 17:04:09.705 GMT	childproc (suppressed)	PID 4320 ended /Users/prove/tools/forkexec/fork: Unsigned (c3a277905556a29f26f8197d203b966)	2016-04-21 17:04:09.705 GMT	childproc (suppressed)	PID 4320 started /Users/prove/tools/forkexec/fork: Unsigned (c3a277905556a29f26f8197d203b966)	2016-04-21 17:04:03.408 GMT	childproc	PID 4318 started /Users/prove/tools/forkexec/forkexec: Unsigned (615486fc748222db2c13bd7770132a)	2016-04-21 17:04:03.407 GMT	childproc	PID 4318 ended /Users/prove/tools/forkexec/forkexec: Unsigned (615486fc748222db2c13bd7770132a)	2016-04-21 17:03:58.394 GMT	childproc (suppressed)	PID 4317 ended /bin/ls Signed (a5b67fd3f99a224f31db0eeaf25a6668a)	2016-04-21 17:03:58.394 GMT	childproc (suppressed)	PID 4317 started /bin/ls Signed (a5b67fd3f99a224f31db0eeaf25a6668a)
2016-04-21 17:04:14.861 GMT	childproc	PID 4322 ended /Users/prove/tools/forkexec/forwriteexec: Unsigned (99dd58b0a0a5c47271e876da79f)																							
2016-04-21 17:04:14.861 GMT	childproc	PID 4322 started /Users/prove/tools/forkexec/forwriteexec: Unsigned (99dd58b0a0a5c47271e876da79f)																							
2016-04-21 17:04:09.705 GMT	childproc (suppressed)	PID 4320 ended /Users/prove/tools/forkexec/fork: Unsigned (c3a277905556a29f26f8197d203b966)																							
2016-04-21 17:04:09.705 GMT	childproc (suppressed)	PID 4320 started /Users/prove/tools/forkexec/fork: Unsigned (c3a277905556a29f26f8197d203b966)																							
2016-04-21 17:04:03.408 GMT	childproc	PID 4318 started /Users/prove/tools/forkexec/forkexec: Unsigned (615486fc748222db2c13bd7770132a)																							
2016-04-21 17:04:03.407 GMT	childproc	PID 4318 ended /Users/prove/tools/forkexec/forkexec: Unsigned (615486fc748222db2c13bd7770132a)																							
2016-04-21 17:03:58.394 GMT	childproc (suppressed)	PID 4317 ended /bin/ls Signed (a5b67fd3f99a224f31db0eeaf25a6668a)																							
2016-04-21 17:03:58.394 GMT	childproc (suppressed)	PID 4317 started /bin/ls Signed (a5b67fd3f99a224f31db0eeaf25a6668a)																							
<p>fork</p>	<p>(OSX and Linux only) Indicates this is a fork process and shows the instance’s parent process, forked with a different Process ID (PID).</p> <p>When a process performs a fork() system call, all activity for that process will continue to be associated with the parent. A new fork event type will be 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 will be recorded</p>																								

Event Type	Details
posix_exec	<p>(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.</p> <p>If at any point 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.</p>
crossproc	<p>Windows only (not supported on Windows XP/2003): Shows occurrences of processes that cross the security boundary of other processes:</p> <ul style="list-style-type: none"> • Description of the OpenProcess API call for the cross process. CB Response records all OpenProcess API calls that request <code>PROCESS_CREATE_THREAD</code>, <code>PROCESS_DUP_HANDLE</code>, <code>PROCESS_SUSPEND_RESUME</code>, <code>PROCESS_VM_OPERATION</code>, or <code>PROCESS_VM_WRITE</code> 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, user name 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	<p>The path and hash of a process that has been blocked by a CB Response process hash ban. When expanded, provides metadata for the process and its binary:</p> <ul style="list-style-type: none"> • 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	<p>(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 computers 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 may provide additional details.</p>

Analysis Preview Page

In the Process Search page (discussed in “Overview of Process Search” on page 176), scroll down to the **Results** table (discussed in “Results Table” on page 187). Click anywhere in a query result row (except for on a hyperlinked item or the “>” icon):

Results Showing 10 of 10,047 Sort by None Create Watchlist Export CSV

Process	Endpoint	Regmods	Filemods	Modloads	Netconns	Children	Tags	Hits
searchindexer.exe c:\windows\system32\searchindexer.exe	a month ago amcity	●	●	●		●		>
searchindexer.exe c:\windows\system32\searchindexer.exe	a month ago amcity	●	●	●		●		>
searchindexer.exe c:\windows\system32\searchindexer.exe	a month ago amcity	●	●	●		●		>
searchindexer.exe c:\windows\system32\searchindexer.exe	a month ago amcity	●	●	●		●		>
searchindexer.exe c:\windows\system32\searchindexer.exe	a month ago amcity	●	●	●		●		>
searchindexer.exe c:\windows\system32\searchindexer.exe	a month ago amcity	●	●	●		●		>
sdiagnhost.exe c:\windows\system32\sdiagnhost.exe	2 days ago amcity	●	●	●		●		>
services.exe c:\windows\system32\services.exe	a month ago amcity	●	●	●		●		>

The Analysis Preview page appears:

Preview ✕

system_installd [Analyze >](#)

Running for 14 days, last activity about 13 days ago [View Binary >](#)

Signed status: Signed Company: Apple Inc. Product: (unknown) Description: (unknown) Publisher:	Hostname: mac-6.mycorp.local Start time: 2016-05-09T15:18:40.79Z Path: /System/Library/PrivateFrameworks/PackageKit.framework/Versions/A/Resources/system_installd Command line: /System/Library/PrivateFrameworks/PackageKit.framework/Resources/system_installd Username: root
---	---

regmods: 0 filemods: 10000 modloads: 0 netconns: 0

Time	Type	Description
Tue May 10 2016 09:38:47 GMT-0500 (Central Daylight Time)	filemod	Deleted /private/var/folders/zz/zyxvpxvq6cstxvn_n0000000000000/C/PKInstallSandboxManager-SystemSoftware/9A341D51-A633-48E2-B977-5739CD305950.activeSandbox/Root/usr/share/man/man3/BC.T_HJAJS
Tue May 10 2016 09:38:47 GMT-0500 (Central Daylight Time)	filemod	Last wrote to /private/var/folders/zz/zyxvpxvq6cstxvn_n0000000000000/C/PKInstallSandboxManager-SystemSoftware/9A341D51-A633-48E2-B977-5739CD305950.activeSandbox/Root/usr/share/man/man3/instr.3x () (Unknown)
Tue May 10 2016 09:38:47 GMT-0500 (Central Daylight Time)	filemod	Created /private/var/folders/zz/zyxvpxvq6cstxvn_n0000000000000/C/PKInstallSandboxManager-SystemSoftware/9A341D51-A633-48E2-B977-5739CD305950.activeSandbox/Root/usr/share/man/man3/BC.T_fPKxZ0
Tue May 10 2016 09:38:47 GMT-0500 (Central Daylight Time)	filemod	Deleted /private/var/folders/zz/zyxvpxvq6cstxvn_n0000000000000/C/PKInstallSandboxManager-SystemSoftware/9A341D51-A633-48E2-B977-5739CD305950.activeSandbox/Root/usr/share/man/man3/BC.T_fPKxZ0
Tue May 10 2016 09:38:47 GMT-0500 (Central Daylight Time)	filemod	First wrote to /private/var/folders/zz/zyxvpxvq6cstxvn_n0000000000000/C/PKInstallSandboxManager-SystemSoftware/9A341D51-A633-48E2-B977-5739CD305950.activeSandbox/Root/usr/share/man/man3/BC.T_fPKxZ0

more events can be found on the [analyze page](#)

Close

The Analysis Preview page provides a quick overview of the following execution details for the process 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.
Hostname	The name of the host 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.
Username	The user on the given host who executed the process. The format is <domain>\<username>.
Publisher	The official publisher of the process executable file.
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.
Time	The full timestamp for a data source (data sources are regmod, filemod, modload, or netconn). The time is displayed in Greenwich Mean Time (GMT).
Type	The type of data source.
Description	Shows information about the event in context with the event type. For example, for filemods, the path of the file that was modified would display in this field.
Analyze	Click to open the Process Analysis page, which provides a more granular analysis of the process executable file. (This is the same page that opens when you click the Process Analysis icon (>) in the Process Search page.)
View Binary	Click to view the detailed binary analysis page for the process executable file. For more information, see "Binary Search and Analysis" on page 212.
More events...	Click the more events can be found on the analyze page link. This opens the Process Analysis page, which contains more process event details. For more information, "Process Analysis Page" on page 190.

Title	Description
Hostname	The name of the host 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.
Username	The user on the given host who executed the process. The format is <domain>\<username>.

Chapter 12

Binary Search and Analysis

This chapter explains how to search for and analyze binary metadata.

Sections

Topic	Page
Overview of Binary Search	213
Entering Search Criteria	213
High-level Result Summaries	215
Related Metadata	217
Binary Search Results Table	217
Binary Preview	218
Binary Analysis	219

Overview of Binary Search

CB Response sensors begin tracking binaries at the moment that they are executed by a process. You can perform a binary search to explore the metadata of a binary.

To search for binaries, in the navigation bar, select **Binary Search**.

Search Binaries

Contains text... Search

+ Add Criteria

Digital Signature Publisher Company Name Product Name File Version File paths Groups Hostnames

Digital Signature (8) Publisher (64) Company Name (78) Product Name (200)

Unsigned (60.5%) Signed (39.4%) Bad Signature (0.0%) Invalid Chain (0.0%)

Microsoft Corporation (32.3%) Dell (22.9%) Microsoft (21.8%) Google (21.6%)

Microsoft Corporation (36.4%) Dell (20.7%) Microsoft (19.7%) Google (19.6%)

Microsoft® Windows® Operatin... Microsoft® .NET Framework (1... Microsoft SQL Server (7.8%) Microsoft (R) Windows (R) Oper...

Sign Time Host Count First Seen Cb Reputation Score

Related Metadata

Results Showing 10 of 19,638 Sort by First seen time

Binary	Time First Seen	Signature Status	Size
FEF08C107812B3684B741C3211BA6B60 usbhub.sys	5 minutes ago	Signed Microsoft Corporation	409.81 KB
BE743083CF7083C406A4390E3AEFE59A fpydisk.sys	5 minutes ago	Signed Microsoft Corporation	24.5 KB


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, on the Binary Search page, 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. In the navigation bar, select **Binary Search** to display the Binary Search page.
2. In the **Search** box:
 - Enter a search string (must be formatted with the correct syntax) or
 - Click **Add Criteria** to display predefined search criteria options:



3. If you select a search criteria option after clicking **Add Criteria**, a window appears where you must specify details for that search criteria option. Repeat this process to use more than one search criteria option. For example, if you select the **OS Type** search criteria option, the following window appears where 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 the AND operator.


4. When you finish entering a search string or selecting 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 224.

Additional Search Page Features

In the top-right corner of the page, an **Actions** menu provides several options:



- **Share** – Use this option to share query strings with other people. 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** – 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.
- **Export CSV** – Use this option to export the first 1000 process search results to a .csv file in a comma-separated value format for reporting, retention, or compliance. Each row will contain 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.


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 binaries that seem suspicious.

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:

- **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 turns the facet gold and 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 matching the current search listed by CB Reputation Score.

Related Metadata

Below the facets and to the left of the table of binary search results, the **Related Metadata** panel appears. In this example, the metadata reflects search results using `calc.exe` for search criteria.

Related Metadata
c:\windows\system32\calc.exe
c:\windows\system32\calc.exe
C:\Windows\system32\calc.exe
C:\Windows\system32\calc.exe
\\server5.myc0.localshare\calc.exe

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 appear. Each row provides details about binary metadata that matches the search criteria.

Results		Showing 10 of 10,035	Sort by First seen time
Binary	Time First Seen	Signature Status	Size
 5B00E3364FA5C185510170EA75EA8884 onedrivestandaloneupdater.exe	a day ago	Signed Microsoft Corporation	2.4 MB
 9D3CCF104C7CF32E1DC581A84EF761E1 usoclient.exe	2 days ago	Signed Microsoft Corporation	38.5 KB
 35E8431ACDDB1F236393CF681738F5FD msvcp110_win.dll	2 days ago	Signed Microsoft Corporation	407.59 KB
 534F396332A76B9F4AAA7C8CFFA886C1 filesyncfalwb.dll	2 days ago	Signed Microsoft Corporation	316.3 KB

Above the search results, you can see how many binaries match the search criteria and filters you selected. You can also select sorting options for the list of binaries.


Search results provide the following information about binaries in the list:

Title	Description
Icon	<p>The icon of the file in which the binary was detected. For example:</p>  <p>Click to display the Binary Preview page, which provides a more detailed summary of information about the binary than the information in this table, but less than the Binary Details page. For more information, see "Binary Preview" on page 218.</p>
Binary MD5 Hash	The MD5 hash value of the binary.
Seen as	The filenames that were seen for binaries that match this hash value and the last time the binary was loaded.
Size	The size of the file that contains the binary.
Signature	Shows whether the binary file is signed or unsigned.
Company	The binary file's company name
	<p>Indicates whether an existing watchlist identified the binary.</p> <ul style="list-style-type: none"> • If the icon is gray, the binary was not identified by a watchlist. • If the icon is green, the binary was identified by a watchlist. <p>Click the icon to open the watchlist.</p>
>	Click to display the Binary Analysis page, which provides details about the binary file. For more information about process analysis, see " Binary Analysis " on page 219.

