

**Carbon Black.**

A graphic element consisting of a dark blue hexagon with a red hexagon inside it, positioned in the center of the page. The background features a dark blue grid of hexagons and several light blue lines forming a geometric pattern.

# **CB Response Server/Cluster Management Guide**

**Server Version: 7.0**

**Document Date: January 2020**

# Copyrights and Notices

Copyright ©2011-2020 VMware, Inc. All rights reserved. Carbon Black is a registered trademark and/or trademark of VMware, Inc. in the United States and other countries. All other trademarks and product names may be the trademarks of their respective owners.

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

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

We acknowledge the use of the following third-party software in the CB Response software product:

- Antlr python runtime - Copyright (c) 2010 Terence Parr
- Backbone routefilter - Copyright (c) 2012 Boaz Sender
- Backbone Upload - Copyright (c) 2014 Joe Vu, Homeslice Solutions
- Backbone Validation - Copyright (c) 2014 Thomas Pedersen, <http://thedersen.com>
- Backbone.js - Copyright (c) 2010–2014 Jeremy Ashkenas, DocumentCloud
- Beautifulsoup - Copyright (c) 2004–2015 Leonard Richardson
- Canvas2Image - Copyright (c) 2011 Tommy-Carlos Williams (<http://github.com/devgeeks>)
- Code Mirror - Copyright (c) 2014 by Marijn Haverbeke [marijnh@gmail.com](mailto: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](mailto:fgnass@neteye dot de)]
- JQuery tablesorter - Copyright (c) Christian Bach.
- JQuery timepicker - Copyright (c) Jon Thornton, [thornton.jon@gmail.com](mailto: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](mailto: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](mailto:benoitc@e-engura.org) and Copyright 2009–2013 (c) Paul J. Davis [paul.joseph.davis@gmail.com](mailto: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](mailto: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.

**VMware Carbon Black**

1100 Winter Street, Waltham, MA 02451 USA

Tel: 617.393.7400

Fax: 617.393.7499

Email: [support@carbonblack.com](mailto:support@carbonblack.com)

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

*CB Response Server / Cluster Management Guide*

Product Version: 7.0

Document Revision Date: November 14, 2019

# Before You Begin

This preface provides a brief orientation to the *CB Response 7.0 Server/Cluster Management Guide*.

## Sections

Topic	Page
<a href="#">What this Document Covers</a>	6
<a href="#">Other Documentation</a>	6
<a href="#">Community Resources</a>	7
<a href="#">Contacting Support</a>	7

## What this Document Covers

This document explains how to manage CB Response servers and clusters. The following table summarizes the contents of this guide:

Chapter	Description
1 <a href="#">Server Overview</a>	Provides an overview of the CB Response server technology stack, daemons, configuration, and logs.
2 <a href="#">Installing the CB Response Server</a>	Explains how to install/initialize a new CB Response server, as well as how to upgrade, troubleshoot, and uninstall the server.
3 <a href="#">Server Backup and Restoration</a>	Explains how to perform various backup and restore procedures.
4 <a href="#">Ports and Protocols</a>	Provides a collection of tables that detail port and protocol information for several different server communications
5 <a href="#">Installing a CB Response Cluster</a>	Introduces CB Response clusters and explains how to configure clusters, add minions to existing clusters, remove minion nodes from clusters, and upgrade cluster nodes
6 <a href="#">Using CBCLUSTER as a Non-Root User</a>	Describes how to use the CBCLUSTER command as a non-root user.

## 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. This was called the *Server Sizing Guide* in previous releases.
- *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* – (this document) 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* – Describes the CB Response product and explains how to use all of its features and perform administration tasks. Beginning with CB Response Server 6.5, this User Guide is also available online through the **Help** menu in the console.

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

## Community Resources

The Carbon Black User Exchange website at <https://community.carbonblack.com> provides access to information shared by Carbon Black customers, employees and partners. It includes information and community participation for users of all 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](mailto:support@carbonblack.com)

**Phone:** 877.248.9098

**Fax:** 617.393.7499

## Reporting Problems

When you call or email technical support, provide the following information to the support representative:

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



# Contents

<b>Copyrights and Notices</b> .....	2
<b>Before You Begin</b> .....	5
What this Document Covers .....	6
Other Documentation .....	6
Community Resources .....	7
Contacting Support .....	7
Reporting Problems .....	8
<b>1 Server Overview</b> .....	12
Server Overview .....	13
Server Configuration .....	15
Server Logs .....	16
Log Overview .....	16
Troubleshooting .....	16
Error in the CB Response Console Interface .....	16
Sensors Are Not Checking in .....	16
Ensure that Everything is Working .....	17
<b>2 Installing the CB Response Server</b> .....	18
Overview .....	19
Firewall and Connectivity Requirements .....	19
Substituting a Server Communication Certificate .....	20
Installing and Initializing a CB Response Server .....	22
Upgrading a CB Response Server .....	31
Using cbupgrade .....	32
Server Upgrades and New Sensor Versions .....	32
Upgrade Policy Settings .....	32
Supporting Multiple Volumes for Event Data .....	33
Naming Conventions .....	33
Using New Data Directories .....	34
Partitioning .....	34
Active and Read-Only Directories .....	34
Partition Purging .....	34
Extending Disk Space on the Fly .....	35
Troubleshooting the Server .....	35
Uninstalling a CB Response Server .....	37
Removing the CB Response Server .....	38
<b>3 Server Backup and Restoration</b> .....	40
Overview .....	41
Restoration Servers .....	41
Backup .....	42
Configuration Backup .....	42
Data Backup .....	44
Restore .....	45

Failed Minion Cluster Restore .....	46
Configuration Restore .....	47
Data Restore .....	49
<b>4 Ports and Protocols .....</b>	<b>51</b>
<b>5 Installing a CB Response Cluster .....</b>	<b>54</b>
Overview .....	55
Cluster Architecture .....	55
Cluster Operation .....	57
Configuring CB Response Clusters .....	58
Sensor Install and Verification .....	60
Best Practices .....	61
Adding Minions to Existing Clusters .....	61
Removing Minions from Existing Clusters .....	62
Best Practices .....	62
Read-Only Minions .....	63
Removing Minions .....	63
Upgrading Cluster Nodes .....	64
Manually Upgrading Cluster Nodes .....	64
<b>6 Using CBCLUSTER as a Non-Root User .....</b>	<b>66</b>
Overview .....	67
Required User Privileges .....	67
Defining Users .....	69

## List of Tasks

### How to . . .

This section lists procedures documented as stepwise tasks in this guide. Some single-step tasks and tasks that are more easily explained in the context of a paragraph are not included here. If you do not see the task you are looking for, see the main [Contents](#) pages.

To add minions to existing clusters: .....	61
To install and initialize a new server: .....	22
To install sensors and verify they appear in CB Response: .....	61
To install the server and restore data on failed minions of an otherwise operational cluster: .....	46
To manually upgrade minion servers in a cluster: .....	65
To mark minions as read-only: .....	63
To perform a configuration backup: .....	42
To perform a configuration restore: .....	47
To perform a data backup: .....	44
To perform a data restore: .....	50
To remove a cluster minion: .....	63
To set up the cluster configuration: .....	59
To uninstall a server (CB Response and related RPMs): .....	38
To uninstall a server (CB Response RPMs only): .....	38
To unmark minions as read-only: .....	63
To upgrade a clustered server: .....	31
To upgrade a standalone server: .....	31
To upgrade cluster nodes: .....	64
To use a symlinked location for event storage: .....	33

## Chapter 1

# Server Overview

This chapter provides an overview of the CB Response server technology stack, daemons, configuration, and logs.

### Sections

Topic	Page
<a href="#">Server Overview</a>	13
<a href="#">Server Configuration</a>	15
<a href="#">Server Logs</a>	16

## Server Overview

This section describes the technology stack on a CB Response server. Five major daemons exist in CB Response server, as described in the following table:

Daemon	Description
cb-nginx	Used as an HTTP reverse proxy to internal daemons.
cb-coreservices	(Python, Gunicorn) All non-data application logic for HTTP transactions.
cb-datastore	(Java/Jetty) All incoming data, including event logs and binary files.
cb-solr	(Java/Jetty) Apache Solr, the primary data store.
cb-postgres	Traditional relational database.
cb-sensorservices	Handles all non-data sensor requests, such as sensor check-ins, registrations, and upgrades.

nginx is the only daemon with public sockets. The remaining daemons are bound to the CB Response server using the default IP address, which is `127.0.0.1`, and can only be accessed locally or by using the nginx reverse proxy.

nginx owns `tcp/80` and `tcp/443` and redirects to `coreservices`, `cb-datastore` or `cb-sensorservices` to the CB Response web `root:q` based on the URL prefix, as described in the following table:

nginx	Redirects to
/	/var/www/cb/
/api/*	coreservices on tcp/5000
/sensor/*	sensorservices on tcp/6500 and 6501
/data/*	cb-datastore on tcp/9000

### Note

**coreservices** handles `/api/*`

All `/api/` URLs are used by the CB Response console interface and by REST clients.

**sensorservices** handles `/sensor/*`

All `/sensor/` URLs are used by the sensors that are pushing data. These URLs are isolated to allow binding a separate nginx server instance to `tcp/443` on a public / DMZ interface for sensors that are outside of the internal network (for example, sensors on laptops used by traveling employees, and work laptops that are at employees' homes) without exposing the `/api/` interfaces externally. You can isolate these URLs by using a simple nginx configuration change, as shown in the example in the file:

```
/etc/cb/nginx/conf.d/cb-multihome.conf.example
```

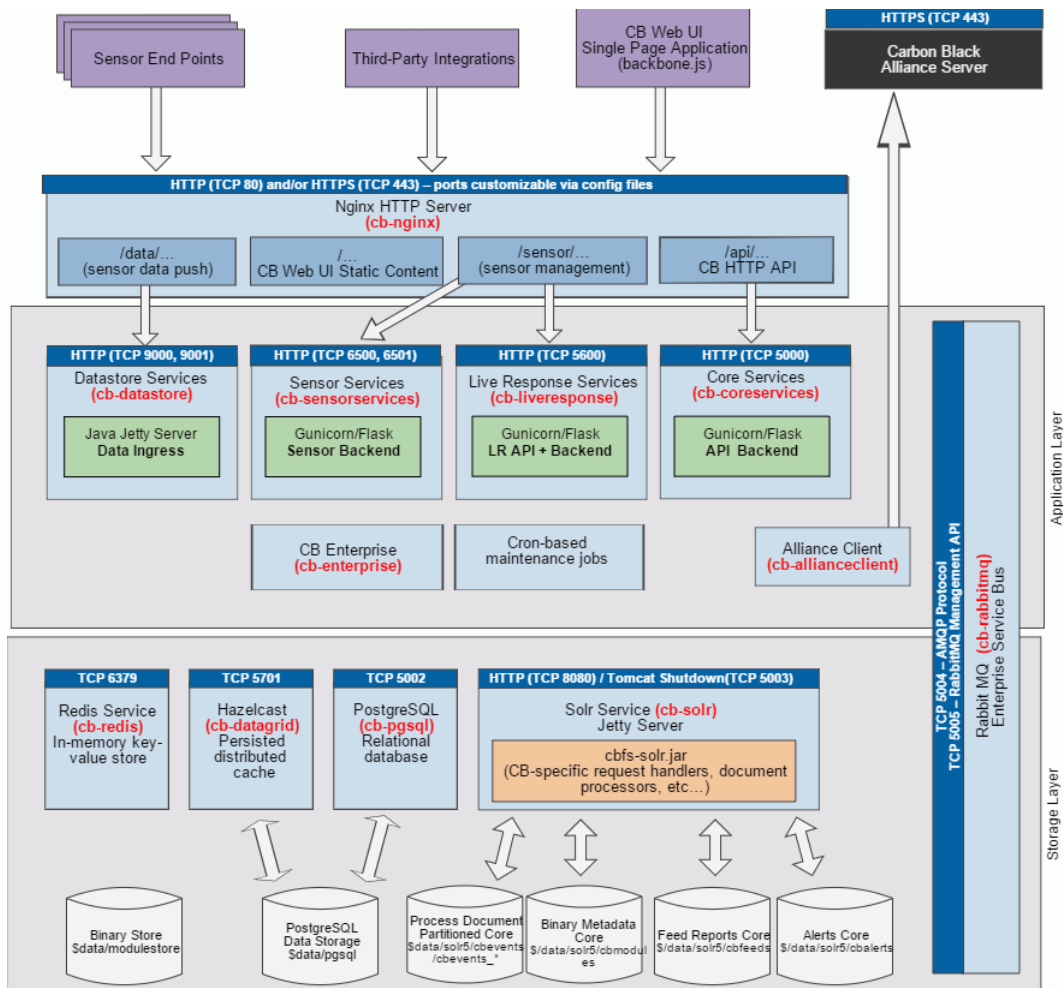
**Note**

Listening ports are configured differently in a clustered setup. See cluster-specific documentation for more details.