Binary Preview

If you click the icon on the left of a row in the binary search results table, the **Binary Preview** page appears:

Binary Preview
✕



MD5: 769F6CFE41437B4E5291B33D0194EA19

SHA-256: E7C87E5B32B5B53932E73D8C22C3AEC4C73078BA1C95B73E71F230272778C870

[View Binary >](#)

Q Related process(es): 1 | [Find related >](#)

Signed status: Signed

Company: [Apple Inc.](#)

Product: [com.apple.Notes.datastore](#)

Description: [\(unknown\)](#)

Publisher:

Feed Information

Close

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 analysis page for the binary. For more information about binary analysis, see “[Binary Analysis](#)” on page 219.
- **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.

Binary Analysis

Use the **Binary Analysis** page to investigate a binary at a deeper level. 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 screenshot displays the Binary Analysis page for a specific MD5 hash. At the top, it shows the MD5 and SHA-256 hashes, a 'Ban this hash' button, and a 'Feed Information' section. Below this, there is a 'Seen as:' section with the file path, first seen date, status (Signed), and publisher name. There are also links for file writers, related processes, and a search option. The main content is organized into several panels: 'General Info' (OS Type, Architecture, Binary Type, Size), 'Frequency Data' (1 computers, 2 processes), 'File Version Metadata' (File Description, File Version, Original Filename, Internal Name, Company Name, Product Name, Product Version, Legal Copyright), 'Digital Signature Metadata' (Result, Publisher, Signed Time, Program Name, Issuer, Subject, Result Code), 'Observed Paths (1)' (single path), and 'Observed Hosts and Sensor Ids (1)' (macbook4, 7077).

MD5: 9396F1B52BAB27595DC42CB2392B56CE
SHA-256: AC961004560D734A9BAAC6F20924603D64D15B4727A6D253A44E4F2B7C446C29

[Ban this hash](#)

Feed Information

Seen as:
 /Applications/Firefox.app/Contents/MacOS/libmozavutil.dylib
First seen at: 2019-03-01T16:31:46.942Z (about 3 days)
Status: Signed
Publisher Name:

Q File writer(s): 0 | [Find writers >](#)
 Q Related process(es): 2 | [Find related >](#)
 Search the web: [Google >](#)

General Info

OS Type	Osx
Architecture	64 bit
Binary Type	Standalone Resource
Size	459.55 KB Download

Frequency Data

1 computers have seen this md5 recently in 2 processes.
[Download full list](#)

File Version Metadata

File Description	(unknown)
File Version	(unknown)
Original Filename	Firefox
Internal Name	(unknown)
Company Name	Mozilla Corporation
Product Name	Firefox
Product Version	48.0.1
Legal Copyright	Firefox 48.0.1

Digital Signature Metadata

Result	Signed
Publisher	
Signed Time	
Program Name	libmozavutil.dylib
Issuer	Apple Inc.
Subject	Mozilla Corporation
Result Code	0x0

Observed Paths (1)

/Applications/Firefox.app/Contents/MacOS/libmozavutil.dylib

Observed Hosts and Sensor Ids (1)

macbook4	7077
----------	------

[Download full list](#)

The **Binary Analysis** page contains data for investigating the binary. See the following sections for details:

- [“Binary Overview”](#) on page 221
- [“Frequency Data”](#) on page 221
- [“Feed Information”](#) on page 221
- [“General Info”](#) on page 222
- [“File Version Metadata”](#) on page 222
- [“Digital Signature Metadata”](#) on page 223
- [“Observed Paths”](#) on page 223
- [“Observed Hosts and Sensor IDs”](#) on page 223

Banning a Hash

Banning a hash terminates a process, if running, and blocks it from running in the future. To ban a hash, click **Ban this hash**.

See [“Banning Process Hashes”](#) on page 163 for more information.

Binary Overview

The **Binary Overview** includes the following information about the binary:

Heading	Description
MD5 Hash Value	MD5 hash value for the binary.
SHA-256 Hash Value	<p>The SHA-256 hash value for the binary.</p> <p>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 capable of generating SHA-256 hashes.</p> <p>For files 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.</p>
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.
Related Process(es)	Number of processes that have used this binary.
Search the web	Performs a Google search for the MD5 hash value of the binary.

Frequency Data

Frequency Data shows how many hosts have observed the binary with this MD5 hash value.

Feed Information

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	<p>Click 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).</p> <p>The zip file contains two files: metadata and filedata.</p> <p>The metadata file is a text file that contains a timestamp and original filename.</p> <p>For example:</p> <pre>Timestamp: 04/27/2016 09:50:56 OrigFilename: \Device\HarddiskVolume1\ProgramData\Microsoft\Windows Defender\Definition Updates\{0B1E4D3A-9612-462F-8067-B0EDCE49CBF2}\mpengine.dll</pre>

File Version Metadata

File Version Metadata information 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.
Legal Copyright	Copyright details for the file, including its publisher.

Digital Signature Metadata

Digital Signature Metadata information about the binary is as follows:


Heading	Description
Result	The status of the binary signature: either Signed or Unsigned.
Publisher	The name of 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.

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 along with the ID number of the sensor.

To download a list of all hosts that have used this binary, click the **Download** icon .

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 more complex queries. The fields, field types, and examples here focus on queries to search for processes and binaries, but most of the syntax descriptions here also apply to alerts and threat reports.

Sections

Topic	Page
Query Syntax Details	225
Fields in Process and Binary Searches	227
Fields in Alert and Threat Report Searches	233
Field Types	235
Searching with Multiple (Bulk) Criteria	242
Searching with Binary Joins	242
Example Searches	245

Query Syntax Details

CB Response supports multiple types of operators and syntax that can be used to form more 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 more 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 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 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 would be “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 using the backslash (\).

For example:

This input: `microsoft\ office\office15\powerpnt.exe`

Is interpreted as: `microsoft office\office15\powerpnt.exe`

See “[path](#)” on page 238 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 (\) as the escape character.

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"
```

would be 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 185 for more information.

Fields in Process and Binary Searches

This section contains a complete list of fields that are searchable in CB Response process and/or binary searches. Some fields are valid in only one of the two, some in both. Any binary-related field used in the process search 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 capable of generating SHA-256 hashes.

For files 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-NotTerminatedCriticalSystemProcess 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.

Field	Process Search	Binary Search	Field Type	Description
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.
crossproc_type	x (def)		processopen remotethread processopentarget remotethreadtarget	<ul style="list-style-type: none"> • 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.

Field	Process Search	Binary Search	Field Type	Description
digsig_result	x	x (def)	sign	If digitally signed, the result. Values are: <ul style="list-style-type: none"> • “Bad Signature” • “Invalid Signature” • “Expired” • “Invalid Chain” • “Untrusted Root” • “Signed” • “Unsigned” • “Explicit Distrust”
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	Values are True or False - Indicates whether process has EMET mitigations configured/ enabled.
has_emet_event	x	-	bool	Values are 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.

Field	Process Search	Binary Search	Field Type	Description
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. Note that only a remote (destination) IP address is searchable regardless of incoming or outgoing.
ipv6addr	x	-	ipv6addr	Network connection to or from this IPv6 address. Note that 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).
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.

Field	Process Search	Binary Search	Field Type	Description
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 .
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.

Field	Process Search	Binary Search	Field Type	Description
sensor_id	x	-	long	The internal CB Response sensor guid of the computer on which this process was executed.
server_added_timestamp	-	x	datetime	The time this binary was first seen by the server.
special_build	x	x (def)	text	Special build string from the class FileVersionInfo .
start	x	-	datetime	Start time of this process in the computer's local time.
tampered	x	x	bool	Values are True or False - indicates when attempts are 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	The 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 253.
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.
created_time	datetime	Creation time of the alert.
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.
md5	md5 (def)	MD5 of the process that triggered the alert.
sha256	sha256 (def)	SHA-256 of the process that triggered the alert (if available).
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 this process.
report_score	float	Report score of the feed that triggered the alert. For more information, see "Threat Intelligence Feed Scores" on page 253.
resolved_time	datetime	Time this alert was triaged by a resolution action.
status	keyword	Status of the alert: one of "resolved", "unresolved", "in progress", "false positive".
username	keyword (def)	Username in whose context the process that triggered the alert event was executed.
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 name (for feeds).
create_time	datetime	Date and time this feed report was created.
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	The numeric value of a feed.

Field	Field Type	Description
feed_name	keyword (def)	The name a feed.
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)	An MD5 IOC value in the feed report.
sha256	sha256 (def)	A SHA-256 IOC value in the feed report.
report_id	keyword	Name or unique identifier of the threat report that is part of the field.
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.

Field Types

domain

Domains are split into labels for query purposes. For example, “foo.com” is split into “foo” 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

Search	If domain is foo.com	If domain is foo.com.au
domain:.com.	no match	match
domain:com.	no match	no match
domain:foo.	match	match
domain:foo.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 presumed 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

will be interpreted as microsoft AND visual AND studio AND 2010.

Searches for any one of these strings will match on the binary. Phrase queries for any two consecutive terms will 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 $\geq X$ and $\leq Y$. 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:ss` 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.
```

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 238), 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
```

For more information on control sets, see <https://support.microsoft.com/en-us/kb/100010>.

- 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 241.

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 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 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, and may 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,executescheduledspc
```

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
executescheduledspcc
```

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 recorded on the server to manage bandwidth and processing costs (see [“Advanced Settings”](#) on page 113). As part of this feature, the process cmdline field for parent processes store also store the cmdlines of any 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 believed 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, no parentheses).
- You must use the “ipaddr:” prefix to successfully use a list of IP addresses in a bulk search.

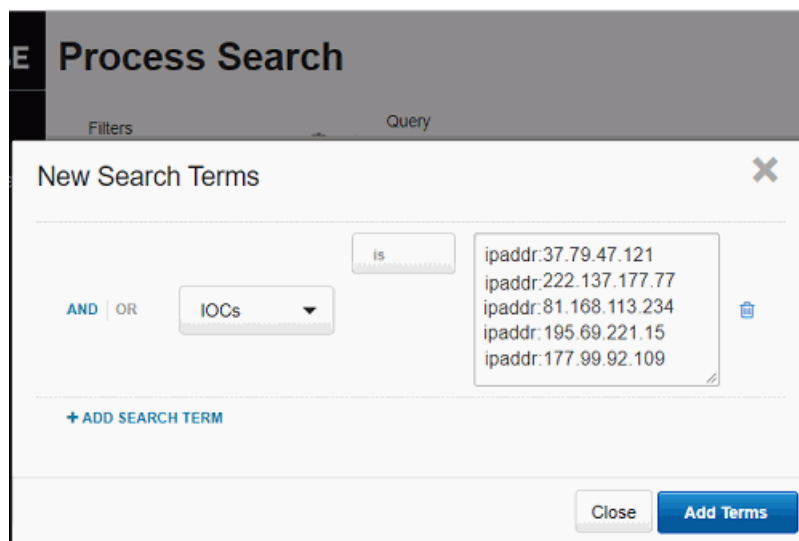
- For most other types of data, such as md5s, prefixes are optional but more efficient. See [“Fields in Process and Binary Searches”](#) on page 227 for a table of search criteria types and their prefixes.

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 will be dealt with as individual numbers rather than the four-part address.

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 that you want to include in your search, click the **Reset Search** button under the search box to make certain you are starting with a fresh search.
2. Click **Add Search Terms** under the search box.
3. In the New Search Terms dialog, use the **Choose Criteria** menu to choose **Bulk IOC > IOCs**.
4. In the text box to right of the IOCs, type or paste the list of IOCs you want to search for, making sure they meet the syntax requirements described in this section. For example:

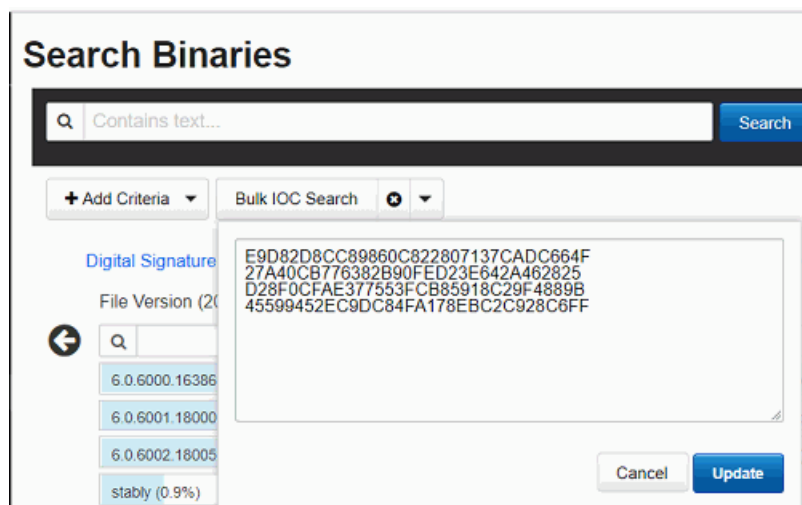


5. Although for most search criteria, you are likely to be interested in records that match one of the items on your list, you also can choose to get results that do not match your terms. You can use the **is / is not** toggle in the dialog to make this choice.
6. If you want to include additional search criteria, you can click the **Add Search Term** link in the dialog to continue defining terms.
7. 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 using the search facet tables or additional manually entered 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 that you want to include in your search, click the **Reset Search Terms** button in the upper right of the page to make certain you are starting with a fresh search.
2. Click the **Add Criteria** button under the search box.
3. In the criteria dialog that appears, under **Bulk search** check the box for **IOCs**.
4. In the text box that appears, type or paste the list of IOCs you want to search for, making sure they meet the syntax requirements described in this section. For example:



5. Click the **Update** link in the dialog to apply the search terms.

Your search is initiated and the results (if any) 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 additional manually entered 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 227.)

In this case, the results returned are process instances 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 backed by an MD5 that is unsigned. 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 appending 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 185 if you need to modify this behavior.

Example Searches

Process Search Examples

Example Query Strings	Result
<code>domain:www.carbonblack.com</code>	Returns all processes with network connections to or from domains matching the given FQDN.
<code>domain:.com</code>	Returns all processes with network connections to or from domains matching <code>*.com</code>
<code>domain:.com.</code>	Returns all processes with network connections to or from domains matching the form <code>*.com.*</code>
<code>domain:www.</code>	Returns all processes with network connections to or from domains matching the form <code>www.*</code>
<code>domain:microsoft</code>	Returns all processes with network connections to or from domains matching <code>*.microsoft</code> OR <code>*.microsoft.*</code> OR <code>microsoft.*</code>
<code>ipaddr:127.0.0.1</code>	Returns all processes with network connections to or from IP address <code>127.0.0.1</code>
<code>ipaddr:192.168.1.0/24</code>	Returns all processes with network connections to or from IP addresses in the network subnet <code>192.168.1.0/24</code>
<code>ipv6addr:fe00:b9:266:2011:28dc:43d4:3298:12e2</code>	Returns all processes with network connections to or from IPv6 address <code>fe00:b9:266:2011:28dc:43d4:3298:12e2</code>

Example Query Strings	Result
ipv6addr:fe00:b9:266:2011::0/50	Returns all processes with network connections to or from IPv6 addresses in the range of network subnet <code>fe00:b9:266:2011::0/50</code>
modload:kernel32.dll	Returns all processes that loaded a module <code>kernel32.dll</code> (accepts path hierarchies).
modload:c:\windows\system32\sxs.dll	Returns all processes that loaded a module matching path and file <code>sxs.dll</code> (accepts path hierarchies).
path:c:\windows\system32\notepad.exe	Also returns all processes with the matching path (accepts path hierarchies).
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 (as shown above). 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.

Example Query Strings	Result
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.
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).

Example Query Strings	Result
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).
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
<code><type>_desc:"database support"</code>	Returns all binaries that have <code><type>_desc</code> field with matching text, where type is file or product.
<code>legal_<type>:Microsoft</code>	Returns all binaries with matching <code>legal_<type></code> field text, where type is trademark or copyright.
<code><type>_build:"Public version"</code>	Returns all binaries with matching <code><type>_build</code> field text, where type is special or private.
<code>is_executable_image:True or False</code>	Boolean search (case insensitive) returning all binaries that are executable or not executable.
<code>is_64bit_:True or False</code>	Boolean search (case insensitive) returning all binaries that are 64-bit or not 64-bit.
<code>watchlist_4:[2014-04-01 TO 2014-09-31]</code>	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 251.

`<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 14, “Threat Intelligence Feeds”](#).

Example Query Strings	Result
<code>alliance_score_<feed>:*</code>	Returns all binaries that have <code><feed></code> score > 0 .
<code>alliance_score__score_<feed>:10</code>	Returns all binaries that have <code><feed></code> score = 10.
<code>alliance_score__score_<feed>:[10 TO 20]</code>	Returns all binaries that have <code><feed></code> score ≥ 10 and ≤ 20 .
<code>alliance_score__score_<feed>:[10 TO *]</code>	Returns all binaries that have <code><feed></code> score ≥ 10 .
<code>alliance_score__score_<feed>:[* TO 10]</code>	Returns all binaries that have <code><feed></code> score < 10 .

Chapter 14

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	
Overview of Threat Intelligence Feeds	252
Managing Threat Intelligence Feeds	253
Enabling, Disabling, and Configuring a Feed	260
On-Demand Feeds from CB Threat Intel	262
Creating and Adding New Feeds	263
Searching for Threat Reports	265

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 may 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
- 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 262 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:

- In the navigation bar, select **Threat Intelligence**.

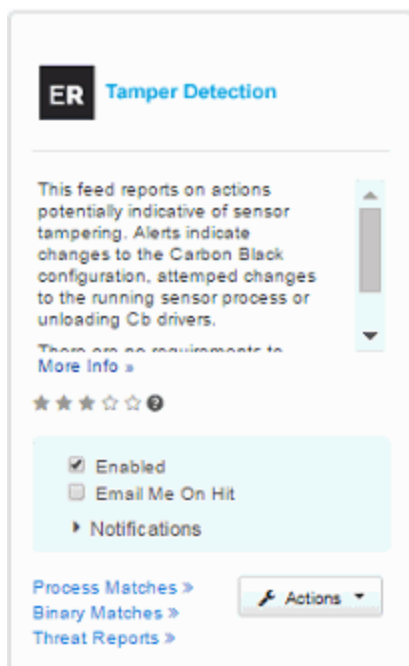
The **Threat Intelligence Feeds** page appears, as in the following example.

The screenshot displays the 'Threat Intelligence Feeds' interface. At the top, there are buttons for 'Threat Report Search' and '+ Add New Feed'. Below this, a grid of eight feed cards is shown, each with a logo, title, description, and configuration options.

Feed Name	Description	Requirements
COMMUNITY	This is a feed containing Carbon Black community produced detection queries. These queries have been publicly posted to the Carbon Black UserExchange site in the Detection Exchange group.	None
Bit9 + CARBON BLACK SOFTWARE REPUTATION TRUST	The Bit9 Software Reputation Service (SRS) feed provides a level of software trustworthiness. It is necessary to share MD5s of observed binaries with the Carbon Black Alliance to use this feed.	MD5s of observed binaries
REPUTATION THREAT	The Cb Reputation Threat feed is sourced by Carbon Black and provides an assessment of the risk associated with hashes in your environment. This feed is customer specific, and the installation must share Binary Hashes & Metadata with Carbon Black.	Binary Hashes & Metadata
National Vulnerability Database	NVD is the U.S. government repository of standards based vulnerability management data. This feed will flag executed applications vulnerable to one or more CVEs with CVSS scores higher than 7.0 from 2013-2015 for Java, Flash Player and Google.	None
abuse.ch	Abuse.ch tracks C&C servers for Ransomware, Zeus, and Feodo malware. This feed combines information from the IP, Domain and Binary blocklists.	None
ADVANCED THREAT	This feed is a list of high-confidence threat indicators, updated periodically. Generally, hits on this feed should be suitable for generating alerts. There are no requirements to share any data to receive this feed.	None
Open Threat Exchange (OTX)	This feed contains intelligence provided by AlienVault's Open Threat Exchange (OTX). It leverages insights into attacks across the community and will show you hostile scanning hosts, malware hosts, and other targeting and security event.	None
Banning Events	This feed reports on Carbon Black process blocking events due to MD5 hash based banning rules on the endpoint. There are no requirements to share any data to receive this feed.	None

Each feed card includes a 'More Info' link, a star rating, and configuration options: 'Enabled' (checked), 'Email Me On Hit' (unchecked), and 'Notifications'. At the bottom of each card are links for 'Process Matches', 'Binary Matches', and 'Threat Reports', along with an 'Actions' button.

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 you want to use. See [“Enabling, Disabling, and Configuring a Feed”](#) on page 260 for information on enabling and configuring a feed.

See [“Creating and Adding New Feeds”](#) on page 263 for information on 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 as needed.

The **Check for new feeds** option on the **Threat Intelligence Feeds** page allows you to check for the most recent feeds. It also removes feeds if their source no longer provides them. (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, choose **Check for new feeds** on the action (down arrow) menu in the top-right corner of the page.
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 with the latest information from their source, use the **Sync All** command.

To sync all threat intelligence feeds on the page:

1. On the **Threat Intelligence Feeds** page, choose **Sync All** on the action (down arrow) menu in the top-right corner of the page.
2. All feeds are synced with the latest data on the **Threat Intelligence Feeds** page.



Note

You can sync feeds individually rather than all at one time. Sync commands for individual feeds are on the **Action** menu in the feed panel. See [“Syncing Threat Intelligence Feeds”](#) on page 256 for information about these options.