In general, sensors first register and check into `sensorservices` by using `nginx`. If sensors have data, after they check in, they post event logs to `cb-datastore` by using `nginx`.

`cb-datastore` caches data for a few minutes before sending a collection of related data to `cb-solr`.

The following diagram illustrates the CB Response server architecture at a high level:



## Server Configuration

Daemon configuration data is generally static, and is stored in flat files that follow typical Linux conventions. Dynamic run-time configuration data is stored in PostgreSQL and configured by using the CB Response console.

Most major configuration is done after the CB Response server is installed, when the `cbinit` script is run. `cbinit` configures a combination of initial settings in both static configuration files and PostgreSQL. For more information about installing and configuring the CB Response server, see [Chapter 2, “Installing the CB Response Server.”](#)

The following table describes the major static configuration files.

Static Configuration File	Description
<code>/etc/cb/cb.conf</code>	Major CB Response enterprise-wide settings.
<code>/etc/cb/solr6/core_conf</code>	This directory stores <a href="#">Apache Solr settings</a> . Each subdirectory has individual solr core configurations.
<code>/var/cb/data/pgsql/postgresql.conf</code>	<a href="#">PostgreSQL settings</a> .
<code>/etc/cb/coreservices-logger.conf</code>	Configuration of log settings for the Python <code>coreservices</code> daemon.
<code>/etc/cb/sensorservices-logger.conf</code>	Configuration of log settings for the Python <code>sensorservices</code> daemon.
<code>/etc/cb/cb-datastore/*</code>	Configuration of log settings for the Java <code>cb-datastore</code> daemon.
<code>/etc/cb/nginx/conf.d/cb.conf</code>	<a href="#">nginx server settings</a> for defining IP addresses and ports to bind to listeners.
<code>/etc/sysconfig/iptables</code>	Standard <a href="#">iptables configuration</a> for the server firewall (for CentOS 6).

The following table describes the secondary static configuration files:

Static Configuration File	Description
<code>/etc/cb/allianceclient-logger.conf</code>	Configuration of log settings for the CB Response Alliance client daemon.
<code>/etc/cb/nginx/cb-nginx.conf</code>	Standard <a href="#">nginx server configurations</a> that are specific to CB Response.
<code>/etc/cb/solr6/solr.in.sh</code>	Jetty servlet settings for Solr

# Server Logs

## Log Overview

The CB Response server uses log files extensively. These files provide both reassurance of system health and a record of activity under error conditions.

The following table describes critical logs that are on the CB Response server:

Log	Description
<code>/var/log/cb/nginx/access.log (and error.log)</code>	nginx HTTP access and error logs for all sensor and API traffic.
<code>/var/log/cb/coreservices/debug.log</code>	Application logic for API traffic.
<code>/var/log/cb/sensorservices/debug.log</code>	Application logic for sensor traffic.
<code>/var/log/cb/datastore/debug.log</code>	Incoming sensor data cache.
<code>/var/log/cb/solr/debug.log</code>	Sensor data storage, indexing, and queries.

## Troubleshooting

### Error in the CB Response Console Interface

Look in the log file: `/var/log/cb/coreservices/debug.log` for a Python stack trace with details. Contact your Carbon Black Technical Support representative for guidance.

### Sensors Are Not Checking in

If you find that a sensor is not checking in, look in the log file: `/var/log/cb/nginx/access.log` for a request from the host in question. For example:

```
164.230.214.13 - - [20/Apr/2017:20:04:52 +0000 (3.811)] "POST /
sensor/checkin/35998 HTTP/1.1" 502 166 "-" "sensors.vibrant-
pies.my.carbonblack.io" ">170.16.20.21:6501" "-" "-"
```

In the output example, there are several fields highlighted in red, and these are useful for diagnosing the issue:

- Sensor Id is reported after the checkin field. It is **35998** in the above case.
- The actual error code is reported in the field following HTTP/1.1 text. In the example above, **502** is the error code.
- In a clustered environment, some requests proxy calls to a different minion in the cluster. This minion's address is reported after the '>' character (**170.16.20.21** in the above case).

If you do find a checkin error, or an error to any other call with the `/sensor/` prefix, check the following log for more information about the error:

```
/var/log/cb/sensorservices/debug.log
```



If you see an error related to requests prefixed with `"/data/"`, check the following log:

```
/var/log/cb/datastore/debug.log
```

In a cluster environment, you will need to look at log files on the node referenced in the nginx error log entry.

### Alternatives

If sensors are not checking in but there are no entries in `access.log`, check `error.log`.

If the sensor SSL client certificates are not signed by the Certificate Authority (CA) in `/etc/cb/certs` and configured in `/etc/cb/nginx/conf.d/cb.conf`, nginx will refuse the request.

If there are no entries in `error.log`, check the status of sensor communications as described in the “Troubleshooting Sensors” section in the *CB Response User Guide*.

### Ensure that Everything is Working

Check the nginx `access.log` for ‘200’ HTTP response codes. 200 codes indicate that communications are working normally.

## Chapter 2

# Installing the CB Response Server

This chapter explains how to install/initialize a new CB Response server, as well as how to upgrade, troubleshoot, and uninstall the server.

### Sections

Topic	Page
<a href="#">Overview</a>	19
<a href="#">Installing and Initializing a CB Response Server</a>	22
<a href="#">Upgrading a CB Response Server</a>	31
<a href="#">Supporting Multiple Volumes for Event Data</a>	33
<a href="#">Troubleshooting the Server</a>	35
<a href="#">Uninstalling a CB Response Server</a>	37

## Overview

This chapter describes the steps for installing the CB Response server. It covers new installations and server upgrades. You can complete the entire process in about ten minutes, assuming reasonable download speed.

The separate *CB Response Server Operating Environment Requirements* guide provides guidelines for hardware and software required for the CB Response server. Your environment must meet these requirements before you begin installation. See the [Carbon Black User Exchange](#) to locate this guide.

A CB Response server installation consists of these main steps:

1. Obtain and install an RPM from Carbon Black. This RPM does not install the CB Response server. It sets up a Yum repository and installs an SSL client certificate that allows the full CB Response server to be downloaded and installed.
2. Install the CB Response server. This is a two-step process that involves running the `yum install` command and the `cbinit` configuration script. The CB Response server is downloaded when you run the `yum install` command.

For more information on `cbinit`, see *Automating cbinit* on the [Carbon Black User Exchange](#).

When you have installed the server, you can then install sensors on the endpoints you intend to monitor. Instructions for installing and upgrading sensors are provided in the “Manage Sensors” chapter in the *CB Response User Guide*.

## Firewall and Connectivity Requirements

Internet connectivity through outbound TCP ports is required on the CB Response Server system for the scenarios described in the following table.

Scenario	Description	Address
CB Response Yum Repository	The RPM installer sets up a Yum repository.	yum.carbonblack.com:443 yum.distro.carbonblack.io:443
CB Response Alliance Server and CB Threat Intel	The Alliance Server and CB Threat Intel provide threat intelligence and can enable further analysis of files on endpoints through CB Threat Intel partners.  To view all threat intelligence data, you need both addresses.	api.alliance.carbonblack.com:443 threatintel.bit9.com:443
CentOS Yum Repository	The standard CentOS Yum repository server used during CB Response server installation to download standard packages	mirror.centos.org:80

## Substituting a Server Communication Certificate

CB Response uses HTTPS and TLS to secure communication between endpoints and the server, and to assure that the endpoint communicates only with the Response server it trusts and the server communicates only with trusted endpoints. 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, the certificate installed during `cbinit` is produced by the server itself. As an alternative to the default legacy certificate, you can substitute your own certificate during the server post-installation initialization process. In either case, the certificate will be named “Legacy” where certificates appear in the console, and it will be protected from deletion.

A server certificate used for sensor communications must meet the following requirements:

- The files you provide must be a valid certificate and key pair (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.
- The certificate 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 certificate, even in standalone configurations, so that the certificate remains 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 a new certificate 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.

The example below shows how you could set up the SAN portion of the certificate. 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
```

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, server-provided 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.

The server certificate and key are copied into the server as `/etc/cb/certs/cb-server.crt` and `/etc/cb/certs/cb-server.key`, and are also stored in the CB Response database.

**Important**

There are other certificate management features available in CB Response, including support for adding multiple certificates after server initialization. Please see “Managing Certificates for Server-Sensor Communication” in the *CB Response User Guide* for complete details.

## Installing and Initializing a CB Response Server

This section describes the steps for installing and initializing a **new** CB Response server. Root-level permissions are required throughout the entire installation/configuration process. You will use `su` or `sudo` to enter the installation/initialize commands.

### Caution

The steps in this section are for a new installation only. If you already have the CB Response Server installed, **do not perform these steps**. Instead, see [“Server Upgrades and New Sensor Versions”](#) on page 32.

**Using the new installation procedure on an existing server will likely result in loss of all data**, including the configuration and event data collected from sensors.

### To install and initialize a new server:

1. Verify that the host machine on which you intend to install CB Response server meets the hardware and software requirements specified in the *CB Response Server Operating Environment Requirements* guide you received from your Carbon Black representative.
2. Verify that the server has Internet connectivity as specified in [“Firewall and Connectivity Requirements”](#) on page 19.
3. Contact Carbon Black Technical Support to procure an installation RPM for the CB Response server.
4. Install the RPM:
  - a. Run the following command using the customer-specific RPM you received:

```
sudo rpm -ivh carbon-black-release-1.0.3-1-  
<customername>.x86_64.rpm
```
  - b. (Optional) Verify that the CB Response [cb] Yum repository was configured correctly. You can run this command to see the contents of the new Yum repository entry for CB Response:

```
cat /etc/yum.repos.d/CarbonBlack.repo.
```

```
[root@cb-enterprise-testing ~]# cat /etc/yum.repos.d/  
CarbonBlack.repo  
  
[CarbonBlack]  
name=CarbonBlack  
baseurl=https://yum.distro.carbonblack.io/enterprise/stable/  
$releasever/$basearch/  
gpgcheck=1  
enabled=1  
metadata_expire=60  
sslverify=1  
sslclientcert=/etc/cb/certs/carbonblack-alliance-client.crt  
sslclientkey=/etc/cb/certs/carbonblack-alliance-client.key
```

- c. (Optional) You should see the CB Response SSL certificates and keys in the following directory:

```
/etc/cb/certs/
```

5. Install the CB Response server:
  - a. Verify that your computer's date and time settings are accurate. Incorrect date/time settings can cause failures in SSL negotiation, which is required for Yum downloads.
  - b. Run the following command:
 

```
sudo yum install cb-enterprise
```

```
[bsmith@localhost yum.repos.d]$ sudo yum install cb-enterprise
```

- c. Install the CentOS GPG key if you are prompted to do so.
  - d. If your environment requires that outbound firewall exceptions be made, ensure that the exceptions documented in ["Firewall and Connectivity Requirements"](#) on page 19 are followed. You must also update `/etc/yum.repos.d/CentOS-Base.repo` to enable the baseurl of <http://mirror.centos.org>.
 

**Note:** Yum supports the use of web proxies. However, Carbon Black is not aware of a way to use Yum with NTLM-authenticated web proxies.
6. When the installation completes, initialize and configure the CB Response server.
  - a. Run the following command -- If you plan to use the server-provided certificate to secure communications with sensors, no arguments are necessary:
 

```
sudo /usr/share/cb/cbinit
```

To substitute your own certificate, add the following arguments to the `cbinit` command, substituting the full path to a certificate file and a key file where shown:

```
--server-cert-file=<certpath> --server-cert-key=<keypath>
```

**Important:** See ["Substituting a Server Communication Certificate"](#) on page 20 for certificate requirements.
  - b. Press **[Return]** to open the EULA. When you are done reviewing it, if you agree to the terms, enter `q` and then enter `yes` to continue the configuration.

```
-----
END USER LICENSE AGREEMENT
-----
Please, review and accept the End User License Agreement before proceeding
with the server setup

Hit 'return' to open the agreement and 'q' when you're done reading it:

Do you accept the license agreement [yes/no]: yes
```

- c. Select a storage location for your data and press **[Return]**.

```
-----
STORAGE LOCATION
-----
Please choose a data storage location with as much space as possible. If needed,
refer to the Carbon Black Data Storage Guidelines document.

Enter path for data storage location [/var/cb/data]:

You picked: /var/cb/data
```

**Note:** Per the *CB Response Operating Environment Requirements* guide, the primary datastore is mapped by default to `/var/cb/data`. If you configured your storage differently, review your current file system mapping (`df -h`) with Carbon Black Support or Professional Services. Incorrect or insufficient disk configurations will prevent CB Response from operating correctly.

- d. Enter an initial Administrator account to log in and configure CB Response. Enter values for **Username**, **First Name**, **Last Name**, **E-Mail**, **Password**, and **Confirm password**:

```

-----
ADMINISTRATOR ACCOUNT
-----
Here you configure your GLOBAL ADMINISTRATOR account.
This account is the most powerful account on the server.