Data Sharing Settings

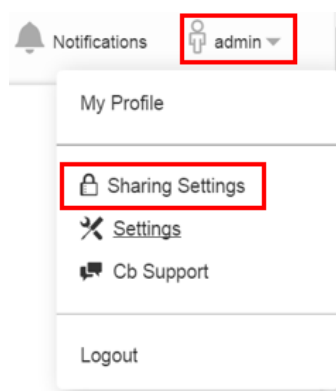
Most of Carbon Black's threat intelligence feed partners provide a list of all of the IOCs they track, and almost all feeds require that you enable communication on the Sharing Settings page. Also, some feeds require that you enable data sharing.

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. Select **username** > **Sharing Settings**.



The Sharing page appears.

2. Under General Sharing Settings, specify sharing settings for Alliance:
 - **Enable Alliance Communication** – When selected, 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** – When selected, enables your server to send threat intelligence statistics to Carbon Black, including alert resolutions, ignored reports, and feed ratings. These statistics are used to 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** – When selected, enables sharing of usage, resource, and sensor statistics with Carbon Black.
 - **Allow Unattended Background Upload of Diagnostics Data** – When selected, enables the CB Response server to do background collection of diagnostics data such as application logs and configuration files in order 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** – When selected, no sensor diagnostics data can be uploaded to Carbon Black.
 - Manual** – When selected, you can upload sensor diagnostics data manually by using a utility installed on the sensor.
 - Automatic** – When selected, 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 down to **Endpoint Activity Sharing**:

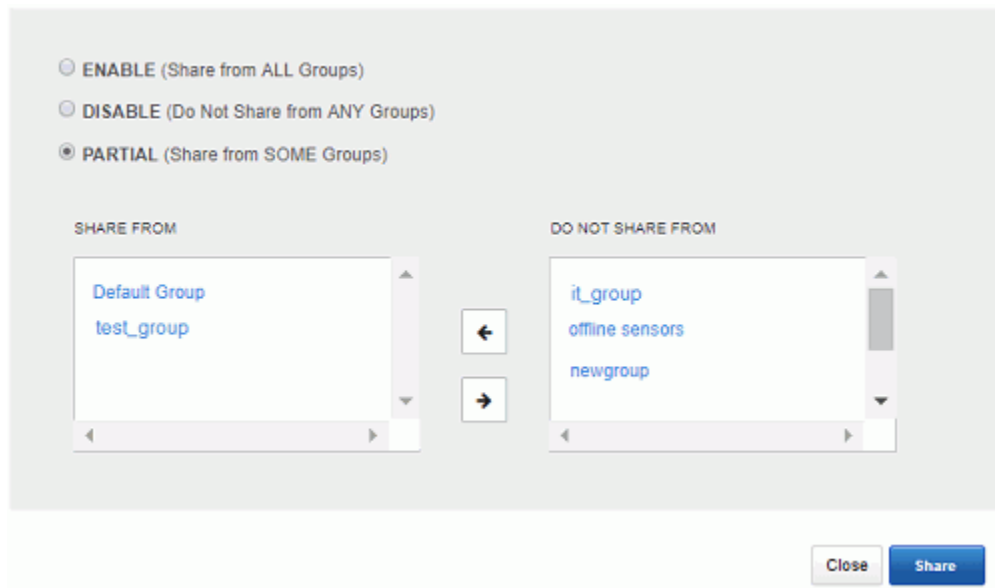
	Carbon Black	Cb Inspection
Binary Hashes & Metadata	PARTIAL	N/A
Complete Binaries	PARTIAL	DISABLED
Response Event Data	PARTIAL	N/A

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 from sensor groups 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**) for a data type 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 described 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.

The screenshot shows the 'Threat Intelligence Feeds' page. At the top, there are buttons for 'Threat Report Search' and 'Add New Feed'. Below this, there are eight panels, each representing a different threat intelligence feed. Each panel includes a logo, a title, a brief description, a 'More Info' link, a star rating, and a set of controls. The controls include a checkbox for 'Enabled', a checkbox for 'Email Me On Hit', and a 'Notifications' link. At the bottom of each panel, there are links for 'Process Matches', 'Binary Matches', and 'Threat Reports', along with an 'Actions' dropdown menu.

To enable and configure a threat intelligence feed:

1. In the navigation bar, select **Threat Intelligence**.
The Threat Intelligence Feeds page opens.
2. Locate the feed to enable, and on that panel, select **Enabled**.
3. Configure the feed using the following options and controls:

Field/Menu	Description
More info	Link to the feed provider's website. It provide technical information about the feed or general information about the provider and its products.
★★★★☆ (Rating)	Rating of this threat intelligence feed by the community of CB Response users. The default for all ratings is three stars. You can click on 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 sorted by severity.
Enabled	If selected, the threat intelligence feed is enabled; otherwise it is disabled. Note: Most feeds also require 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 256.
Email Me on Hit	If selected, IOCs from this feed that reference a process or binary 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 303. Only Global Administrators or cloud Administrators can change this setting.
Notifications menu	Additional notification options: <ul style="list-style-type: none"> • Create Alert – If this check box is selected, indicators from this feed that reference a process or binary recorded on this CB Response server cause a console alert. See "Enabling Console Alerts" on page 293. • Log to Syslog – If this check box is selected, IOCs from this feed that reference a process or binary recorded on this CB Response server are included in Syslog output from this CB Response server. See the <i>CB Response Integration Guide</i> on the Carbon Black User eXchange for further 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," for information.
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," for more information.
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 265 for more information.

Field/Menu	Description
Actions menu	<p>The Actions menu includes the following commands:</p> <ul style="list-style-type: none"> • Create Watchlist – Creates a watchlist, which is a saved search whose results will be processes or binaries matching IOCs reported by this feed. • 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. In the navigation bar, select **Threat Intelligence**.
2. On the Threat Intelligence Feeds page, identify the feed to disable.
3. Deselect the **Enabled** check box within that feed panel to disable it.

If you disable a feed, its reports remain on the server and data coming in will be tagged against the locally existing IOCs that it reported. However, the following things will not occur 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 patterns of compromise associated with them, 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 the data for a process 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

In order 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, and can be enabled in the Response Event Data row in the Endpoint Activity Sharing section of the Sharing page. See [“Data Sharing Settings”](#) on page 256.

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 itself must match the feed structure (or schema) defined in the "Feed Structure" section of the CBFAPI documentation.

You have several options regarding the amount of 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. In 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, select 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.
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: <ul style="list-style-type: none"> Username Password Public Cert Private Key

- To add a feed manually, select 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 will be 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: <ul style="list-style-type: none"> Username Password Public Cert Private Key

- When you finish entering the settings for this feed, 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 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):

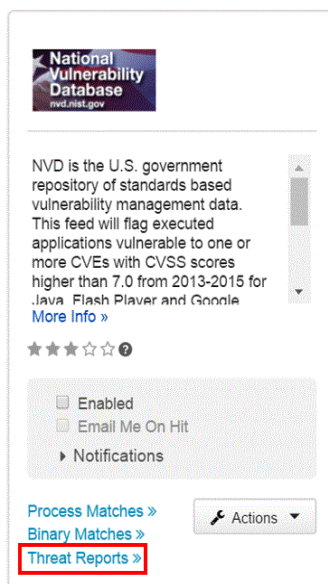
1. In the navigation bar, click **Threat Intelligence**.
2. On the Threat Intelligence Feeds page, click **Threat Report Search**.



3. In the Search Threat Reports page, you can enter search criteria to search for the reports in which you are interested. See [“Threat Report Searches and Results”](#) on page 266 for more details about the page.

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 whose reports you want to view.



2. The Search Threat Reports page displays reports from the selected feed.
You can further refine the search if needed 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 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 screenshot shows the 'Search Threat Reports' interface. At the top, there is a search bar with the text 'Contains text...' and a 'Search' button. To the right, there are links for 'Cb Support', 'Notifications', and a user profile 'admin2'. Below the search bar is a '+ Add Criteria' button and a '9 Matching Reports' indicator. The interface features four filter sections: 'Feed Name (6)' with a dropdown menu showing 'cbtamper (44.4%)', 'binarytestfeed (11.1%)', 'cbbanning (11.1%)', and 'cbemet (11.1%)'; 'Feed Category (0)'; 'Report Score' with a bar chart showing a peak at score 100; and 'Report Creation Time' with a bar chart showing reports from 07/2014 to 07/2016. Below the filters, there is a 'Show 9 of 9' indicator and a 'Sort by Severity' dropdown. The main content is a table with the following data:

	Description	Indicators	Report Score	Ignore	
	hashtestfeed Hash test feed Updated 3 years ago	MD5s: 2, SHA-256s: 2, IPs: 0, Domains: 0, Queries: 0	100	No <input type="checkbox"/>	Details >
	cbbanning Carbon Black Process Blocking Updated 5 years ago	MD5s: 0, SHA-256s: 0, IPs: 0, Domains: 0, Queries: 0	90	No <input type="checkbox"/>	Details >

The Search Threat Reports page presents report data in formats similar to the presentation of other data in the Process Search and Binary Search pages, as follows.

Column	Description
Description	This column includes: <ul style="list-style-type: none"> • 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: <ul style="list-style-type: none"> • 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 269.
Details link	Opens a Threat Report Details page for the report in this row. See “Threat Report Details” on page 268.

Threat Report Details

When you click on the **Details** link in the far right column of a report in the Threat Report Search Reports table, a new page shows details for that threat report, if available.

Search Threat Reports Cb Support | Notifications | admin2

CVE-2016-4240 flash (cvss_score: 10.0)
National Vulnerability Database

Report Details

ID CVE-2016-4240

Link <http://nvd.nist.gov/vuln/detail?vulnId=CVE-2016-4240>

Updated Tue Jan 10 2017 21:16:48 GMT-0500 (Eastern Standard Time)

IPs none

MD5s 13

SHA-256s none

Domains none

Queries none

Report Tags
None

Feed Description

NVD is the U.S. government repository of standards based vulnerability management data. This feed will flag executed applications vulnerable to one or more CVEs with CVSS scores higher than 7.0 from 2013+ for Java, Flash Player and Google Chrome applications.

Report Description

Adobe Flash Player before 18.0.0.306 and 19.x through 22.x before 22.0.0.209 on Windows and OS X and before 11.2.202.632 on Linux allows attackers to execute arbitrary code or cause a denial of service (memory corruption) via unspecified vectors, a different vulnerability than CVE-2016-4172, CVE-2016-4175, CVE-2016-4179, CVE-2016-4180, CVE-2016-4181, CVE-2016-4182, CVE-2016-4183, CVE-2016-4184, CVE-2016-4185, CVE-2016-4186, CVE-2016-4187, CVE-2016-4188, CVE-2016-4189, CVE-2016-4190, CVE-2016-4217, CVE-2016-4218, CVE-2016-4219, CVE-2016-4220, CVE-2016-4221, CVE-2016-4233, CVE-2016-4234, CVE-2016-4235, CVE-2016-4236, CVE-2016-4237, CVE-2016-4238, CVE-2016-4239, CVE-2016-4241, CVE-2016-4242, CVE-2016-4243, CVE-2016-4244, CVE-2016-4245, and CVE-2016-4246.

Report Score

100 HIGH THREAT -100 100

Ignore this Report? No

Type	Indicator	
MD5	19C4715DB06D0650955935CD143256183	Process Search >> Binary Search >>
MD5	1F0207FA58544C83B506F1DE00D90259	Process Search >> Binary Search >>

The information on the Threat Report Details page varies depending on the feed source and type of indicator in the report. The following table describes the fields in this page.

Field	Description
Title	The feed name and the unique ID of the report.