Be sure to put a valid e-mail address if you want to take full advantage
of Carbon Black's notification system.

      Username: cbadmin
      First Name: CB
      Last Name: ADMIN
      E-Mail: cbadmin@carbonblack.com
      Password:
      Confirm password:

Verify Account Information:
      Username: cbadmin
      First Name: CB
      Last Name: ADMIN
      E-Mail: cbadmin@carbonblack.com

Is this correct [Y/n]: Y

```

- e. Press **[Enter]** and then validate the account information by entering **y**.
- f. In the **Sensor Communications** section, you define the address that the sensors will use to communicate back to the CB Response server:
- ```

Would you like to keep the default [Y/n]: n
Use SSL [Y/n]: Y
Hostname [192.168.117.141]: cbr.company.com
Port [443]: return

```

If the Verify Account Information looks correct, **Y**

**Note:** The IP address of the server will be accessed via the default SSL port 433. A best practice is to use a DNS record that points to this IP address.

Work with Carbon Black Support or Professional Services to make sure you understand the external connectivity options supported by CB Response Server

```

-----
SENSOR COMMUNICATIONS
-----
You need to configure the address that the sensors will talk to. This needs
to be an ip-address or domain name that is reachable by the sensor machines.
This can be different per sensor-group and can be changed later, but it is
easiest if you put in the valid address now.

Default sensor group server URL: https://192.168.117.141:443

Would you like to keep the default [Y/n]: n
Use SSL [Y/n]: Y
Hostname [192.168.117.141]: cbr.company.com
Port [443]:

New default sensor group server URL: https://cbr.company.com:443

Is this correct [Y/n]: Y

```



- g. Review all prompts and configure sharing settings in accordance to your company's security policies. The recommended settings are provided here. You can change these settings at any time by accessing the CB Response console and selecting your **username** > **Sharing Settings** in the top-right corner.
- Do you want to enable communication with the Carbon Black Alliance? - Y  
This enables the program to be supplemented with updated threat intelligence from CB Threat Intel and the extended network of CB Threat Intel partners.
  - Do you want your server to submit statistics and feedback information to Carbon Black? - Y  
This enables the server to submit health statistics back to CB Response. These are used by Carbon Black Support and Professional Services to determine how the allocated server is performing with our application.
  - Do you want the default sensor group to submit hashes to Carbon Black Alliance? - N  
See the "Threat Intelligence Feeds" chapter in the *CB Response User Guide* for more information on sharing hashes with CB Response.

- Continue with current sharing settings? - Y

```
-----
CARBON BLACK ALLIANCE
-----

Do you want to enable communication with the Carbon Black Alliance and Threat
Intelligence Servers?

This option controls the main switch which allows the Carbon Black Enterprise
Response Server to connect out and establish communications with Carbon Black
cloud infrastructure. This is required for a number of features which include
downloading threat intelligence feeds and reporting diagnostics data.
Enabling this switch does not enable any data transmission; each data stream
is controlled by separate, individual settings.

Your server must be able to communicate to
  https://api.alliance.carbonblack.com:443
  https://threatintel.bit9.com:443

Do you want to enable communication with the Carbon Black Alliance? [Y/n]: Y
-----
HELP IMPROVE YOUR CARBON BLACK EXPERIENCE
-----

We are constantly looking for ways to make the Carbon Black user experience
better. Please help us achieve this goal by allowing automatic reporting of
usage, resource, and sensor statistics to our technology and support teams.

You can later change your mind, too, by going here:

  >>> Administration -> Sharing Settings

Do you want your server to submit statistics and feedback information back to Carbon Black? [Y/n]:

Be notified of any binary that could be a potential threat. Information such as
the filename, MD5 hash and parent process will be shared with the Carbon Black
Alliance partners, including VirusTotal.

All information is anonymized to the extent reasonably practicable before being
shared with Carbon Black Alliance partners. The applicable terms and conditions
are set forth in and subject to your Carbon Black License Agreement. For further
information on what information is collected and shared by the Carbon Black
Alliance Server, please
visit https://www.carbonblack.com/solutions/carbon-black/collaboration/.

You can change this setting at any time in the server web console:

  >>> Administration -> Sharing Settings

If you enable this, you will then be prompted to either enable or disable the
uploading of unknown binaries.

Do you want the default sensor group to submit hashes to Carbon Black Alliance? [Y/n]: N

You have chosen to share data with Carbon Black Alliance
Please review your choices to make sure they are correct

Continue with current sharing settings? [Y/n]: Y
```

- h. The **SSL Certificates** section is automated and requires no user input. If you used arguments pointing to valid certificate and key files when you ran `cbinit`, the certificate from your organization is substituted for the default certificate created by the server. See [“Substituting a Server Communication Certificate”](#) on page 20 and the *CB Response User Guide* for more information.

Run the following script to create an encrypt backup of your certificates. The exact certificates are critical to disaster recovery efforts.

```
/usr/share/cb/cbssl backup --out <backup_file_name>
```

```
-----
SECURITY - SSL CERTIFICATE GENERATION
-----

Generating self-signed HTTPS Server certificate...

Generating self-signed HTTPS Sensor CA certificate...

Carbon Black Enterprise Response Server uses a SSL certificate to establish
secure communications between sensors and the server.

Should the certificate and/or its private key be lost, sensors will no longer
be able to communicate with the server.

We recommend backing up the SSL certificate files at this time by running:

  /usr/share/cb/cbssl backup --out <backup_file_name>

IMPORTANT: Backup file must be securely stored.  Anyone with access to the
information contained in that file will be able to compromise the security
of sensor-server communications and potentially compromise the security of
the computers on which the sensors run.

Continue [return]: 
```

- i. In the **IP Tables** section, answer `Y`. This opens port 433 in the server's IP tables.

```
-----
SECURITY - IPTABLES CONFIGURATION
-----

Carbon Black Enterprise Response Server listens on a number of TCP/IP ports. If
iptables firewall is running on the host machine, iptables must be configured
to allow incoming connections on these ports.

To get a list iptables rules that need to be added to current host
configuration, you can run '/usr/share/cb/cbcheck iptables -l' at any time and
apply the rules manually.  Alternatively, server setup and configuration
tools can take over management of iptables configuration and apply updates
whenever they are needed.

Would you like Carbon Black Enterprise Response Server to manage iptables [Y/n]: Y
Applying iptables rules:
  -I INPUT 5 -p tcp -m state --state NEW -m tcp --dport 443 -j ACCEPT
iptables: Saving firewall rules to /etc/sysconfig/iptables:[ OK ]
```

- j. The **POSTGRESQL Database Setup** section is automated and requires no user input.

```
-----  
SETTING UP POSTGRESQL DATABASE  
-----  
  
Initializing Carbon Black Server PostgreSQL Instance...  
  
The files belonging to this database system will be owned by user "cb".  
This user must also own the server process.  
  
The database cluster will be initialized with locale "en_US.UTF-8".  
The default text search configuration will be set to "english".  
  
Data page checksums are disabled.  
  
creating directory /var/cb/data/pgsql ... ok  
creating subdirectories ... ok  
selecting default max_connections ... 100  
selecting default shared_buffers ... 128MB  
creating configuration files ... ok  
creating template1 database in /var/cb/data/pgsql/base/1 ... ok  
initializing pg_authid ... ok  
setting password ... ok  
initializing dependencies ... ok  
creating system views ... ok  
loading system objects' descriptions ... ok  
creating collations ... ok  
creating conversions ... ok  
creating dictionaries ... ok  
setting privileges on built-in objects ... ok  
creating information schema ... ok  
loading PL/pgSQL server-side language ... ok  
vacuuming database template1 ... ok  
copying template1 to template0 ... ok  
copying template1 to postgres ... ok  
syncing data to disk ... ok  
  
Success. You can now start the database server using:  
  
    /usr/pgsql-9.3/bin/postgres -D /var/cb/data/pgsql  
or  
    /usr/pgsql-9.3/bin/pg_ctl -D /var/cb/data/pgsql -l logfile start  
  
waiting for server to start... done  
server started  
Creating alliance model DB schema...  
Creating core model DB schema...  
waiting for server to shut down... done  
server stopped
```

- k. In the **Setup Complete** section, enter **Y** to start the services.

```
-----
SETUP COMPLETE!
-----

Server setup has COMPLETED successfully.

Do you want to start the services [Y/n]: Y
Starting cb-supervisord:           [ OK ]
Starting cb-pgsql:                 [ OK ]
Starting cb-redis:                 [ OK ]
Starting cb-rabbitmq:             [ OK ]
Starting cb-solr:                  [ OK ]
    Waiting for cb-solr to build the terms dictionary.
    Depending on index size this may take a while...
Starting cb-coreservices:         [ OK ]
Starting cb-datastore:            [ OK ]
Starting cb-liveresponse:         [ OK ]
Starting cb-allianceclient:      [ OK ]
Starting cb-enterprise:           [ OK ]
Starting cb-nginx:                [ OK ]
-----

THANK YOU FOR INSTALLING CARBON BLACK ENTERPRISE RESPONSE!
-----
```

**Note:** To confirm sensor-to-server communications are functioning properly:

1. Open Google Chrome and then launch your server:

`https://<your_cber_server_url>`

2. Download a sensor and install it on an endpoint.

For more information on installing and managing sensors, see the “Manage Sensors” chapter in the *CB Response User Guide*.

7. Configure your firewall if you have not already done so. There are many ways to configure your firewall. The following is just an example for CentOS 6.
  - a. Open port 443 if you did not allow the `cbinit` script to manage iptables for you.

```
[bsmith@localhost yum.repos.d]$ sudo vim /etc/sysconfig/iptables
# Firewall configuration written by system-config-firewall
# Manual customization of this file is not recommended.
*filter
:INPUT ACCEPT [0:0]
:FORWARD ACCEPT [0:0]
:OUTPUT ACCEPT [0:0]
-A INPUT -m state --state ESTABLISHED,RELATED -j ACCEPT
-A INPUT -p icmp -j ACCEPT
-A INPUT -i lo -j ACCEPT
-A INPUT -m state --state NEW -m tcp -p tcp --dport 22 -j ACCEPT
-A INPUT -j REJECT --reject-with icmp-host-prohibited
-A FORWARD -j REJECT --reject-with icmp-host-prohibited
# New additions to the IPTABLES for carbon black
-A INPUT -m state --state NEW -m tcp -p tcp --dport 443 -j ACCEPT
COMMIT
```

- b. (Optional) Open port 80 to allow use of web interface and sensor communications through an unsecured channel. This is not required and only recommended for exploration or troubleshooting. Connections to the web interface through port 80 are redirected to port 443.
8. Log into the CB Response server web user interface at <https://<your server address>> and use the username and password that you set up in the `cbinit` script.

### Note

Google Chrome is the only supported browser for this release. Although not supported, internal testing indicates that Firefox, Opera, and IE10 or higher should work. However, IE browsers must not be in compatibility mode, and servers in the same subnet as the browser are automatically connected in this mode.

When the CB Response server is installed, configured, and initialized, it should be accessible through the web interface on port 443 with a self-signed certificate. If you attempt to access the web interface through HTTP on port 80, the connection is redirected to port 443.

The next step, especially in a test environment, is to download and install one or more sensors to begin collecting data. Sensor installation is described in the “Manage Sensors” chapter in the *CB Response User Guide*.

## Upgrading a CB Response Server

If you are upgrading the server, the procedure varies depending on whether:

- You are upgrading a standalone server or a clustered server.
- The database schema or CB Threat Intel feed data must be migrated after the new server version is installed.

These steps require SSH or console access to the server and minions with root privileges.

### To upgrade a standalone server:

1. On the server, stop the CB Response services:  

```
sudo service cb-enterprise stop
```
2. (Optional) Clean the yum cache of metadata and packages:  

```
yum clean all
```
3. Update the CB Response services:  

```
sudo yum upgrade cb-enterprise
```
4. Restart the CB Response services:  

```
sudo service cb-enterprise start
```

### To upgrade a clustered server:

1. On the Master server, navigate to the `cb` install directory (defaults to `/usr/share/cb`) and stop the CB Response services:  

```
sudo cbcluster stop
```
2. (Optional) Clean the yum cache of metadata and packages:  

```
yum clean all
```
3. Update the CB Response services on each Master and Minion server node:  

```
sudo yum update cb-enterprise
```
4. Run the `cbupgrade` utility (see the section below for more information.):  

```
/usr/share/cb/cbupgrade
```
5. On the Master server, restart the CB Response services:  