Report Details	<p>This section includes:</p> <ul style="list-style-type: none"> 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 IPs – the number of suspicious IP addresses MD5s – the number of suspicious MD5s SHA-256s – the number of suspicious SHA-256s Domains – the number of suspicious domains Queries – the number of queries in the report; this is always 1.

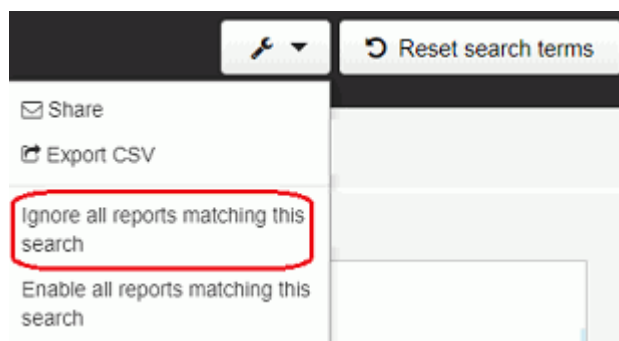
Field	Description
Report Tags	One or more descriptive strings from the feed provider to help explain what the report is about. For example, tags can describe a specific threat, a threat category, a targeted industry, a known threat actor, or geographic information. Not all reports have tags.
Report Indicators	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.
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.
Ignore this Report?	Ignore any future instances of this report so that they do not trigger alerts.

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 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 specify that you want any future instances of reports by the same name to be ignored.

To mark reports as ignored, use one of the following options:

- On the Search Threat Reports page, on the **Actions** menu, you can 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 301.

Chapter 15

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	271
Creating Investigations	274
Creating Investigations	274
Adding Events to Investigations	275
Removing Events from Investigations	276
Adding Custom Events to Investigations	276
Deleting Investigations	277

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 to your systems. 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.

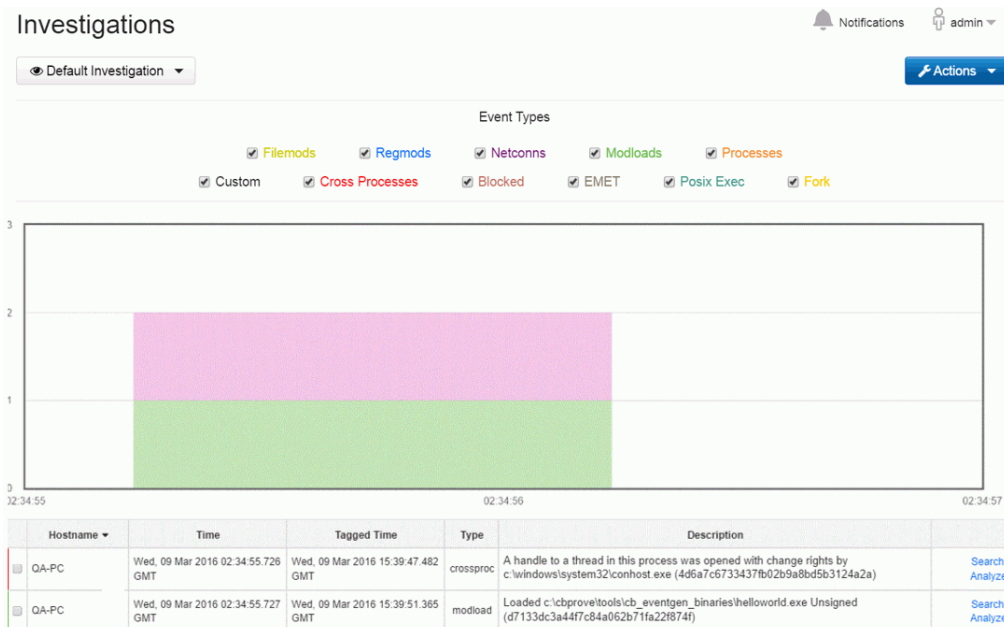
The first time that you open the **Investigations** page, the default investigation appears.

Viewing Investigations

To view investigations:

- In the navigation bar, select **Investigations**.

The following figure shows an example of a default investigation:



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 276.
 - **Add Custom Event** – See [“Adding Custom Events to Investigations”](#) on page 276.
 - **Add Investigation** – See [“Creating Investigations”](#) on page 274.
 - **Delete Investigation** – See [“Deleting Investigations”](#) on page 277.
 - **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 from the rows at the bottom of the page to a csv file and downloads it to your computer.

Event Types

Select/deselect the 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 202.


From left to right on the **Investigations** page, 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 will generate a Blocked event. Color-coded as brown.
- **EMET** – Represents a specific type of event the sensor captures that deal with Microsoft’s Enhanced Mitigation Experience Toolkit (EMET) software that can possibly be installed on a Windows endpoint. Color-coded as gray.
- **Fork** – (OS X and Linux only) The instance’s parent process, forked with a different Process ID (PID). Color-coded as yellow orange.
- **Posix_Exec** – (OS X and Linux only) The instance’s process that is loaded and the new binary image. Color-coded as green.

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 appear 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.




Events Table

The events table shows the events contained in the investigations. A colored bar appears on the left border of each row to indicate the event type. The following table describes the information that appears:

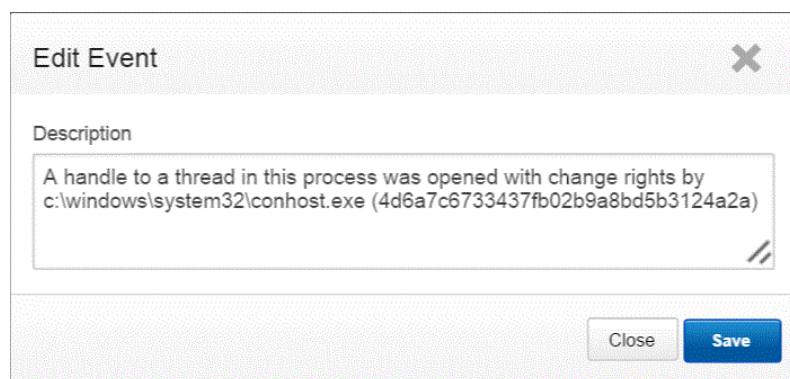
Column	Description
Hostname	The name of the host on which the event occurred.
Time	The date and time that the event occurred.
Tagged Time	The time that the event was tagged for this investigation.
Type	The event type (filemod, regmod, netconn, modload, child process, fork, posix_exec, custom, crossproc, 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 each row, an **Edit** icon appears.

Hostname	Time	Tagged Time	Type	Description	
QA-PC	Wed, 09 Mar 2016 02:34:55 726 GMT	Wed, 09 Mar 2016 15:39:47 482 GMT	crossproc	A handle to a thread in this process was opened with change rights by c:\windows\system32\conhost.exe (4d6a7c6733437fb02b9a8bd5b3124a2a)	 Search > Analyze >
QA-PC	Wed, 09 Mar 2016 02:34:55 727 GMT	Wed, 09 Mar 2016 15:39:51 365 GMT	modload	Loaded c:\cbprove\tools\cb_eventgen_binaries\helloworld.exe Unsigned (d7133dc3a447c84a062b71fa22f874f)	Search > Analyze >

Clicking the **Edit** icon displays 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 made to a description are visible within the investigation, but do not appear in the process execution data when viewed outside of the **Investigations** page.

Child Processes

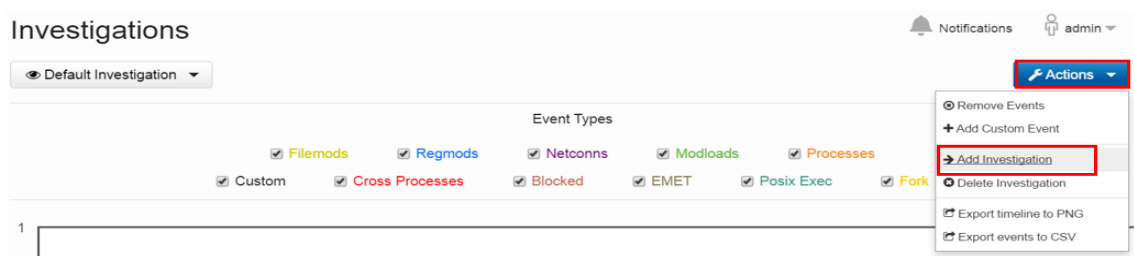
Rows that represent child processes contain a **Search** (magnifying glass) icon. This option displays a preview of what you would see if you chose to open the **Process Analysis** page for the child process.

Creating Investigations

You can create customized investigations and then add events to these as discussed in [“Adding Events to Investigations”](#) on page 275.

To create an investigation from the Respond menu:

1. In the navigation bar, select **Investigations**.
2. The **Default investigation** appears.
3. Click **Actions > Add Investigation**.



- The **Add Investigation** window appears:



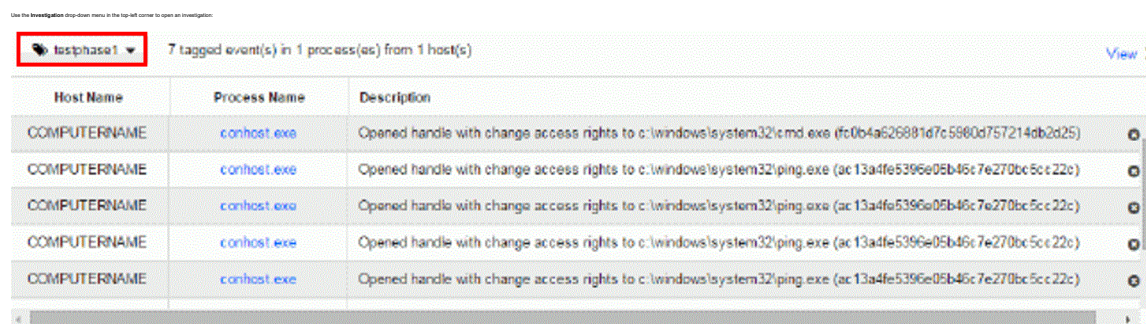
- In the **Name** field, enter a name for the investigation and click **Save**.
Note: The name must be alpha-numeric. Special characters are not allowed.
- The new investigation appears in the **Investigations** window but will be empty until you add events to it. See [“Adding Events to Investigations”](#) on page 275.

Adding Events to Investigations

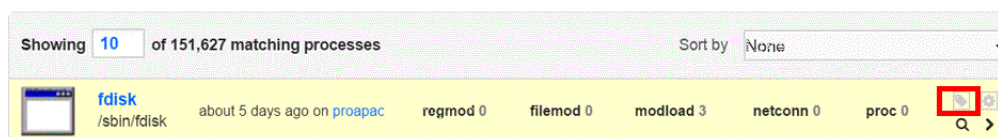
While you are performing searches for process executions or binary files, you can use the **Investigations** icon to have an investigation continually open (see [“Creating Investigations”](#) on page 274). Events that you tag in your search results are added to this investigation.

To add events from process or binary searches to investigations:

- Click the **Investigations** icon (discussed in [“Creating Investigations”](#) on page 274) to keep the investigation open.



- Select a process from the **Process Search** page (see [“Overview of Process Search”](#) on page 176) and open the **Process Analysis** page ([“Process Analysis Page”](#) on page 190).
- Click the **Tag** icon in an event row that you would like to add to the investigation (the **Tag** icon changes from gray to blue). The events are automatically added to the investigation that you have open.




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. You can remove an event from an investigation in one of two ways:

- Use the full **Investigations** page (accessed by selecting **Investigations** in the console menu).
- Use the smaller version of the **Investigations** page by clicking the **Investigation** icon (discussed in “[Creating Investigations](#)” on page 274).

To remove an event from an investigation from the Respond menu:

1. In the navigation bar, select **Investigations**.
2. The **Default investigation** appears.
3. Use the **Investigation** drop-down menu in the top-left corner to open an investigation from which to remove an event.
4. Click **Actions > Remove Events**.



5. The event is removed from the list at the bottom of the page.

Adding Custom Events to Investigations

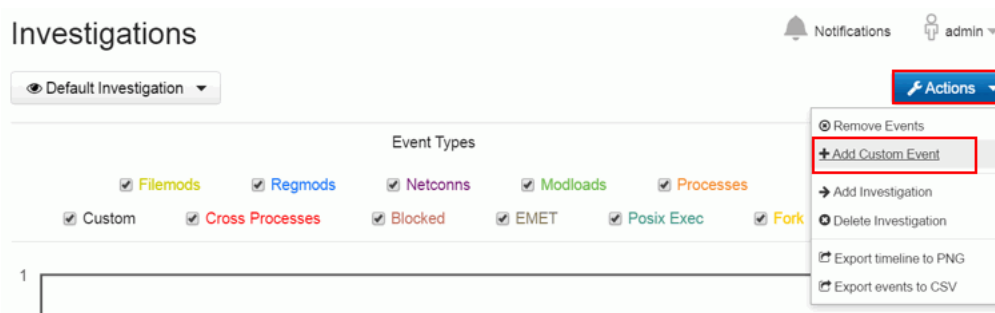
You can create a custom event that you can use to:

- 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 where you want it to in the timeline.

To create custom events:

1. In the navigation bar, select **Investigations**.
2. Click **Actions > Add Custom Event**.



3. The **Add Custom Event** window appears:

The screenshot shows a dialog box titled 'Add Custom Event' with a close button (X) in the top right corner. Inside the dialog, there is a 'Description' field with a text input area and a small icon in the bottom right corner. Below the description field is a 'Start Time' field with a date and time picker icon. At the bottom of the dialog, there are two buttons: 'Close' and 'Save'.

4. In the **Description** field, type a description for the event.
5. In **Start Time**, enter the date and time for the event.
6. Click **Save**.

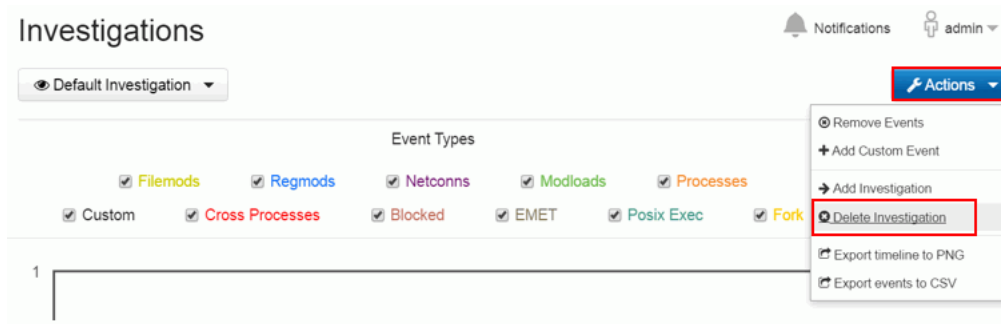
Deleting Investigations

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 how those processes appear in other pages.

To delete investigations:

1. In the navigation bar, select **Investigations**.
2. In the **Investigations** menu, select the investigation to delete.
The investigation opens.

3. Click **Actions > Delete Investigation**.



4. Confirm your decision to remove the investigation by clicking **OK**.

Chapter 16

Watchlists

This chapter describes creating, using and managing watchlists. Watchlists are saved searches visible to all users and run periodically against the process or binary data in CB Response.

Sections

Topic	Page
Overview	280
Viewing Watchlists and their Results	280
Built-in and Community Watchlists	283
Creating Watchlists	284
Managing Watchlists	289
Editing Watchlists	291
Deleting Watchlists	291

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 one of several configurable ways.

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 used to calculate the severity assigned to IOCs from a feed.

In addition to this chapter, 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 14, “Threat Intelligence Feeds”](#)
- [Chapter 17, “Console and Email Alerts”](#)

Viewing Watchlists and their Results

Click **Watchlists** in the main console navigation bar 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 on its name in the table.

The Watchlists Table

The screenshot displays the Watchlists page. At the top, there is a search bar labeled "Search...". Below it, on the left, is a table of watchlists. The table has columns for watchlist name, status, and type. The watchlists listed are:

Watchlist Name	Status	Type
Autoruns	Disabled	Process
Chrome helper instances	Enabled	Process
Chrome running on lab mac	Enabled	Process
Filemods to Webroot	Enabled	Process

On the right side, the details for the "Chrome helper instances" watchlist are shown. It includes a description: "Tell me when users run chrome helper applications." and a query string: "cb.uriver=1&rows=10&facet=false&facet.field=process_name&facet.field=group&facet.field=hostname&facet.field=parent_name&facet.field=path_full&facet.field=process_md5&cb.query_source=ui&start=0&cb.fq.process_name=Gooc".

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 the watchlists that appear in the table:

- At the top of the Watchlists page, you can use the **Search** box to search for watchlists by name.
- Immediately above the table of Watchlists on the left, filters and sorting controls can be used to modify what is shown in the table. In the **Show** field, you can choose 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 289 for information on how you might use these features to more effectively manage watchlists.

The Watchlist Details Panel

Newly Loaded Modules Disable Delete

Last successful watchlist execution was 6 minutes ago, duration was 0.03 s.

DESCRIPTION

ON HIT

Email Me

Create Alert

Log to Syslog

QUERY