```
sudo cbcluster start
```

### Caution

On pre-6.1 versions of CB Response, it was possible to change the home directory of the `cb` service account, which is `/var/cb` by default. This is the account under whose auspices server installations occur. This is no longer supported and will cause upgraded servers to fail with an `nginx` error.

**The home directory of the `cb` service account must be `/var/cb`.** This is where upgrades place runtime configuration files for `nginx`, so if you changed the `cb` service home directory, these files will not be found. If you edited the home directory in `/etc/passwd`, restore it to `/var/cb` before upgrading.

## Using cbupgrade

Upgrades to the CB Response server will occasionally require you to use a utility called `cbupgrade` after executing `yum update cb-enterprise` to migrate the database schema or CB Threat Intel feed data. The operator will be notified of this requirement when attempting to start the `cb-enterprise` services. In a clustered-server configuration, you must run the `cbupgrade` tool on all nodes before restarting the cluster. When running this utility in a clustered environment, be sure to answer `No` when asked to start the CB Response services. You will need to use `cbcluster` to start the clustered server.

See [“Upgrading Cluster Nodes”](#) on page 64 for a description of `cbupgrade` options.

## Server Upgrades and New Sensor Versions

CB Response server releases generally include updated sensor versions for one or more operating systems. Servers and sensors can be upgraded independently, and sensors can be upgraded by sensor groups rather than all at once.

Decide if you would like the new sensor to be deployed immediately to existing sensor installations, or if you want to install only the server updates first. Carbon Black recommends a gradual upgrade of sensors to avoid any unacceptable impact on network and server performance and strongly recommends reviewing your Sensor group Upgrade Policies before upgrading your server to avoid inadvertently upgrading all of the sensors across your environment at once. For detailed information on Sensor Group Upgrade Policy, see the Sensor Group section of the *CB Response User Guide*

### 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.



## Supporting Multiple Volumes for Event Data

This section explains how customers can add more storage to their existing CB Response deployment after upgrading to the latest release. This involves adding multiple Solr data directories for cbevents cores. These directories can be added as mount points into new storage arrays, so that you can easily add more disk space. This can be easily configured. Basically, if you need more disk space, you attach a new volume, mount it into the Solr data directory, and the server starts using it automatically.

### Naming Conventions

Solr uses new `cbevents` directories (mount points) if their name is prefixed with:

`cbevents*`

or

`_cbevents*`.

#### Note

The `cbevents` directory (without the suffix) is the default directory but does not need to remain on the original data partition. You can remove it if needed.

The following is an example of a valid multi-volume configuration:

```
[root@ip-172-31-14-184 solr5]# df
Filesystem      1K-blocks    Used Available Use% Mounted on
/dev/xvda1      32895856 10341996  20860168  34% /
tmpfs           31389104         0  31389104   0% /dev/shm
/dev/xvdb       206293688 28033360 167758184  15% /data
/dev/xvdf       206293688 35809668 159981876  19% /data/solr5/cbevents2
/dev/xvdg       226936188 63976332 151409136  30% /data/solr5/cbevents3
```

In this example, the default data drive is mounted to `/dev/xvdb`, and `/data` is configured as the data root inside of `cb.conf`. In addition, two more volumes are added and mounted to `/data/solr6/cbevents2` and `/data/solr6/cbevents3`.

#### Caution

The system assigns the correct `user:group` upon `cb-enterprise` restart. If you created the mount points on a live server, ensure that the user assigned to the CB Response server has write permissions on the mounted directory. Failure to do so causes the system to ignore the new mount points.

Another option to expand `cbevents` storage is to use symlink.

#### To use a symlinked location for event storage:

1. Create a mount point in another location in the file system, such as `/data2`.
2. Create a symlink to the `cbevents*` directory inside the `solr6` directory that points to the mounted directory. For example:

```
ln -s /data2 /var/cb/data/solr6/cbevents2
```

3. Ensure that the CB Response user has write permissions in the mounted directory (`/data2`).

## Using New Data Directories

This section discusses partitioning and purging relating to new data directories.

### Partitioning

New data directories are used when the next partition occurs (every three days by default) or sooner if the current data disk is at risk of becoming full. The server uses simple heuristics in calculating when to partition and where to place the new event partition:

1. A new partition is created in the `cbevents*` directory with the most free space at the time of partitioning.
2. If the current data volume is more than 95% full and additional partitions exist that have more than 5% free space available, the server immediately partitions.

You can control this threshold using the following configuration parameter:

```
SolrTimePartitioningFreeSpaceThresholdPerc
```

- Rule 1 ensures that new volumes are used in a balanced fashion. As old data is aging out (being purged), some partitions free up. This ensures optimal use of free space.
- Rule 2 ensures that the system uses fragmented disk space efficiently in case many `cbevents*` directories exist. For example, assume you have five volumes, and each has 20% free space. This could result in none of the volumes fitting into the three-day partition. The system will continue trying to use one of the partitions (up to its maximum available space) before moving to the next one. As a result, the server might end up with smaller partitions. However, this scenario should be rare.

### Active and Read-Only Directories

Any `cbevents` directories prefixed with `cbevents*` will be used to create new cbevent partitions.

Any `cbevents` directories prefixed with `_cbevents*` will be used as read-only. They can be used to load existing partitions, but new partitions will not be created on it. This approach can be used when retiring old volumes. Old partitions will eventually be purged based on time. A second use for `_cbevents` prefix is for directories that are used only for "cold" partitions (old partitions that will only be loaded on demand).

### Partition Purging

The system purges partitions based on disk space, time, or the maximum number of allowed partitions.

When purging based on disk space, a purging algorithm considers the overall amount of free disk space. For example, three 100 GB volumes exist, each with 30 GB of free space, gives you a total of 90 GB of free space and a total disk space of 300 GB. The total event data size is the sum of index sizes on all three volumes. (This could be less than 210 GB since the main data volume may also contain store files and other data.)

The following shows how the current purging thresholds (in `cb.conf`) are interpreted when multiple volumes exist:

- `MaxEventStoreSizeInPercent` - Purge the oldest partition when the total sum of all event core sizes exceeds the given percentage of a total disk space (on all volumes).

- `MaxEventStoreSizeInMB` - Purge the oldest partition when the total event store size (on all volumes) exceeds the given threshold.
- `MinAvailableSizeInMB` - Purge the oldest partition when the total free disk space (on all volumes) falls below the given threshold.

## Extending Disk Space on the Fly

You can add disk space on the fly without having to restart their CB Response server. New directories are automatically used when a new partition occurs, avoiding any server downtime. As mentioned above, make sure that directory is given correct read+write permissions to the cb user.

## Troubleshooting the Server

The following table shows the CB Response server logs that are found in `/var/log/cb` and organized into subdirectories by component. CB Response Server Logs:

| Component      | Description                                                                                                                                                                                                                                                                               |
|----------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| allianceclient | The Alliance client communicates with the CB Response Alliance server.                                                                                                                                                                                                                    |
| audit          | Contains log files for the following activities: banning, sensor isolation, and live response. If <code>EnableExtendedApiAuditLogging</code> is enabled in <code>cb.conf</code> , this directory also includes a user activity log file based on user-generated API calls in the console. |
| cbfs           | Was the location of the datastore engine in earlier versions of CB Response but is no longer used in versions 5.0.0 and later.                                                                                                                                                            |
| cbfs-http      | Contains log files of the second generation Java datastore engine.                                                                                                                                                                                                                        |
| cli            | Contains events pertaining to the CB Response service commands used at the server console level.                                                                                                                                                                                          |
| coreservices   | Provides access to functionality via web APIs to both the web interface and sensors. Nearly all interface issues should result in log entries for coreservices.                                                                                                                           |
| sensorservices | Provides entry-point for sensor registrations and checkins. Look for issues here if there are problems with sensor connectivity                                                                                                                                                           |
| datastore      | Used for core event data processing and managing incoming sensor data.                                                                                                                                                                                                                    |
| enterprise     | Used for event logging of the CB Response service.                                                                                                                                                                                                                                        |
| job-runner     | The CB Response server uses cron jobs to provide various scheduled maintenance, data trimming, and similar tasks.                                                                                                                                                                         |
| liveresponse   | Used to hold CB Live Response session-related events.                                                                                                                                                                                                                                     |
| nginx          | The reverse proxy and SSL termination point for the CB Response server.                                                                                                                                                                                                                   |

| Component     | Description                                                                                                                                                                    |
|---------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| notifications | The location of the syslog output for feeds and watchlists.                                                                                                                    |
| pgsql         | The CB Response server uses Postgres SQL to store administrative data. Event data gathered from the sensors is not stored in Postgres.                                         |
| rabbitmq      | The logging location for the rabbitmq component of the CB Response server.                                                                                                     |
| redis         | The logging location for the redis component of the CB Response server.                                                                                                        |
| services      | The logging location for the start/stop services of the CB Response server.                                                                                                    |
| solr          | Used for indexes and stores data.                                                                                                                                              |
| supervisord   | The supervisord process utility is used to manage CB Response server processes, handling startup and shutdown dependencies between the various server components and services. |

The following table shows the scripts found in `/usr/share/cb`, most of which are diagnostics scripts; it only includes scripts that are used for diagnostics

| Component       | Description                                                                                                                                                                          |
|-----------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| cbbanning       | Assists in managing the CB Response server banning features. To get a list of available commands, run this command:<br><code>cbbanning commands</code>                               |
| cbstats         | This utility provides access to the statistics collected by the CB Response server.                                                                                                  |
| cbsyslog        | Provides an interface for testing CB Response's notifications syslog output.                                                                                                         |
| cbpost          | This utility is used to send file(s) to the Alliance server; typically used during interaction with Carbon Black Technical Support.                                                  |
| py_runtime_info | Generates a runtime report that shows the stack trace, process memory map, and open file descriptors for the running CB Response processes.                                          |
| cbfeed_scrubber | Helps clean up feed tags on existing Solr documents.                                                                                                                                 |
| cbinit          | Used to configure a combination of initial settings during a CB Response server installation.                                                                                        |
| cbdiag          | Dumps verbose troubleshooting information, including logs and configuration, to a gzip archive. This file can be analyzed offline or provided to Carbon Black with support requests. |
| sql_stats       | Contains outputs of various SQL database statistics; typically used during troubleshooting.                                                                                          |
| cb_solr         | Used for indexes and stores data.                                                                                                                                                    |

| Component             | Description                                                                                                                                                                                                                                                 |
|-----------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| cbget                 | This utility is used to download or list files from Alliance server; typically used during interaction with Carbon Black Technical Support.                                                                                                                 |
| sensor_report         | Generates a report that shows the status of every sensor communicating with CB Response server. Optionally, it can be used to identify specific sensors that might require the attention of IT support personnel.                                           |
| cbcheck               | Assists in troubleshooting CB Response server installation. To get a list of available commands, run this command:<br><code>cbcheck commands</code><br>to learn more about a specific command, run this command:<br><code>cbcheck &lt;command&gt; -h</code> |
| cbcluster             | Used to manage clusters (not a diagnostic tool).                                                                                                                                                                                                            |
| cb_rabbitmq-server.sh | This is a system utility and should never be run manually.                                                                                                                                                                                                  |
| cbrabbitmqctl         | A command-line interface that provides access to the CB Response rabbitmq service.                                                                                                                                                                          |
| pgsql_diag.sh         | Prints diagnostic info about the CBER Postgres database                                                                                                                                                                                                     |
| cbpasswd              | Resets user's password. Can only be run as root.                                                                                                                                                                                                            |

## Uninstalling a CB Response Server

This section describes the procedure for uninstalling a CB Response server from RHEL/CentOS.

You must remove the following CB Response packages:

- carbon-black-release
- cb-datagrid
- cb-datastore
- cb-enterprise
- cb-solr
- cb-swagger
- cbui
- libselenium-cb-python
- python-cb-coreservices
- python-cb-response-venv

## Removing the CB Response Server

Using the Yum utility, you can uninstall a CB Response server in one of two ways. You can remove or erase the packages that constitute the CB Response server or remove CB Response packages plus all other third-party packages upon which it depends. Instructions for both methods are provided here.

### To uninstall a server (CB Response RPMs only):

1. Stop the CB Response services with one of the following commands:

```
sudo /usr/share/cb/cbcluster stop
```

or

```
sudo service cb-enterprise stop
```

2. Use the following Yum utility command to remove CB Response packages listed above.