```
cb.uriver=1&q=is_executable_image:false&sort=server_added_timestamp desc
```

HITS OVER TIME

RESULTS [Search >](#)

BINARY	PATH	SIGNATURE	FIRST SEEN
83D16ECD1E00C87BC1365A6FC43BC...	/bin/mount	Unsigned	a day ago
635E77AE4EEB2FFB4D1EB495D286F663	/sbin/runlevel	Unsigned	a day ago

The Watchlist Details panel on the right side of the Watchlists page shows details for the watchlist currently selected in the table to the left. It includes the following information:

- the 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 execute every 10 minutes, but if a previous watchlist session is still running, the next one will be delayed and will try to start periodically (every 10 minutes).
- the Query it uses to match events to the watchlist
- the On Hit settings that determine how (or if) you are notified when an event matches the query
- a graph showing the number of hits on this watchlist over time
- a table of results showing details for each hit

Note

For each watchlist run, the number of matching events that are tagged is limited to 100, even if more events actually match the watchlist. This limit is in place to prevent performance issues and eliminate the potential for excessive numbers of notification emails that are unlikely to add useful information.

If you click the **Search** link above the table of results, the results of the query are shown 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 in the background on your server. New results matching the search query do not result in any notification or record that they would have triggered a hit for the watchlist.

When you click the **Delete** button and confirm your choice, the watchlist is permanently removed.

Built-in and Community Watchlists


You can create your own watchlists, but CB Response provides access to two different sources of pre-configured watchlists:

- The Watchlists page in the console includes a list of default watchlists, any of which you can enable for your own use.
- The Carbon Black User eXchange provides a forum for sharing watchlists.

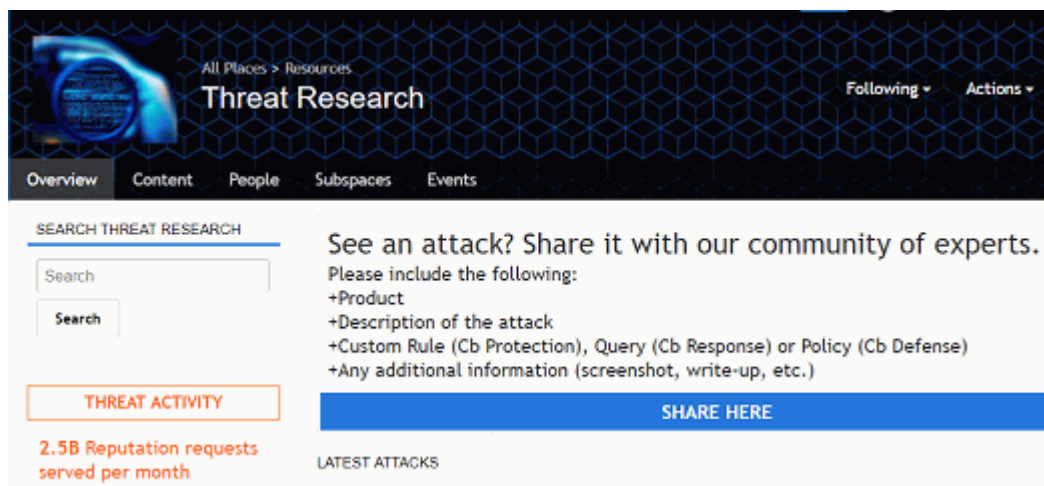
The default watchlists on the Watchlists page includes include the following:

- **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**

In addition to the Watchlists built into the server, in the top-right corner of the **Watchlists** page, you can click the **Community Watchlists** button to access the [Carbon Black User eXchange](#).



The [Carbon Black User eXchange](#) is a central portal where watchlist users can publish and discuss watchlists that might eventually be included as a feed from CB Threat Intel. You might want to use one of these in your own environment.



Creating Watchlists

If the default and community watchlists do not offer the features you need, you can create your own customized watchlists. You can create watchlists from the **Watchlists**, **Process Search**, **Binary Search**, or **Threat Intelligence Feeds** pages. The information you must provide varies depending upon which of these locations you start from.

To create watchlists from Process Search or Binary Search pages:

1. From the navigation bar, select either **Process Search** or **Binary Search** to open the appropriate search page.
2. Enter the query for the processes or binaries for which you want to create a watchlist. The syntax you use for a watchlist query should match what you would use for 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 be sure to examine the results before proceeding. For more information on editing queries, see [“Editing Watchlists”](#) on page 291.

For more information on performing searches, see:

- [Chapter 11, “Process Search and Analysis,”](#)
- [Chapter 12, “Binary Search and Analysis,”](#)
- [Chapter 13, “Advanced Search Queries.”](#)

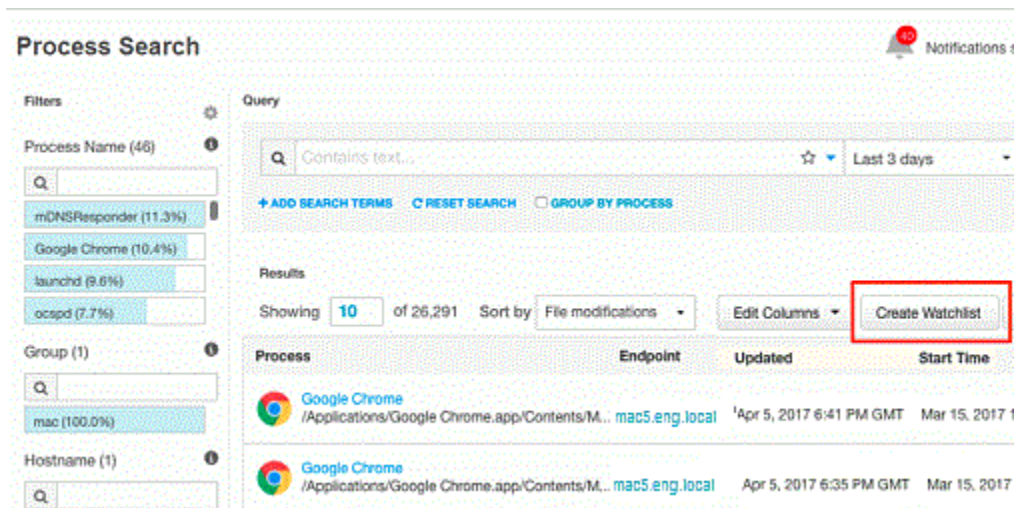
If you are using multiple MD5 or SHA-256 hash values for search criteria to create a watchlist, you must enclose the values in parentheses ().

For example:

```
(md5:45cc061d9581e52f008e90e81da2cfd9
md5:829e4805b0e12b383ee09abdc9e2dc3c
md5:ac9fa2ba34225342a8897930503ae12f
md5:5f7eaaaf5d10e2a715d5e305ac992b2a7)
```

If you do not enclose the list in parentheses, the only value that will be tagged for the watchlist is the last value in the list.

3. Click **Create Watchlist**:



4. The **Add Watchlist** window opens.

The 'Create Watchlist' dialog box is shown. It has a title bar with 'Create Watchlist' and a close button (X). The fields are:

- Watchlist Name ***: A text input field.
- Description**: A text area.
- Query ***: A text area with a note: 'The query is stored URL-encoded, and displayed in decoded form for readability. Both forms are accepted here.' The example query 'q=process_name:example' is entered.
- Try it out >**: A link.
- Options**: Three checkboxes: 'Email Me', 'Create Alert', and 'Log to Syslog'.
- Watchlist Type:** Two radio buttons: 'Process' (selected) and 'Binary'.
- Buttons**: 'Create' and 'Cancel'.

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 **Search Query** field, notice that the URL is the query that is currently open. You cannot edit this field.
8. Select the **Email Me** check box to receive email notifications when hits occur that match your search.
9. 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 17, "Console and Email Alerts."](#)
10. Select the **Log to Syslog** check box to log all hits that match the search in this watchlist to `syslog`. Syslogs are written to `/var/log/cb/notifications/`. In this case, the log filenames have the format `cb-notifications-<watchlist ID>.log`.
11. Click **Save changes**.

To create watchlists from the Watchlists page:

1. In the navigation bar, select **Watchlists**.
2. Click the **Create Watchlist** button at the top right of the page. The Add Watchlist dialog appears.

Create Watchlist X

Watchlist Name *

Description

Query *

The query is stored URL-encoded, and displayed in decoded form for readability. Both forms are accepted here.

q=process_name:example

[Try it out >](#)

Email Me

Create Alert

Log to Syslog

Watchlist Type:

Process

Binary

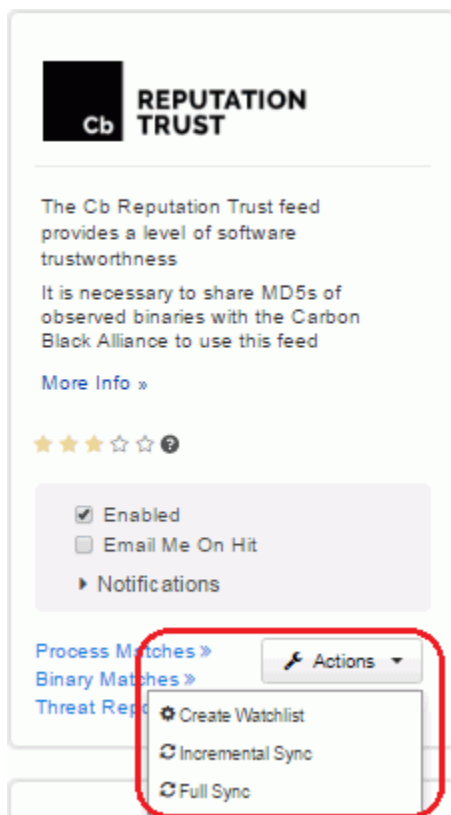
Create **Cancel**

3. Enter a **Name** and **Description**.
4. In the **Watchlist Type** field, choose the **Process** or **Binary** radio button.
5. In the **Query** field, enter the query you would like to use for this watchlist. As with watchlists created from the Process Search or Binary Search page, be sure to follow the rules for creating a properly functioning query.

- Caution:** As with searches outside of a watchlist, use of leading wildcards is discouraged because of performance issues.
6. Click the **Try it out** link to run the query in a separate browser window to see whether it works, and if so, what the results are.
 7. Select the **Email Me** check box to receive email notifications when hits occur that match your search.
 8. 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 17, "Console and Email Alerts."](#)
 9. Select the **Log to Syslog** check box to log all hits that match the search in this watchlist to `syslog`. Syslogs are written to `/var/log/cb/notifications/`. For watchlists, log filenames have the format `cb-notifications-<watchlist ID>.log`.
 10. Choose the radio button for the **Watchlist Type** you want to create: Process or Binary.
 11. Click **Save changes**.

To create watchlists from the Threat Intelligence Feeds page:

1. In the navigation bar, select **Threat Intelligence**.
The **Threat Intelligence Feeds** page opens.
2. Select the feed for which you would like to create a watchlist.
3. From the **Actions** menu, select **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 **Feed Score Criteria** section, use the various fields to enter the score criteria for the severity of IOCs to track.
8. From the **Type** menu, select **Process** or **Binary** depending on your search type.
9. Select the **Email Me** check box to receive email notifications when hits occur that match your search.
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 17, "Console and Email Alerts."](#)
11. Select the **Log to Syslog** check box to log all hits that match the search in this watchlist 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 should provide you with valuable information about conditions that matter in your environment. To make this happen, you might need to fine tune them for your environment, based on their performance and the quality of the information they provide.

You can monitor that 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 you expect, you can edit, disable, or delete it.

Watchlist Status

Watchlists will 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, 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 of the last execution appears if that execution was successful. See [“Slow or Error-producing Watchlists”](#) for more information on how this might be used in troubleshooting.

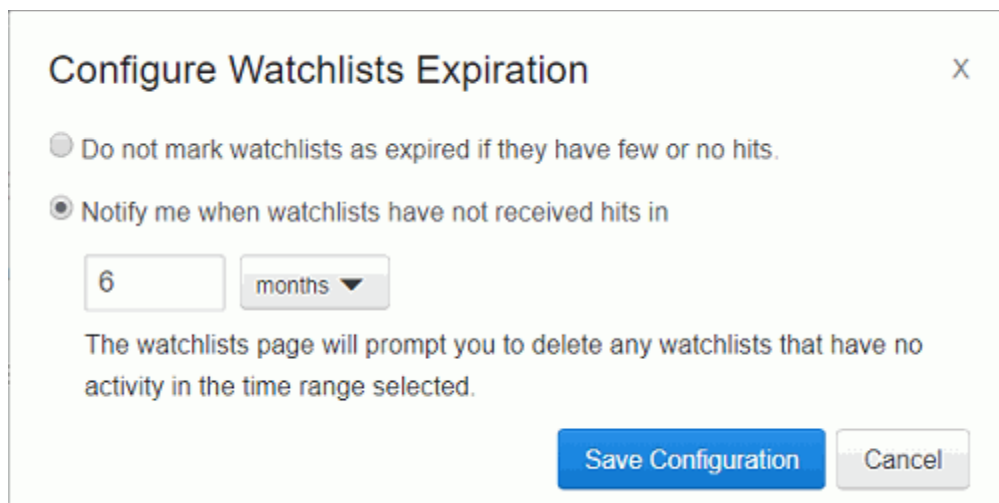
Watchlist Expiration

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 needs to 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.

There is a single watchlist expiration configuration that applies to all watchlists.

To configure watchlist expiration:

1. In the navigation bar, select **Watchlists**.
2. Click the **Configure Expiration** button at the top right of the page. The Configure Watchlists Expiration dialog appears.



Configure Watchlists Expiration X

Do not mark watchlists as expired if they have few or no hits.

Notify me when watchlists have not received hits in

6 months

The watchlists page will prompt you to delete any watchlists that have no activity in the time range selected.

Save Configuration Cancel

3. By default, watchlists are marked as “Expired” if they have received no hits over a six-month period. If you want to use a different time period for expiration or are reconfiguring 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 using 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. Since you normally are not 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.

Once 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. You can experiment with changes and use the **Try it out** button in the Watchlist Details to determine whether the query will work well outside of a watchlist.

If you have not been able 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. In the navigation bar, select **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 as appropriate, and then click **Save Changes**.

Note: Saving a modified watchlist query replaces the watchlist (that is, it 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 in this watchlist 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. From the navigation bar, select **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**.
4. Click **OK** in the confirmation window to delete the watchlist.

Chapter 17

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	293
Enabling Console Alerts	293
Viewing Alert Activity on the HUD Page	295
Managing Alerts on the Triage Alerts Page	296
Enabling Email Alerts	303

Overview of Alerts

You can create alerts that will 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 16, “Watchlists,”](#) for more information.
- **Threat intelligence feed hits** – Threat intelligence feeds can be configured to send an alert when that feed reports an IOC that has been seen on endpoints reporting to your CB Response server. See [Chapter 14, “Threat Intelligence Feeds,”](#) for more information.

Triggered alerts are reported in two locations in the CB Response console:

- The **HUD** page (accessed from the navigation bar) contains a summary showing the number of unresolved alerts, the number of hosts with unresolved alerts, and other alert-related data, including the alerts for each host. See [“Viewing Alert Activity on the HUD Page”](#) on page 295 for more information.
- The **Triage Alerts** page (accessed from the navigation bar) contains more details about triggered alerts and provides a filter and search interface to find alerts matching 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 296 for more information.

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 303 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 these to monitor endpoints for detected IOCs. You can also select the most important watchlists to monitor and add console alerts to them. Then, you can then view and manage these key watchlist and feed hits in one place, the **Triage Alerts** page.

To enable console alerts for a watchlist:

1. From the navigation bar, select **Watchlists**.
2. In the left panel of the **Watchlists** page, select the watchlist for which you want to create an alert. If you can't locate the watchlist, use the **Search** box at the top of the panel to locate it.
3. In the right panel, select the **Enable Watchlist** and **Create Alert** check boxes.

The watchlist should begin generating alerts.

Threat Intelligence Feed Alerts

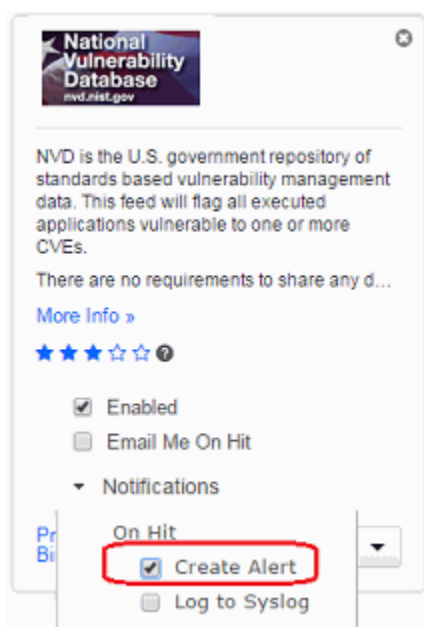
Threat intelligence feeds provide information that helps you identify malware and its sources. CB Response integrates with third-party and internal feeds (such as the CB Threat Intel Reputation and CB Response Tamper Detection) that identify hosts on which tamper attempts have occurred.

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. In the navigation bar, select **Threat Intelligence**.
2. Select the **Notifications** drop-down menu and select the **Create Alert** check box for each feed panel in which you want to enable console alerts.



To disable console alerts for a threat intelligence feed:

1. In the navigation bar, select **Threat Intelligence**.
2. Select the **Notifications** drop-down menu, and uncheck the **Create Alert** check box for each feed panel in which you want 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 reporting to your CB Response server. For more information, see [“Using the Head-Up Display Page”](#) on page 307.

UNRESOLVED ALERTS Go to Alerts >

Search...

Resolved False Positive In Progress Unresolved

SCORE	SOURCE	HOST	FEED	SIGNATURE	CAUSE	PATH	TIME
51	Test list	k8sp	My Watchlists	svchost.exe	svchost.exe	c:\windows\system32\svcho...	17 minutes ago
51	Test list	k8sp	My Watchlists	svchost.exe	svchost.exe	c:\windows\system32\svcho...	an hour ago
51	Test list	k8sp	My Watchlists	svchost.exe	svchost.exe	c:\windows\system32\svcho...	an hour ago
51	Test list	k8sp	My Watchlists	svchost.exe	svchost.exe	c:\windows\system32\svcho...	an hour ago
51	Test list	k8sp	My Watchlists	svchost.exe	svchost.exe	c:\windows\system32\svcho...	an hour ago
51	Wiain	k8sp	My Watchlists	services.exe	services.exe	c:\windows\system32\servic...	an hour ago
51	Wiain	k8sp	My Watchlists	lsass.exe	lsass.exe	c:\windows\system32\lsass...	an hour ago
51	Wiain	k8sp	My Watchlists	lsn.exe	lsn.exe	c:\windows\system32\lsn.exe	an hour ago

Showing 1 to 8 of 80 1 2 3 4 5

By default, the **Unresolved Alerts** panel of the **HUD** page displays all unresolved alerts for a sensor. You can also display resolved, false positive, and in-progress alerts by clicking on one of the following buttons 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 associated with the alert, such as threat intelligence and watchlist feeds. Clicking a link in this column opens the associated page, such as the Watchlist or Threat Intelligence page.
Host	Displays the host associated with the alert. Clicking a link in this column opens the Sensors page.
Feed	Displays the feed type associated with this alert.
Signature	If this column contains a Signed/Unsigned status, this indicates that the alert is associated with a binary. If this column is empty, this indicates that the alert is associated with a process.
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.
Path	Displays the path to the affected binary/process.
Time	Displays the time when the alert occurred.

The **Unresolved Alerts** panel also contains a **Go to Alerts** 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 on one or more of your endpoints, incident responders must:

- Determine the seriousness of the alert.
- Determine whether or not the alert indicates a sufficiently severe threat.
- Find a way to resolve a serious threat.

This might involve using CB Response features, such as:

- Endpoint Isolation
- Live Response
- Banning

It might also require using other tools.

Given the high volume of threat reports in the current environment, 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 the alert statuses.

To open the Triage Alerts page:

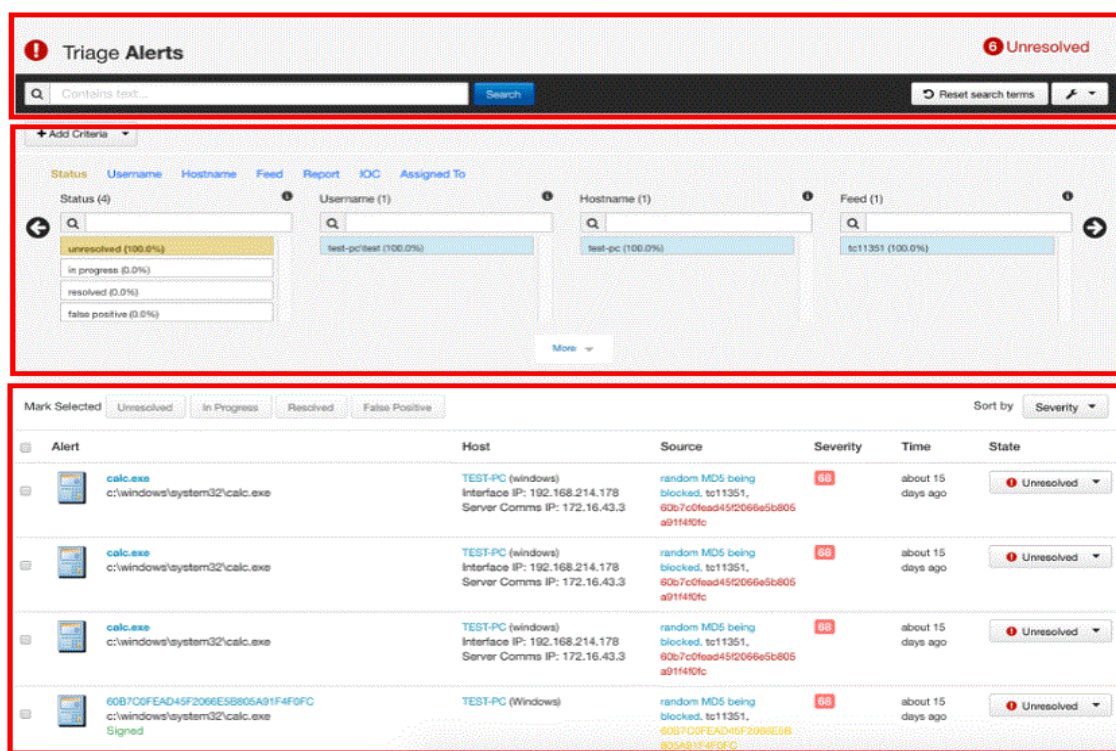
In the navigation bar, select **Triage Alerts**.

Note

You also can navigate to the Triage Alerts page from the **HUD** page by clicking **Go to Alerts** in the **Unresolved Alerts** panel. See [“Viewing Alert Activity on the HUD Page”](#) on page 295 for more details.

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 a series of filters that are category-specific lists (**Status**, **Username**, and so on). These filters show the percentage of alerts matching different values in each category and allow you to filter the view to show alerts matching one or more values.
- The bottom section contains the **Alerts** table, which contains details for alerts matching the search criteria entered in the first two sections.



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 or not the binary for a process has been signed
- Whether or not a binary for which an alert has been reported is considered “Trusted” by the CB Response Alliance

The **Alerts** table has several tools for adjusting the table display:

- **Alerts per page** – You can choose the number of alerts to show per page by changing the number in the **Show** box at the bottom-left corner of the **Alerts** table.
- **Sort order** – You can sort the **Alerts** table using the **Sort By** button in the top-right corner of the **Alerts** table:
 - **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.

Mark Selected Unresolved In Progress Resolved False Positive Sort by **Severity** ▼

<input type="checkbox"/>	Process	Host	Source	Severity	Time	State
<input type="checkbox"/>	chrome.exe c:\program files (x86)\google\chrome\application\chrome.exe	ALFONSE (windows) Interface IP: 192.168.174.135 Server Comms IP: 10.36.8.28	Carbon Black Process Blocking, cbbanning	61	about 6 days	<input type="checkbox"/> Unresolved ▼