```
sudo yum remove <package1> <package2> <packageN>
```

3. Manually remove the following `cb` directories:

- /var/www/cb/
- /var/run/cb/
- /var/log/cb/
- /var/lib/cb/
- /var/cb/
- /usr/share/cb/
- /etc/cb/

You can also use the Yum utility to remove CB Response packages as well as all other third-party packages upon which it depends.

### Warning

This procedure might remove packages that are also required by other software applications. Use caution when performing this procedure.

### To uninstall a server (CB Response and related RPMs):

1. Stop the CB Response services with one of the following commands:

```
sudo /usr/share/cb/cbcluster stop
```

or

```
sudo service cb-enterprise stop
```

2. Access the `yum.conf` file with the following command:

```
vi /etc/yum.conf
```

3. Add this line to `yum.conf`:

```
clean_requirements_on_remove=1
```

4. Enter the following:

```
yum erase cb-enterprise
```

5. Enter the following:

```
yum remove carbon-black-release
```

**6. Manually remove the following `cb` directories:**

- `/var/www/cb/`
- `/var/run/cb/`
- `/var/log/cb/`
- `/var/lib/cb/`
- `/var/cb/`
- `/usr/share/cb/`
- `/etc/cb/`

## Chapter 3

# Server Backup and Restoration

This chapter explains how to perform various backup and restore procedures.

### Sections

| Topic                    | Page |
|--------------------------|------|
| <a href="#">Overview</a> | 41   |
| <a href="#">Backup</a>   | 42   |
| <a href="#">Restore</a>  | 45   |



## Overview

This chapter provides procedures for backing up and restoring a CB Response server. The procedures are designed to ensure minimal data loss in the event of a catastrophic failure.

Backup files and data should be stored on a different server than the one that is used for daily operations.

All procedures in this document are performed at the command prompt and require root-level access.

### Notes

- `/var/cb` is assumed to be the default installation data path throughout this appendix. If you have your datastore root configured in a different location, use that in place of what is provided here.
- Network integrations will not be fully backed up. Any network integration needing a bridge or connector installed must be installed on the new restore server before restoring the configurations. Configuration items for network integrations will be restored throughout this appendix.

## Restoration Servers

For all restoration options, it is assumed that a new server install has been performed with the same number of master and minion systems, each with the same number of Solr data shards in place on each system, and configured to use the same hostname and IP addresses.

The same server version must be used between backed up and restored systems.

To install a new system, see [“Installing the CB Response Server”](#) on page 18.

All installation steps involve running `cbinit` (and `cbcluster add-node` for clustered systems). These must be completed before performing a system restoration.

# Backup

This section highlights two backup and restore methods:

- A configuration-only backup covers all configuration and sensor metadata for the system. It is designed to be minimal in storage size and allows the quickest return to operation in the event of a catastrophic server failure. It does not include any data that the sensors submit (event or binary data) or any feed or alert information from a running server. The configuration backup step is required for both backup and restore methods. For more information, see [“Configuration Backup”](#) on page 42.
- A data backup is a continuation of the configuration backup and includes all data stored by the server. Backing up this data ensures a full recovery of the system. Off-server storage requirements are much larger due to the full data set that is being backed up. Also, the time it takes to perform a restoration will likely be longer due to copying and unarchiving large amounts of data to a new system.

## Note

`/cbdata/_backup/ServerName` is used in this procedure as an example for a backup data storage location. Determine the correct location for storing your backups and note the size requirements if also performing a data backup.

You must run all commands on the master and minion systems unless otherwise noted. Perform all steps on all standalone servers. For more information, see [“Data Backup”](#) on page 44.

## Configuration Backup

A configuration backup backs up only the files and data required for a system restore that has functioning sensors, without recorded data. You can perform this configuration backup by itself or use it as the first step in a full backup. It is the quickest way to capture and restore data and consumes the smallest amount of disk space. Configuration backups can be performed while CB Response is running.

### To perform a configuration backup:

1. Change directories to the backup location:

```
cd /cbdata/_backup/ServerName
```

2. Backup the configuration files by running these commands:

- a. Host file:

```
tar -P --selinux -cvf cbhosts.tar /etc/hosts
```

- b. Yum files:

```
tar -P --selinux -cvf cbyum.tar /etc/yum.repos.d
```

- c. Firewall configuration:

- For RHEL v6 (iptables):

```
tar -P --selinux -cvf cbiptables.tar /etc/sysconfig/iptables
```

- For RHEL v7 (firewalld):

```
tar -P --selinux -cvf cbfirewalld.tar /etc/firewalld
```

**d. SSH configuration and keys:**

```
tar -P --selinux -cvf cbssh.tar /etc/ssh/
```

**e. CB Response configuration:**

```
tar -P --selinux -cvf cbconfig.tar /etc/cb/
```

**f. Rsyslog configuration:**

```
tar -P --selinux -cvf cbrsyslog.tar /etc/rsyslog.conf
```

**g. Rsyslog.d configuration:**

```
tar -P --selinux -cvf cbrsyslogd.tar /etc/rsyslog.d/
```

**h. RabbitMQ cookie:**

```
tar -P --selinux -cvf cbrabbitmqcookie.tar /var/cb/.erlang.cookie
```

**i. RabbitMQ node configuration:**

```
tar -P --selinux -cvf cbrabbitmqnode.tar /var/cb/data/rabbitmq
```

**j. (Optional) SSH authorization keys:**

**Note:** Perform this step only if you use trusted keys between systems in a cluster environment.

```
tar -P --selinux -cvf cbrootauthkeys.tar /root/.ssh/authorized_keys
```

**k. (Master only) Syslog CEF template:**

```
tar -P --selinux -cvf cbceftemp.tar /usr/share/cb/syslog_templates
```

**l. (Optional - master only) CB Response installer backups:**

**Note:** Perform this step only if you manually installed additional versions of the sensor.

```
tar -P --selinux -cvf cbinstallers.tar /usr/share/cb/coreservices/installers/
```

**m. (Optional - master only) Custom syslog templates:**

**Note:** Perform these steps only if you use custom syslog templates that are not stored in /user/share/cb/syslog\_templates.

**i. Locate all custom syslog templates paths in use by searching the following:**

/etc/cb/cb.conf file for any instance of SyslogTemplate=

For example:

```
WatchlistSyslogTemplateBinary
```

and

```
FeedIngressSyslogTemplateBinary
```

**ii. Take note of the file path after =.**

For example:

```
WatchlistSyslogTemplateBinary=/var/custom/syslog/
```

```
watchlist_binary_custom.template
```

The path for this example would be /var/custom/syslog

**iii. Tar each custom path that is identified using this command:**

```
tar -P --selinux -cvf syslog_custom1.tar /var/custom/syslog
```

### 3. (Master only) Perform a backup of the Postgres Database:

**Note:** Perform this configuration only if you are performing a configuration-only backup. Otherwise, skip this step and execute the step for backing up a Postgres Database in “[Data Backup](#)” on page 44.

#### a. Backup configuration:

```
pg_dump -C -Fp -f psqldump_config.sql cb -p 5002 \
--exclude-table-data=allianceclient_comm_history \
--exclude-table-data=allianceclient_uploads \
--exclude-table-data=allianceclient_pending_uploads \
--exclude-table-data=banning_sensor_counts \
--exclude-table-data=binary_status \
--exclude-table-data=cb_useractivity \
--exclude-table-
data=detect_dashboard_average_alert_resolution_history \
--exclude-table-data=detect_dashboard_binary_dwelling_history \
--exclude-table-data=detect_dashboard_host_hygiene_history \
--exclude-table-data=investigations \
--exclude-table-data=maintenance_job_history \
--exclude-table-data=moduleinfo_events \
--exclude-table-data=mutex_watchlist_searcher \
--exclude-table-data=sensor_activity \
--exclude-table-data=sensor_comm_failures \
--exclude-table-data=sensor_driver_diagnostics \
--exclude-table-data=sensor_event_diagnostics \
--exclude-table-data=sensor_licensing_counts \
--exclude-table-data=sensor_queued_data_stats \
--exclude-table-data=sensor_resource_statuses \
--exclude-table-data=server_storage_stats \
--exclude-table-data=storefiles \
--exclude-table-data=tagged_events
```

#### b. Backup users and groups:

```
pg_dumpall -p 5002 --roles-only -f psqlroles.sql
```

### 4. Copy the data from the backup location off to a remote location.

## Data Backup

A data backup captures all recorded data that is stored on a server and is required to complete a full restore of a functioning system. You must complete a “[Configuration Backup](#)” on page 42 before performing the steps in this section. The data captured here can be very large, depending on the amount of data that is retained by the system or cluster. Full backups require that CB Response is in a stopped state.

### To perform a data backup:

#### 1. Stop all CB Response services by running these commands:

- In a clustered server environment, on the master only, run:
 

```
/usr/share/cb/cbcluster stop
```
- or-
- In a standalone server environment, run:
 

```
service cb-enterprise stop
```

2. Change directories to the backup location:  

```
cd /cbdata/_backup/ServerName
```
3. Backup the Solr database by running this command:  

```
tar -P --selinux -cvf cbsolr.tar /var/cb/data/solr6/
```
4. Backup the module store:  

```
tar -P --selinux -cvf cbmodulestore.tar /var/cb/data/modulestore/
```
5. (Master only) Backup Postgres Database:
  - a. Start the Postgres service:  

```
service cb-pgsql start
```
  - b. Backup the Postgres database:  

```
pg_dump -C -Fp -f psqldump_full.sql cb -p 5002
```
  - c. (This is a duplicate step from “[Configuration Restore](#)” on page 47.) Backup users and groups:  

```
pg_dumpall -p 5002 --roles-only -f psqroles.sql
```
  - d. Stop the Postgres service:  

```
service cb-pgsql stop
```
6. Start the CB Response server by running these commands:
  - a. In a clustered server environment, on the master only, run:  

```
/usr/share/cb/cbcluster start
```
  - b. In a standalone server environment, run:  

```
service cb-enterprise start
```
7. Copy the data from the backup location to a remote location.

## Restore

To restore a system, the following prerequisites must be met:

- A fresh CB Response server (or an old snapshot) installation must be available on which to restore the backup files.
- The server must have the same configuration for the master and minions (for a clustered environment) as the system that was backed up.
- The same hostname and IP addresses must be used.
- The new server(s) must be installed on the same version of the server(s) from which the backups were taken (for example, v7.0.0).

If detailed configuration items are required to complete the installation, use the following files, which are generated in “[Configuration Backup](#)” on page 42, to obtain the configurations from the backed-up server:

- The IP address used for each system. In `cbhosts.tar`, this is in `/etc/hosts`.
- The number of systems used and the master/minion assignment. In `cbconfig.tar`, this is in the `/etc/cb/cbcluster.conf` file.

To install a new system, see “[Installing the CB Response Server](#)” on page 18. You must complete all installation steps including those that involve running `cbinit` (and `cbcluster add-node` for clustered environments) before performing a server restore. The configuration items chosen when running `cbinit` are overwritten after restoration of the configuration files is complete.

## Failed Minion Cluster Restore

The procedure in this section helps you set up a fresh server install to restore one or more failed minions to an otherwise operational cluster.

The following prerequisites must be met before performing this procedure:

- The master server (at least) must still be functioning and backup data for the failed minions must be available.
- The servers on which the restored minions will be installed should be at the same base operating system level as the rest of the cluster, with no Carbon Black software installed.

### To install the server and restore data on failed minions of an otherwise operational cluster:

1. Stop the cluster:

```
/usr/share/cb/cbcluster stop
```

2. Start the Postgres service:

```
service cb-pgsql start
```

3. Determine the `node_id` for missing minion(s) by querying Postgres:

```
psql -p 5002 cb -c "select * from  
cluster_node_sensor_addresses;"
```

4. Delete the row from the `cluster_node_sensor_address` table, where the `node_id` is equal to the failed minion(s):

```
psql -p 5002 cb -c "delete from cluster_node_sensor_addresses  
where node_id = 2;"
```

where 2 is the number from the missing minion

5. Stop Postgres service:

```
service cb-pgsql stop
```

6. Edit the `/etc/cb/cluster.conf` file as follows:

- a. In the `[Cluster]` section of the file, subtract `N` from `NodeCount` and `NextSlaveAutoInc` lines where `N` is the number of failed minions being restored.

- b. Delete the `[SlaveN]` section for each failed minion.

- c. Save the file.

7. Add the minion(s) like you would as part of a normal install by running:

```
/usr/share/cb/cbcluster add-node
```

and providing the appropriate IP address or hostname.

8. The master will install `cb-enterprise` on the minion and allow for normal operation.

9. Copy over the backup data generated from the backup process and follow the restore procedures in the next sections.

10. Start the cluster:

```
/usr/share/cb/cbcluster start
```

## Configuration Restore

This section contains a procedure for restoring the configuration files and data files. The restore should match the backup procedures performed in “Backup” on page 42.

`/cddata/_backup/ServerName` is used in this procedure as an example for a backup data storage location. Determine the correct location for storing your backups and note the size requirement if you are also performing a data backup.

The following prerequisites must be met before performing this procedure:

- For standalone server environments, all commands in the following sections must be run except for those specifically identified as being for cluster environments only.
- For clustered server environments, all commands in the following sections must be run for the master and minion servers unless otherwise noted.
- All restoration activities require that CB Response be placed in a stopped state.

### To perform a configuration restore:

1. Stop CB Response by running these commands:

- (Master only) In a clustered server environment, run:  
`/usr/share/cb/cbcluster stop`
- *- or -*
- In a standalone server environment, run:  
`service cb-enterprise stop`

2. Delete any stored ssh keys (files may not exist):

```
rm /root/.ssh/known_hosts
```

3. Change directories to the backup location:

```
cd /cbdata/_backup/ServerName
```

4. Restore configuration files as follows:

a. Hosts file:

```
tar -P -xvf cbhosts.tar
```

b. Yum files:

```
tar -P -xvf cbyum.tar
```

c. Firewall configuration:

- For RHEL v6 (iptables):

```
tar -P -xvf cbiptables.tar
```

- For RHEL v7 (firewalld):

```
tar -P -xvf cbfirewalld.tar
firewall-cmd --reload
```

d. SSH keys:

```
tar -P -xvf cbssh.tar
```

e. CB Response configuration:

```
tar -P -xvf cbconfig.tar
```

i. Clear the server.token:

```
rm /etc/cb/server.token
```

ii. Grab the new server token:

```
/usr/share/cb/virtualenv/bin/python -c "from
cb.alliance.token_manager import SetupServerToken;
SetupServerToken().set_server_token('/etc/cb/server.token')"
```

- f. Configure Rsyslog:
 

```
tar -P -xvf cbrsyslog.tar
```
  - g. Configure Rsyslog.d:
 

```
tar -P -xvf cbrsyslogd.tar
```
  - h. Configure Rabbitmq cookie:
 

```
tar -P -xvf cbrabbitmqcookie.tar
```
  - i. Configure Rabbitmq node:
 

```
tar -P -xvf cbrabbitmqnode.tar
```
  - j. (Optional) SSH authorization keys:
 

```
tar -P -xvf cbrootauthkeys.tar
```
  - k. (Master only) Syslog CEF template:
 

```
tar -P -xvf cbceftemp.tar
```
  - l. (Optional - Master only) CB Response installer backups:
 

```
tar -P -xvf cbinstallers.tar
```
  - m. (Optional - Master only) Custom syslog templates; for each custom template TAR:
 

```
tar -P -xvf syslog_custom1.tar
```
5. (Master only) Restore Postgres database:

**Note:** If you are performing a full data restore, skip this procedure.

- a. Start the Postgres database:

```
service cb-pgsql start
```

- b. Drop the old database:

```
dropdb cb -p 5002
```

- c. Restore roles:

**Note:** An error indicates that the account is already created.

```
psql template1 -p 5002 -f psqlroles.sql
```

- d. Restore the data dump:

```
psql template1 -p 5002 -f psqldump_config.sql
```

- e. Restore the database sequence values. Copy and paste the following into a file called `psqlcbvalues` on the server:

```
SELECT
pg_catalog.setval('allianceclient_comm_history_id_seq', 1,
false);
SELECT
pg_catalog.setval('allianceclient_pending_uploads_id_seq',
1, false);
SELECT pg_catalog.setval('allianceclient_uploads_id_seq', 1,
false);
SELECT pg_catalog.setval('cb_useractivity_id_seq', 1,
false);
SELECT
pg_catalog.setval('detect_dashboard_average_alert_resolution
_history_id_seq', 1, false);
SELECT
pg_catalog.setval('detect_dashboard_binary_dwell_history_id_
seq', 1, false);
SELECT
pg_catalog.setval('detect_dashboard_host_hygiene_history_id_
seq', 1, false);
SELECT pg_catalog.setval('investigations_id_seq', 1, false);
```



```

SELECT pg_catalog.setval('maintenance_job_history_id_seq',
1, false);
SELECT pg_catalog.setval('moduleinfo_events_id_seq', 1,
false);
SELECT pg_catalog.setval('sensor_comm_failures_id_seq', 1,
false);
SELECT pg_catalog.setval('sensor_driver_diagnostics_id_seq',
1, false);
SELECT pg_catalog.setval('sensor_event_diagnostics_id_seq',
1, false);
SELECT pg_catalog.setval('sensor_licensing_counts_id_seq',
1, false);
SELECT pg_catalog.setval('sensor_queued_data_stats_id_seq',
1, false);
SELECT pg_catalog.setval('sensor_resource_statuses_id_seq',
1, false);
SELECT pg_catalog.setval('server_storage_stats_id_seq', 1,
false);
SELECT pg_catalog.setval('tagged_events_id_seq', 1, false);

```

- f. **Restore psqldbvalues:**  

```
psql cb -p 5002 -f psqldbvalues
```
  - g. **Rebuild default investigations:**  

```
psql cb -p 5002 -c "INSERT INTO investigations VALUES
('1','Default Investigation',to_timestamp((select value from
cb_settings where key='ServerInstallTime'),'YYYY-MM-DD
hh24:mi:ss'),NULL,to_timestamp((select value from
cb_settings where key='ServerInstallTime'),'YYYY-MM-DD
hh24:mi:ss'),'Automatically Created at Installation Time');"
```
  - h. **Remove all previous query-based feeds:**  

```
psql cb -p 5002 -c "delete from watchlist_entries where
group_id <> '-1';"
```
  - i. **Stop the Postgres database:**  

```
service cb-pgsql stop
```
6. **Start the CB Response services:**
- Note:** If performing a **data** restoration, skip these steps.
- (Master only) In a clustered server environment, run:  

```
/usr/share/cb/cbcluster start
```

- or -
  - In a standalone server environment, run:  

```
service cb-enterprise start
```

## Data Restore

This section describes how to restore the configuration files and data files. The restore should match the backup procedures in [“Backup”](#) on page 42.

`/cddata/_backup/ServerName` is used in this procedure as an example for a backup data storage location. Determine the correct location for storing your backups and note the size requirement if you are also performing a data backup.

The following prerequisites must be met before performing this procedure:

- You must perform the procedure in [“Configuration Restore”](#) on page 47 before performing a data restore.

- For clustered server environments, all commands in the following sections must be run for the master and minion servers unless otherwise noted.
- For standalone server environments, all commands in the following sections must be run as well.

**To perform a data restore:**

1. Make sure that the CB Response server is in a stopped state after performing the [“Configuration Restore”](#) on page 47.

2. Delete and restore the Solr data:

```
rm -rf /var/cb/data/solr6/cbevents/0/data
tar -P -xvf cbsolr.tar
```

3. Restore the module store by running these commands:

```
rm -rf /var/cb/data/modulestore/*
tar -P -xvf cbmodulestore.tar
```

4. (Master only) Restore the Postgres database by running these commands:

- a. Start the Postgres service:

```
service cb-pgsql start
```

- b. Drop the old database:

```
dropdb cb -p 5002
```

- c. Restore roles:

**Note:** You will receive an error indicating that the account is already created.

```
psql template1 -p 5002 -f psqlroles.sql
```

- d. Restore the Postgres data:

```
psql template1 -p 5002 -f psqldump_full.sql
```

- e. Stop the Postgres service:

```
service cb-pgsql stop
```

5. Start the CB Response services:

- a. (Master only) In a clustered server environment, run:

```
/usr/share/cb/cbcluster start
```

- b. In a standalone server environment, run:

```
service cb-enterprise start
```

## Chapter 4

# Ports and Protocols

This chapter provides port and protocol information for several different server communications.

### Note

Underlying IP addresses of any servers identified can change. Avoid listing any specific IP addresses; instead, configure your firewalls to use DNS names.

| Communication                                          | Port      | Protocol             | Comment                                                                                                                                                                  |
|--------------------------------------------------------|-----------|----------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Management Station to Cb Response Server               | TCP 22    | SSH                  | A management station is a machine from which system administrators can SSH into the CB Response Server and address any required administrative tasks or troubleshooting. |
| Management Station to Cb Response Server               | TCP 443   | HTTPS (configurable) |                                                                                                                                                                          |
| Sensor to Cb Response Server                           | TCP 443   | HTTPS (configurable) | NA                                                                                                                                                                       |
| Master Cb Response Server to Minion Cb Response Server | TCP 22    | SSH                  | MinionApiPort in the cb.conf file allows configuration of the HTTPS port configuration.                                                                                  |
|                                                        | TCP 443   | HTTPS (configurable) |                                                                                                                                                                          |
|                                                        | TCP 4369  | RabbitMQ             |                                                                                                                                                                          |
|                                                        | TCP 5701  | datagrid             |                                                                                                                                                                          |
|                                                        | TCP 6379  | REDIS                |                                                                                                                                                                          |
|                                                        | TCP 6500  | sensorservices       |                                                                                                                                                                          |
|                                                        | TCP 6501  | sensorservices       |                                                                                                                                                                          |
|                                                        | TCP 8080  | SOLR                 |                                                                                                                                                                          |
|                                                        | TCP 9000  | CB data store        |                                                                                                                                                                          |
|                                                        | TCP 25004 | RabbitMQ             |                                                                                                                                                                          |

| Communication                                                      | Port      | Protocol       | Comment                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|--------------------------------------------------------------------|-----------|----------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Minion<br>CB Response<br>Server to Master<br>CB Response<br>Server | TCP 4369  | RabbitMQ       | NA                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|                                                                    | TCP 5002  | POSTGRES       |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|                                                                    | TCP 5600  | liveresponse   |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|                                                                    | TCP 5701  | datagrid       |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|                                                                    | TCP 6379  | REDIS          |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|                                                                    | TCP 6500  | sensorservices |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|                                                                    | TCP 6501  | sensorservices |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|                                                                    | TCP 8080  | SOLR           |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|                                                                    | TCP 25004 | RabbitMQ       |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| Minion<br>CB Response<br>Server to Minion<br>CB Response<br>Server | TCP 4369  | RabbitMQ       | NA                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|                                                                    | TCP 5701  | datagrid       |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|                                                                    | TCP 6500  | sensorservices |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|                                                                    | TCP 6501  | sensorservices |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|                                                                    | TCP 8080  | SOLR           |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|                                                                    | TCP 25004 | RabbitMQ       |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| CB Response<br>Server to a CB<br>Alliance Server                   | TCP 443   | HTTPS          | <p>For URLs that can accept CB Alliance Server communications, see one of the following:</p> <ul style="list-style-type: none"> <li>• <a href="https://api.alliance.carbonblack.com">api.alliance.carbonblack.com</a></li> </ul> <p>Points to the CB Alliance Server and has a single IP behind it that will change for various reasons over time.</p> <ul style="list-style-type: none"> <li>• <a href="https://api2.alliance.carbonblack.com">api2.alliance.carbonblack.com</a></li> </ul> <p>Points to the CB Alliance Server and the single IP behind it may be different or the same as <a href="https://api.alliance.carbonblack.com">api.alliance.carbonblack.com</a>. In addition, it may change for various reasons over time.</p> <p>Note that the IPs behind these servers are subject to change.</p> |

| Communication                          | Port    | Protocol | Comment                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|----------------------------------------|---------|----------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| CB Response Server to CB Threat Intel  | TCP 443 | HTTPS    | For a URL that accepts CB Threat Intel communications, see <a href="https://threatintel.bit9.com">threatintel.bit9.com</a> . This URL has multiple elastic IPs behind it, and it points to the "next-gen" CB Threat Intel infrastructure. Note that the IPs behind this URL are subject to change.                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| CB Response Server to YUM Repositories | TCP 443 | HTTPS    | For an API on this type of communication, see <a href="https://yum.distro.carbonblack.io">yum.distro.carbonblack.io</a> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|                                        | TCP 80  | HTTP     | <p>For an API on this type of communication, see <a href="https://mirror.centos.org">mirror.centos.org</a> or any other kind of enabled repository.</p> <p>Note that CB Response uses the default CentOS configuration when installing base CentOS packages.</p> <p>Any communication supporting the installation or update of YUM repositories utilizes <a href="https://mirror.centos.org">mirror.centos.org</a> first, but that host is used only to identify the current mirror server list. After that, <i>any</i> of the available mirror servers might be chosen to download the actual packages.</p> <p>If this default CentOS behavior is a problem, ask your system administrator to change the CentOS configuration.</p> |

## Chapter 5

# Installing a CB Response Cluster

This chapter introduces CB Response clusters and explains how to configure clusters, add minions to existing clusters, remove minion nodes from clusters, and upgrade cluster nodes.

### Note

Some folders and scripts use the term “slave”, which is another term for “minion”.

### Sections

| Topic                                                   | Page |
|---------------------------------------------------------|------|
| <a href="#">Overview</a>                                | 55   |
| <a href="#">Configuring CB Response Clusters</a>        | 58   |
| <a href="#">Adding Minions to Existing Clusters</a>     | 61   |
| <a href="#">Removing Minions from Existing Clusters</a> | 62   |
| <a href="#">Upgrading Cluster Nodes</a>                 | 64   |

## Overview

A CB Response cluster is a group of servers fulfilling certain roles, operating as a single CB Response instance. Clusters are used to support a larger number of sensors and data than a single standalone server can support. This chapter provides details about setting up a CB Response cluster.

CB Response collects a large volume of data at varying degrees of velocity. As volume and velocity increase, the server begins to reach the limit of what a single server can support. In order to scale to support increasing volume and velocity, the infrastructure must grow horizontally. This means adding additional servers to the infrastructure. A cluster of CB Response servers must be configured when the desired number of sensors, activity level, or amount of retention exceeds a certain threshold, which negatively impacts the performance of a single CB Response server instance.

Horizontal scaling is used by Solr, which is the core component that CB Response uses to store data. Solr is a big data solution that uses distributed indexes and distributes queries to those indexes. This increases query performance for high volume and velocity indexing. To support optimal scaling performance, the CB Response infrastructure is modeled around this concept.

See the *CB Response Operating Environment Requirements (OER)* document for more information on server performance and hardware requirements.

## Cluster Architecture

CB Response cluster servers have one of two roles:

- *Master* – Also known as a head-node
- *Minion* – Also known as an index(ing) node

Each cluster membership role fulfills specific functions for a single CB Response instance. Together, the nodes in a cluster form a functioning CB Response server with specific functionalities and internal components configured to perform their membership roles.

The following table provides a matrix for CB Response services and certain components compared to each role:

| Service/Component | Standalone | Master | Minion |
|-------------------|------------|--------|--------|
| cb-nginx          | Yes        | Yes    | Yes    |
| cb-datastore      | Yes        | Yes    | Yes    |
| cb-coreservices   | Yes        | Yes    | Yes    |
| cb-sensorservices | Yes        | Yes    | Yes    |
| cb-liveresponse   | Yes        | Yes    | No     |
| cb-allianceclient | Yes        | Yes    | Yes    |
| cb-datagrid       | Yes        | Yes    | Yes    |
| cb-enterprise     | Yes        | Yes    | Yes    |
| - ThreatIntel     | Yes        | Yes    | No     |

| Service/Component                        | Standalone | Master | Minion |
|------------------------------------------|------------|--------|--------|
| - Statistics                             | Yes        | Yes    | Yes    |
| cb-redis                                 | Yes        | Yes    | Yes    |
| cb-pgsql                                 | Yes        | Yes    | No     |
| cb-solr                                  | Yes        | Yes    | Yes    |
| - Process data<br>(cbevents core shards) | Yes        | No     | Yes    |
| - Binary Info data<br>(cbmodules core)   | Yes        | Yes    | Yes*   |
| - Feed data<br>(cbfeeds core)            | Yes        | Yes    | No     |
| cb-rabbitmq                              | Yes        | Yes    | Yes    |
| Binary (modulestore)                     | Yes        | No     | Yes    |

\* replicated by Master

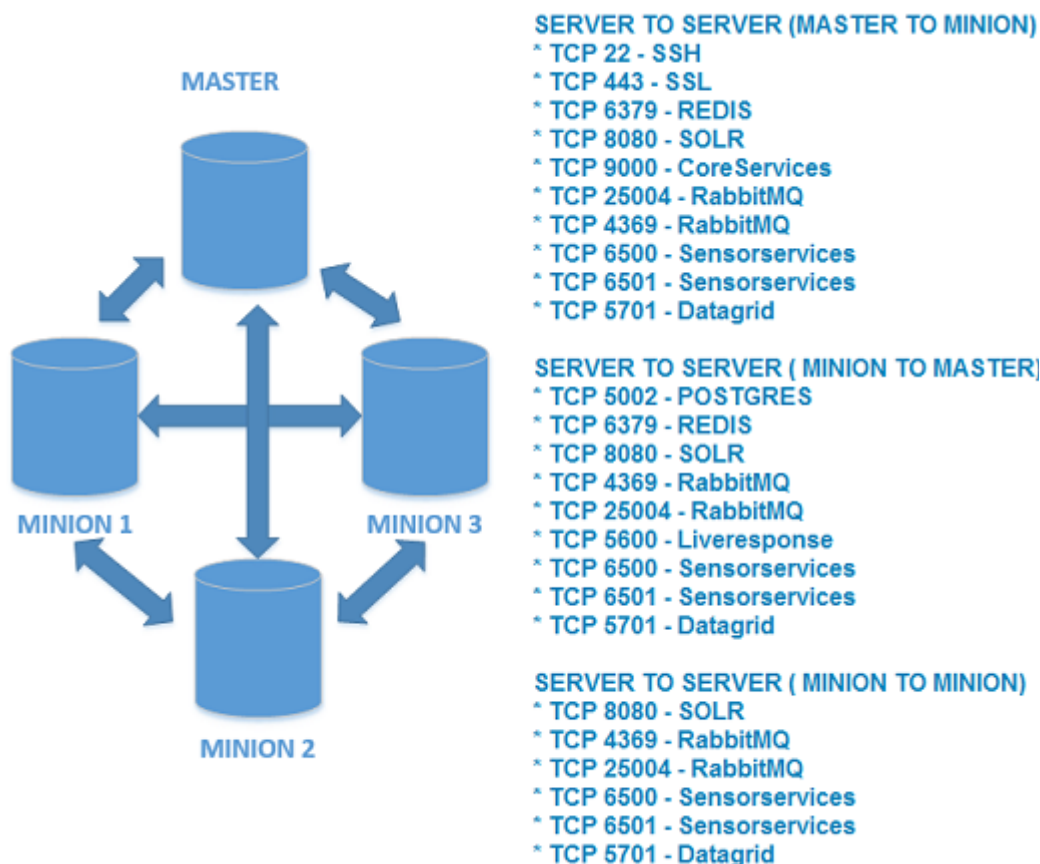
Solr clustering is performed by breaking up a single Solr core into multiple cores called shards. Shards, which are identified by a numerical integer starting at 0, are then evenly distributed to the indexing nodes (minions) for individual management and distributed querying.

For smaller cluster implementations, the master can also perform the role of an indexing server and fulfill all roles similar to a standalone server. However, for larger organizations with more than 60,000 sensors or four minions, the cluster must have a dedicated master that is not performing minion duties.

When CB Response is clustered, internal nodal communication must occur so that the application behaves as a single instance. Most internal components can operate in this distributed capacity by making standard calls to the other components and/or nodes.



The following diagram illustrates inter-nodal cluster communications:



In order for RabbitMQ to function properly as a single system, it must also be clustered. RabbitMQ is an application message bus used to exchange data (or messages) between the processes, application, and servers. Creating a RabbitMQ cluster allows CB Response to exchange and queue internal (server process and application) and inter-nodal (node-to-node) messaging for proper operation as a single CB Response instance.


## Cluster Operation

A CB Response cluster must operate with multiple servers in the same way as a standalone instance. As a result, each role must be responsible for certain aspects of a CB Response instance.

The master contains the user interface and is the primary front end for the API and most integrations. As users navigate the web console, standard API endpoint calls will be made to the master, which in turn queries the appropriate back-end storage location. When process data is queried, the master distributes the query to the minions and renders the aggregated results. However, if a binary, threat intel, or alert search is performed, it only queries its local Solr cores.

The master also contains the only instance of the PostgreSQL database that houses most of the application-specific configuration and certain state information that is used for sensor management. The master is in charge of managing the sensor configuration and communications, which include the CB Live Response capabilities. Managing sensors not only requires sensor state information but also minion state metrics to independently distribute minion bandwidth to allow its assigned sensors to submit data.

The minions serve as the primary ingestion point for sensor data. When a sensor is told to send data to the minion by the master, the minion receives the data and begins to ingest the data for storage. A minion will store all process-related data in its cbevents Solr core(s), where it manages its indexes separately from the rest of the cluster. This data is retrieved when a distributed query is initiated by the master. Binary data is forwarded to the master for storage and management. This is only done once per unique binary per sensor. Unlike the binary metadata, the copies of the binaries are stored locally on the minion. The copies are only submitted and stored once per CB Response cluster or instance. The distribution of the binary storage depends on the sensor to which that binary was submitted. This data flow is illustrated in the following figure:



## Configuring CB Response Clusters

This section describes how to configure a CB Response cluster containing a master node and any number of minion nodes. This process involves two generic CB Response server installs, which have no client software configured on them initially. The goal is to create a CB Response cluster with `cb-master` as the master node and `cb-minion` as the minion node.

You must execute the initial `cbinit`, which runs on the master node.

If the node is a master and the cluster has more than three nodes, add the following parameter to instruct the master node *not* to store events:

```
--no-solr-events
```

**Note**

The following instructions assume that you have installed the CB Response RPM and have run the `yum install cb-enterprise` command on the master node. The minion is simply a generic install of CentOS. For more information, see [“Installing and Initializing a CB Response Server”](#) on page 22.

**To set up the cluster configuration:**

1. On the master node, issue the `cbinit` command with the correct `--no-solr-events` flag defined, based on your installation. In this example, we will instruct the master node to not store events.

```
/usr/share/cb/cbinit --no-solr-events
```

You will use the following CB Response cluster management command-line tool options to initiate the cluster configuration:

```
[root@cb--?master~]# /usr/share/cb/cbcluster
```

Usage: `cbcluster` COMMAND [CMD OPTIONS]

Available commands are:

`help` – Display this help screen

`start` – Start the cluster

`stop` – Stop the cluster

`status` – Get the running status of the cluster

`add-node` – Add a minion node to the cluster

`change-node` – Change parameters of the existing cluster node

`remove-node` – Remove a minion node from the cluster

2. From `cb-master`, run the following command to initiate the cluster configuration (see [“Best Practices”](#) on page 61):

```
[root@cb-master~]# /usr/share/cb/cbcluster add-node
```

3. Enter the following information when prompted:

- For the hostname or IP address of the remote node, enter the IP address of the server to become a minion node. For example: `172.xx.xxx.xxx`.
- For the password of the server that will become a minion node, enter the root password for the server in question; do not enter the CB Response password.

The CB Response software is now installed on each minion. The yum repo configuration used on the master is used on each node added.

4. When the minion nodes have all been configured, start the cluster services:

```
[root@cb-master~]# /usr/share/cb/cbcluster start
```

5. You can now view the results in the config file:

```
[root@cb-master~]# cd /etc/cb/
```

```
[root@cb-master~]# less cluster.conf
```

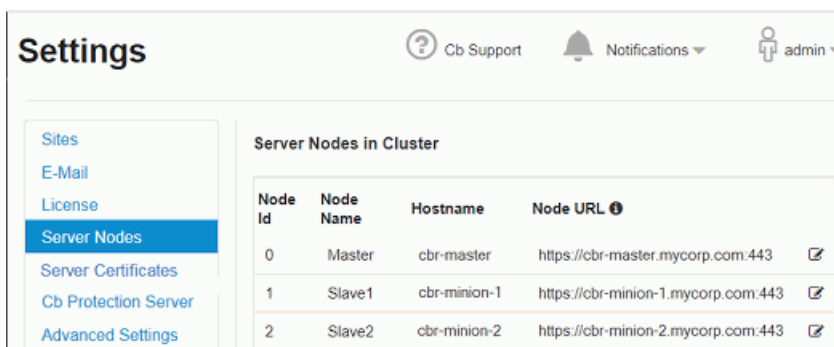
```
#####
#####
#
# /etc/cb/cluster.conf:
# This file contains CB Response server cluster configuration,
# which includes the list of participating nodes and Solr
# shards present on every one of those nodes.
#
# NOTE: The contents of this file are being managed by
# /usr/share/cb/cbcluster command line tool and any changes
# made here may be overwritten next time that tool is used.
#
#####
#####