<input type="checkbox"/>	chrome.exe c:\program files (x86)\google\chrome\application\chrome.exe	ALFONSE (windows) Interface IP: 192.168.174.135 Server Comms IP: 10.36.8.28	Carbon Black Process Blocking, cbbanning	61	about 6 days	<input type="checkbox"/> Unresolved ▼
<input type="checkbox"/>	chrome.exe c:\program files (x86)\google\chrome\application\chrome.exe	ALFONSE (windows) Interface IP: 192.168.174.135 Server Comms IP: 10.36.8.28	Carbon Black Process Blocking, cbbanning	61	about 6 days	<input type="checkbox"/> Unresolved ▼
<input type="checkbox"/>	svchost.exe c:\windows\system32\svchost.exe	ALFONSE (windows) Interface IP: 192.168.174.135 Server Comms IP: 10.36.8.28	Suspicious svchost user, bitSuspiciousindicators	54	about 6 days	<input type="checkbox"/> Unresolved ▼
<input type="checkbox"/>	anacron /etc/cron.hourly/anacron	dev5.mycorp.local (linux) Interface IP: 192.168.174.136 Server Comms IP: 10.36.8.28	unsigned, My Watchlists	51	about 1 month	<input type="checkbox"/> Unresolved ▼
<input type="checkbox"/>	on_ac_power /usr/bin/on_ac_power	dev5.mycorp.local (linux) Interface IP: 192.168.174.136 Server Comms IP: 10.36.8.28	unsigned, My Watchlists	51	about 1 month	<input type="checkbox"/> Unresolved ▼
<input type="checkbox"/>	bash /bin/bash	dev5.mycorp.local (linux) Interface IP: 192.168.174.136 Server Comms IP: 10.36.8.28	unsigned, My Watchlists	51	about 1 month	<input type="checkbox"/> Unresolved ▼
<input type="checkbox"/>	dbus-daemon-launch-helper /lib64/dbus-1/dbus-daemon-launch-helper	dev5.mycorp.local (linux) Interface IP: 192.168.174.136 Server Comms IP: 10.36.8.28	unsigned, My Watchlists	51	about 1 month	<input type="checkbox"/> Unresolved ▼
<input type="checkbox"/>	dbus-daemon-launch-helper /lib64/dbus-1/dbus-daemon-launch-helper	dev5.mycorp.local (linux) Interface IP: 192.168.174.136 Server Comms IP: 10.36.8.28	unsigned, My Watchlists	51	about 1 month	<input type="checkbox"/> Unresolved ▼
<input type="checkbox"/>	bash /bin/bash	dev5.mycorp.local (linux) Interface IP: 192.168.174.136 Server Comms IP: 10.36.8.28 dev5.mycorp.local	unsigned, My Watchlists	51	about 1 month	<input type="checkbox"/> Unresolved ▼

Show of 4,598 First ◀ 1 2 3 4 ... ▶▶

Alerts Table Data

The row for each alert in the table is divided into columns as follows:

Alert	Host	Source	Severity	Time	State
calc.exe c:\windows\system32\calc.exe	TEST-PC (windows) Interface IP: 192.168.214.178 Server Comms IP: 172.16.43.3	random MD5 being blocked, tc11351, 60b7c0fead45f206e5b805a91f4f0fc	99	about 15 days ago	Unresolved
calc.exe c:\windows\system32\calc.exe	TEST-PC (windows) Interface IP: 192.168.214.178 Server Comms IP: 172.16.43.3	random MD5 being blocked, tc11351, 60b7c0fead45f206e5b805a91f4f0fc	99	about 15 days ago	Unresolved
calc.exe c:\windows\system32\calc.exe	TEST-PC (windows) Interface IP: 192.168.214.178 Server Comms IP: 172.16.43.3	random MD5 being blocked, tc11351, 60b7c0fead45f206e5b805a91f4f0fc	99	about 15 days ago	Unresolved
60B7C0FEAD45F206E5B805A91F4F0FC Signed c:\windows\system32\calc.exe	TEST-PC (Windows)	random MD5 being blocked, tc11351, 60b7c0fead45f206e5b805a91f4f0fc	90	about 15 days ago	Unresolved

- **Alert** – The following details:
 - An icon representing 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, and the icon is of the host type.
 - The directory path where the process or binary is installed.
 - Whether or not the certificate for this process or binary is signed.
 - 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 with 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 produced by CB Response, based on underlying alert data. Clicking the **Severity** number in a particular row shows additional details by which the severity is calculated:
 - **Feed rating**
 - **Report scores**
 - **Confidence**
 - **Criticality**

The **Severity** numbers are color-coded as follows with red being severe threats and green being low threats. A score of 100 represents the most severe alerts.

Source	Severity	Time	State
Carbon Black Detected EMET Event, cbemet	61	about 3 months ago	Unresolved
Carbon Black Detected EMET Event, cbemet	61	about 3 months ago	Unresolved
Tamper detection of CB sensor-owned files, cbtamper	48	about 2 months ago	Unresolved
Tamper detection of CB sensor-owned files, cbtamper	48	about 2 months ago	Unresolved

- **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 whose resolutions are in progress.
 - **Mark as False Positive** – Select this for alerts that were not true threats.

You can select alerts using the checkbox to the left of the alert row and then change the **State** to any of the above selections. For more information, see [“Managing Alert Status”](#) on page 300.

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. In the navigation bar, select **Triage Alerts**.
2. On the **Triage Alerts** page, enter the search string and/or filter criteria for the alerts with statuses you want to change.
3. From the **Actions** menu (located in the top-right corner of the **Triage Alerts** page), select the appropriate **Mark all** menu option for the status you want to assign.

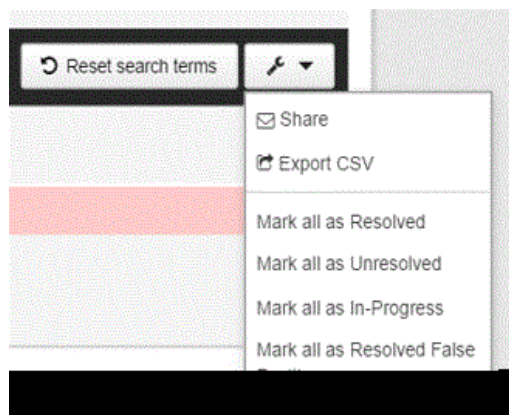
4. Click **OK** in the confirmation window to change the status of all of the alerts on the page.

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 other pages that are not displayed. Once you change the status, there is no “undo” command. Be especially careful about changing alert statuses when the view is unfiltered (i.e., showing all alerts).

To change the status of one alert:

1. In the navigation bar, select **Triage Alerts**.
2. In the **Alerts** table, select the check box to the left of the alert with 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

Keep in mind that alerts with statuses that you change 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 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.

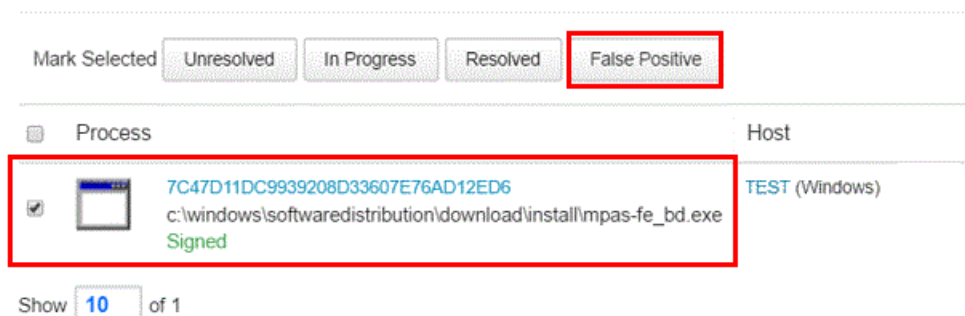
CB Response also provides a feature that allows you to ignore future instances of a false positive alert from a threat feed. You can choose to ignore an individual alert or specify that all alerts matching your search criteria be ignored in the future.

Note

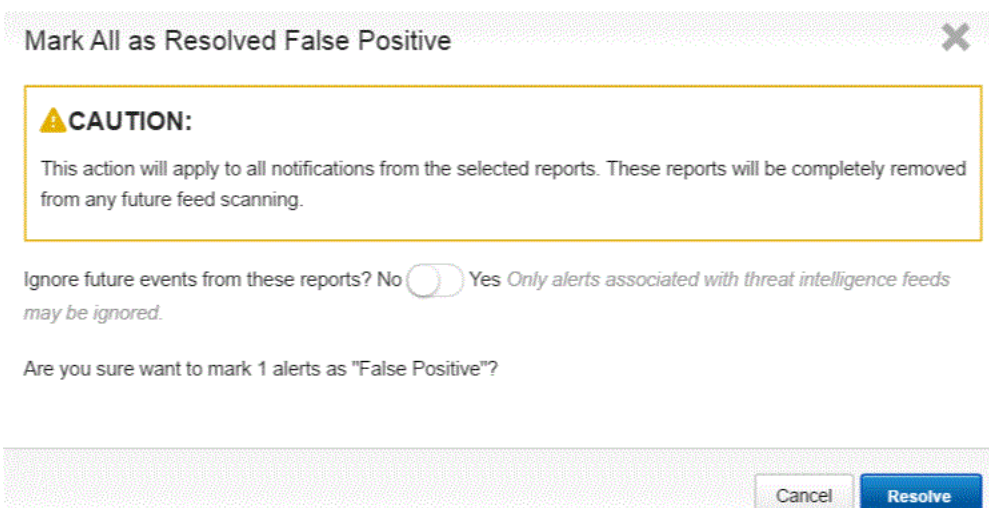
Only threat feed alerts can be designated as alerts to ignore. Alerts from watchlist matches are always triggered, since watchlists are assumed to use criteria specifically chosen by one or more of your CB Response users.

To ignore the triggering event for an alert:

1. In the navigation bar, select **Triage Alerts**.
2. In the **Alerts** table, select the check box to the left of the alert with the triggering event you want to ignore.
3. Click the **False Positive** button.



4. In the **Mark All as Resolved False Positive** window, you can ignore future events from this report by moving the slider button to **Yes**.



5. To resolve the alert and ignore future events from it, click the **Resolve** button.

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. If an event is significant enough, you can then go to the console to investigate and resolve it. 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 mail server.
- Use the CB Response External Mail Server.
- Opt out of email alerts.

If you use the CB Response External Mail Server, the following information is sent through the server and stored by CB Response:

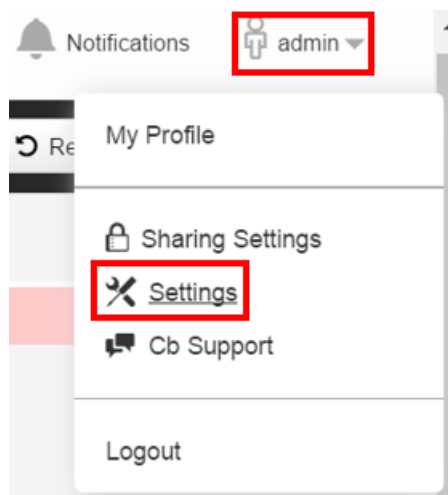
- Your server ID
- The time of the email
- The name of the watchlist or feed are that triggered the hit

Important

Carbon Black strongly recommends that you use your own email server, because email sent through the CB Response External Mail 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, select **username > Settings**.



3. On the **Settings** page, click **E-Mail** in the left panel to display the **Alerting via Email** page.

 A screenshot of the 'Settings' page. The left sidebar has 'E-Mail' selected. The main content area is titled 'Alerting via Email'. It has two radio buttons: 'Use Carbon Black External Mail Server (Secure HTTPS POST to api.alliance.carbonblack.com)' which is selected, and 'Use My Own Mail Server'. Below these are input fields for 'SMTP Server', 'Port', 'Username (Email Address)', and 'Password'. There is a 'Connection Type' section with radio buttons for 'Secure Connection using TLS' (selected), 'Secure Connection using SSL', 'Plaintext Connection (Insecure)', and 'I do not want to receive email alerts from Carbon Black'. A 'Save Changes' button is at the bottom right.

4. Select the **Use My Own Mail Server** radio button.
5. Provide the following information for the mail server you want to use:
 - a. **SMTP Server** – The address of the SMTP server you will use.
 - b. **Port** – The port for email service.
 - c. **Username** – The login account required to Log into the server.
 - d. **Password** – The password needed to login as the specified user.
 - e. **Connection Type** – The security protocol to use for this connection (TLS or SSL) or Plaintext Connection if you do not want the email to be secure.
6. When you have finished entering server configuration settings, click the **Save Changes** button.

All email alerts for all console users should now be routed through this server.

Enabling Specific Email Alerts

Once you have an email server configured, any watchlist or feed in the CB Response console can be configured to send email alerts when it gets a hit on a CB Response sensor.

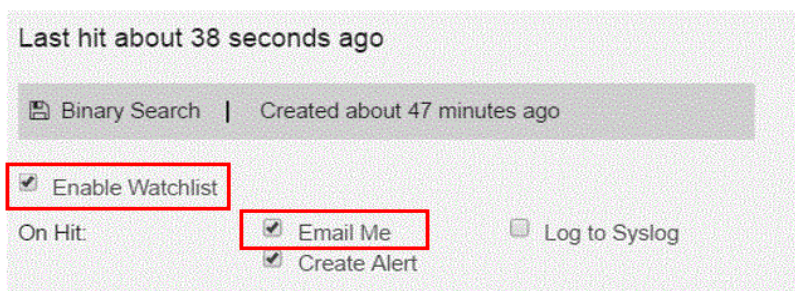
You can turn on/off email alerts for individual watchlists and feeds as needed, for example, if you find that a watchlist or feed is creating too much email traffic. Keep in mind that email alerts for any specific watchlist or feed are enabled on a per-user basis.

Note

If you have upgraded from a previous release of CB Response, any email alerts you had enabled for watchlists and threat intelligence feeds will remain enabled after the upgrade.

To enable email alerts for a watchlist:

1. In the navigation bar, select **Watchlists**.
2. In the left menu, select the watchlist for which you want 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. With the watchlist details showing, if you want to begin receiving alerts immediately, be sure the **Enable Watchlist** box is checked.
4. Check the **Email Me** check box to enable email alerts for this watchlist for the current console user.



To enable email alerts for a threat intelligence feed:

1. In the navigation bar, select **Threat Intelligence**.
2. For each feed for which you want to activate email alerts, check the **Email Me on Hit** check box. Email alerts should now be enabled for each feed.

National Vulnerability Database
nvd.nist.gov

NVD is the U.S. government repository of standards based vulnerability management data. This feed will flag executed applications vulnerable to one or more CVEs with CVSS scores higher than 7.0 from 2013-2015 for Java, Flash Player and Google
[More Info »](#)

☆☆☆☆

Enabled
 Email Me On Hit
▶ Notifications

[Process Matches »](#) [Actions ▼](#)
[Binary Matches »](#)
[Threat Reports »](#)

Chapter 18

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	308
Endpoint Hygiene Panel	309
Event Monitor Panel	309
Query Duration Panel	310
Resolution Time Panel	311
Saved Searches Panel	311
Sensors Panel	311
Unresolved Alerts Panel	313

Overview of HUD

The **HUD** 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 of the page.

Customizing the HUD Page

To reposition HUD panels:

- Hold down your left mouse button on a panel and drag the panel to the desired position.

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 the panels that are contained on the **HUD** page. However, the **Unresolved Alerts** panel is described in [“Viewing Alert Activity on the HUD Page”](#) on page 295.

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.

QUERY DURATION

Filter by query source

All UI Watchlist Feed Report API

DURATION	SOURCE	USERNAME	QUERY	
1 m	API	alice	alliance_score_abusech=[65 TO *]	Copy
59 s	API	bob	alliance_score_abusech=[65 TO *]	Copy
53 s	UI	alice	*.vbs	Copy
35 s	UI	bob	filemod*.pdf	Copy
32 s	UI	alice	*.vbs	Copy

Showing 1 to 5 of 19 [1](#) [2](#) [3](#) [4](#) [5](#)

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** – These are automated queries. Watchlist queries are created by CB Response users and run every 10 minutes.
- **Feed Report** – These are 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. It includes links to the **Process Search** page, where you can view saved search results. See "[Saved Searches](#)" on page 180.



Sensors Panel

Use the **Sensors** panel to monitor and manage sensor hosts. See "[Managing Sensors](#)" on page 92.

SENSORS [View all >](#)

Default Group ▼

Show Uninstalled Sensors

<input type="checkbox"/>	HEALTH	▲ HOST	STATUS	HEALTH MESS...	ACTIVITY	SENSOR VERS...
<input type="checkbox"/>	100		Online	Healthy	expected in 1...	6.1.10.10169

Showing 1

By default, uninstalled sensors do not display in this panel. Select the checkbox to show uninstalled sensors.

Also 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 83.
- Isolate; see [“Isolate Host”](#) on page 192.
- 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 318.
Health Message	Displays a health message that relates to the sensor health score. See “Sensor Health Score Messages” on page 318.
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 100.
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.

Unresolved Alerts		View all »	by: TIME SEVERITY		
	B4A134CCF183F4D5FCF9C6A7D6F38203	 cleopolds-Mac.I...	My Watchlists	18 days ago	
	676782354B29658528680A8C05BDE947	 cleopolds-Mac.I...	My Watchlists	18 days ago	
	0ACC295682B376D5D05158E4025DBBEE	 cleopolds-Mac.I...	My Watchlists	18 days ago	
	D167D97B1832A806F153EB23956F738B	 cleopolds-Mac.I...	My Watchlists	18 days ago	

See [“Viewing Alert Activity on the HUD Page”](#) on page 295.

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	315
Sensor Group Feature Support	317

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*
Global VDI Support	Yes	Yes	Yes
Hash Banning	Yes	Yes	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	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**	No	Yes**

Feature	Windows	Linux	OSX
SHA256 hashes in events (Collection)	Yes***	No	Yes***
Support for FIPS	Yes	No	No
Tamper Detection	Yes	No	No
Tamper Protection	No	No	No

Notes

*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.

**TLS cert swapping support is for sensor versions Windows 6.2.3-win and Mac 6.2.5-osx and above.

***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	319
macOS Health Events	323
Linux Health Events	326

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.