[Cluster]
NodeCount=2
NextSlaveAutoInc=2

[Master]
Host=172.16.100.110
HasEvents=False
User=root

[Slave1] Host=172.16.100.111
HasEvents=True
User=root
```

6. Log into CB Response.
7. In the top-right corner of the console, select your **username > Settings**:
8. In the Settings panel, select **Server Nodes** to view the server nodes in a cluster:

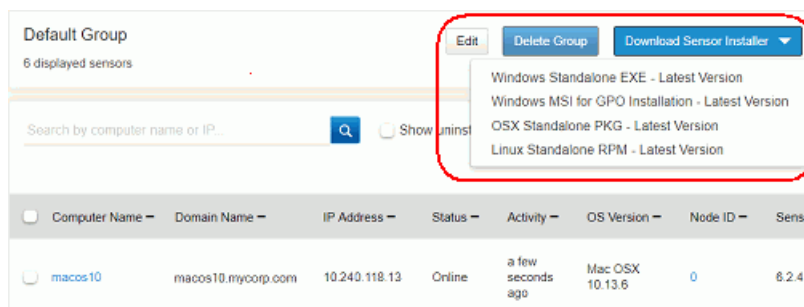


## Sensor Install and Verification

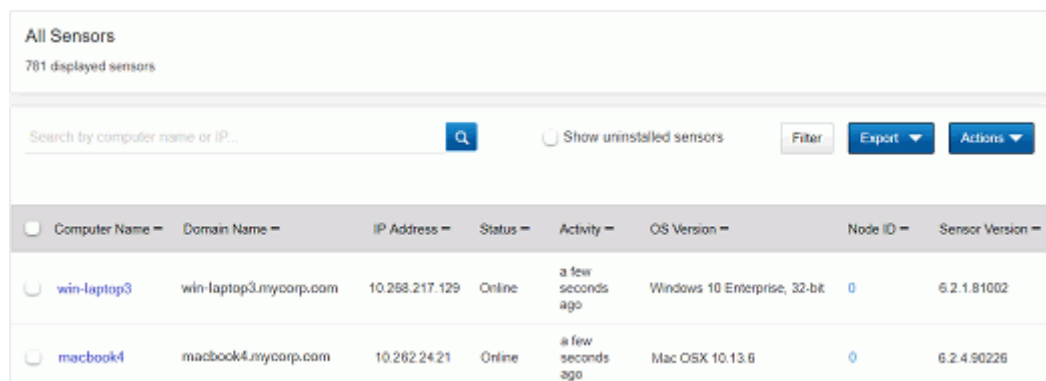
The next procedure involves adding sensors to the cluster node(s) and verifying afterward that the installed sensors appear in the CB Response console on the master node. All sensors report to the master and have an assigned node for reporting process events and binary data.

**To install sensors and verify they appear in CB Response:**

1. Log into CB Response.
2. In the left-navigation menu, select **Sensors**.
3. In the Groups list, click on the name of the group you want the new sensor to belong to. If you have not created any other groups yet, click on Default Group.
4. Select **Download Installer** and select the sensor installer that matches your OS platform and (for Windows) installation method.



5. Install the sensors on the cluster node(s). See the “Installing Sensors” chapter in the *CB Response User Guide* if you are unfamiliar with the steps for installing sensors.
6. After you have installed all sensors on the cluster node(s), verify that they appear in the CB Response console on the master node by again navigating to **Sensors** in the left navigation menu:



## Best Practices

When configuring a CB Response cluster, Carbon Black recommends that you refer to the *CB Response Operating Environment Requirements* guide, which describes performance and scalability considerations in deploying CB Response.

## Adding Minions to Existing Clusters

**To add minions to existing clusters:**

1. Verify that master and minion are on the same version of CB Response.
2. Login to the master server and stop the cluster services:

```
/usr/share/cb/cbcluster stop
```

3. Check the status to ensure the cluster services have stopped:

```
/usr/share/cb/cbcluster status
```

4. Run the cluster add node command:

```
/usr/share/cb/cbcluster add-node
```

5. On the master server, start up the cluster:

```
/usr/share/cb/cbcluster start
```

## Removing Minions from Existing Clusters

Removing unneeded minions from a cluster reduces the overall size of the cluster and simplifies the cluster deployment. This improves the cluster performance by reducing network overhead.

### Best Practices

Be sure to adhere to these best practices before attempting to remove minions from an existing cluster. Consult the *CB Response Operating Environment Requirements (OER) Guide* to ensure that the reduced cluster has the capacity to support the total number of endpoints.

- The remaining minions must have enough disk space to contain the full data.
  - Calculate the daily usage of the disk per sensor and re-calculate the required disk space to new minions for full retention. Here is an example:  
Assume you want to collapse the six-minion cluster that has 40K endpoints down to three minions and have 30-day retention. CB Response v6.0+ supports this because each minion will still have less than the maximum of 18,750 endpoints. You calculate the current daily disk usage per sensor to be 20MB (dividing total cluster data volume storage by current daily retention and number of endpoints). The new total storage requires a minimum of  $20\text{MB} * 40\text{K (endpoints)} * 30 \text{ (days)} = 24 \text{ TB}$  (or 8 TB per node).  
Add 20-30% to the available disk space on top of the calculated amount to allow for increased sensor activity in the future.
  - Additional extra disk space must be provided for desired cold storage in days.
- The master cluster node must have enough disk space to contain binary files for all removed minions. Binary files typically take less space than events but should still be taken into account. You can calculate the required extra space on the master by logging into CB Response, navigating to **Server Dashboard** in the left navigation menu, and viewing the **Storage Statistics** for the minions. Add up the total binary size on the minions that will be removed to obtain the maximum additional space required on the master node.
- The remaining minions must have enough CPU and memory space to support the sensor load (per OER).

## Read-Only Minions

Before removing minions from an existing cluster, you must convert them to a read-only state. Read-only minions are searchable throughout the API and user interface but are not used for sensor checkins or event/datastore pushes. While minions are in the read-only state, the system copies the binary files to the master.

The event data is not added to the minions and existing data is purged periodically as in normal operations. As a result, the read-only minions become completely inactive after the retention period. After the desired data retention period expires, you can safely remove the read-only minions.

Read-only minions can be reverted to active minions at any time if needed.

## Removing Minions

Perform these steps to mark/unmark minions as read-only and to remove a cluster minion. If the `node_id` parameter is required, it can be found inside `/etc/cb/cluster.conf` for a given node.

### To mark minions as read-only:

1. Stop the cluster.  

```
/usr/share/cb
```
2. For each minion that you want to remove, run the following command:  

```
/usr/share/cb/cbcluster change-node -N {node_id} -R True
```
3. (Optional) If you have an eventless master node and want it to contain events, run the following command:  

```
/usr/share/cb/cbcluster change-node -E True
```
4. Start the cluster.

### To unmark minions as read-only:

1. Stop the cluster.
2. For each minion that you want to remove, run the following command:  

```
/usr/share/cb/cbcluster change-node -N {node_id} -R False
```
3. (Optional) If you want to convert master node back to eventless, run the following command:  

```
/usr/share/cb/cbcluster change-node -E False
```
4. Start the cluster.

### To remove a cluster minion:

1. Stop the cluster.
2. For each minion that you want to remove, run the following command:  

```
/usr/share/cb/cbcluster remove-node -N {node_id}
```
3. Start the cluster.

Removed minions are no longer part of the cluster and can be shut down and de-provisioned.

**Notes**

- Only read-only minions can be removed.
- The system can remove a minion only if all storefiles have been copied to the master node. If this is not the case, an error is printed stating how many more files must be moved; you can try to remove the minion at a later time.

## Upgrading Cluster Nodes

Currently, there are two possible routes for upgrading cluster nodes:

- Running `/usr/share/cb/cbcluster start` on the master node after it has been upgraded. If there are no RPMNEW files on any cluster nodes indicating modified configuration files, this will upgrade all cluster nodes. If there are new RPMNEW files, `cbcluster start` will stop and prompt you to manually upgrade the nodes.
- Upgrading each cluster node manually with `cbupgrade`. You will need to do this in any case when there are new RPMNEW files.

**Important**

If an RPMNEW file is encountered on a node during upgrade, you must reconcile the current configuration files on the node with the new configuration information, and you must remove the RPMNEW file from that node, before `cbcluster start` will complete successfully.

**To upgrade cluster nodes:**

1. Login to the master node and stop the cluster:

```
/usr/share/cb/cbcluster stop
```

2. After the cluster is completely stopped, restart it:

```
/usr/share/cb/cbcluster start
```

If there are no RPMNEW files, the cluster nodes will be upgraded without further steps. If there are RPMNEW files, you will need to upgrade each node manually using the next procedure.

## Manually Upgrading Cluster Nodes

The procedure for manually upgrading cluster nodes is similar to the procedure for upgrading a standalone CB Response server (discussed in [“Upgrading a CB Response Server”](#) on page 31). Additional steps are also required to ensure that the cluster nodes are properly communicating with each other.



**To manually upgrade minion servers in a cluster:**

1. After the master node has been upgraded, login to each of the other machines in the cluster and upgrade `cb-enterprise`:

```
yum upgrade cb-enterprise
```

2. Login to each of the cluster nodes and upgrade their data schema to the latest version:

```
/usr/share/cb/cbupgrade [--proceed-on-rpmnew] [--non-interactive]
```

With no argument, this command stops running if CB Response `.RPMNEW` files are detected, indicating a modified configuration file, and reports the names of those files. You must manually resolve differences between the `RPMNEW` and the existing configuration file of the same name, and delete the `RPMNEW` file, before `cbupgrade` will run completely.

With the `proceed-on-rpmnew` argument, `cbupgrade` will run to completion but also reports the `RPMNEW` files it discovered so they may be addressed after upgrade. This can be helpful if you need or want a seamless upgrade and are comfortable editing the configuration files later.

Note that when the `--proceed-on-rpmnew` option is used with `--non-interactive` option, you will not be prompted or have to press any keys for the upgrade to proceed, but the `cb-enterprise` services will not be started automatically as they would if the command was run without the `non-interactive` option.

**Warning**

At the end of the `cbupgrade` process, if you are prompted to start services, do NOT start the services.

3. With a new software version, communications port requirements may change. As a result, you must check if the firewall settings need to be updated.
  - a. If you are managing the firewall yourself, run this command to identify which rules need to be added:

```
/usr/share/cb/cbcheck firewall -l
```
  - b. If you want the firewall settings to be adjusted automatically, run this command to have the utility apply those settings:

```
usr/share/cb/cbcheck firewall --apply
```
4. After all nodes have been upgraded (but not started), start the entire cluster at once by logging into the master node and issuing this command:

```
/usr/share/cb/cbcluster start
```

This step is important to perform after an upgrade. This is due to the fact that when the cluster is started using the `cbcluster` tool, it will redistribute configuration files that must be synchronized across all nodes within the cluster. If a new software version introduces configuration changes, this will ensure that each minion node has the updates.

## Chapter 6

# Using CBCLUSTER as a Non-Root User

This chapter describes how to use the CBCLUSTER command as a non-root user.

### Sections

| Topic                                    | Page |
|------------------------------------------|------|
| <a href="#">Overview</a>                 | 67   |
| <a href="#">Required User Privileges</a> | 67   |
| <a href="#">Defining Users</a>           | 69   |

## Overview

`cbcluster` is a command line utility used to manage and configure CB Response clusters on the master node. It supports the following options:

- `cbcluster start`
- `cbcluster stop`
- `cbcluster status`
- `cbcluster add-node`

These commands run in *root* context and require users to have *sudo* privileges in order to invoke `/usr/share/cb/cbcluster` on the master node. Additionally, the `cbcluster add-node` command, which is used to add a new minion node to a cluster, connects and remotely executes on the minion a series of commands for cluster configuration.

### Notes

- Please ensure that the Master and Minion are on the same version of CB Response prior to running the `cbcluster add-node` command. If the versions are not the same, you will see an error message.
- The `cbcluster reshard` command is deprecated and no longer supported.

As of CB Response version 5.1.1, it is possible to define a non-root user as the remote user for minion communication and execution. Previously, when adding a minion node to a cluster, the `cbcluster` utility required availability of root user on the minion node. In this release, this requirement has been relaxed, and minion nodes can be configured as a non-root user.

This chapter describes the required privileges for the non-root user and its use during cluster setup.

## Required User Privileges

Before invoking `cbcluster` to connect to a minion as a non-root user, the remote user on the minion needs to have certain assigned privileges:

- SSH access to the minion node
- Sudo privileges for the commands listed below; the user **MUST** be configured to run with `NOPASSWD`.

To use the `cbcluster` commands with a non-root user, add the following entries to your `sudoers` file:

```

## Required sudo privileges on minion to run cbcluster add-node
Cmnd_Alias HOSTNAME = /bin/hostname
Cmnd_Alias CB_INIT = /usr/share/cb/cbinit
Cmnd_Alias YUM_INSTALL_CB = /usr/bin/yum install cb-enterprise -y
Cmnd_Alias YUM_INSTALL_RSYNC = /usr/bin/yum install rsync -y
Cmnd_Alias MKDIR_ETC_CB = /bin/mkdir /etc/cb --mode=755
Cmnd_Alias MKDIR_ETC_CB_CERTS = /bin/mkdir /etc/cb/certs --mode=755
Cmnd_Alias COPY_ALLIANCE_CERT = /usr/bin/rsync --remove-source-files -
-verbose /tmp/.cb_tmp/carbonblack-alliance-client.crt /etc/cb/certs/
carbonblack-alliance-client.crt
Cmnd_Alias COPY_SERVER_CERT = /usr/bin/rsync --remove-source-files --
verbose /tmp/.cb_tmp/cb-server.crt /etc/cb/certs/cb-server.crt
Cmnd_Alias COPY_CLIENT_CA_CERT = /usr/bin/rsync --remove-source-files
--verbose /tmp/.cb_tmp/cb-client-ca.crt /etc/cb/certs/cb-client-
ca.crt
Cmnd_Alias COPY_ALLIANCE_KEY = /usr/bin/rsync --remove-source-files -
-verbose /tmp/.cb_tmp/carbonblack-alliance-client.key /etc/cb/certs/
carbonblack-alliance-client.key
Cmnd_Alias COPY_SERVER_KEY = /usr/bin/rsync --remove-source-files --
verbose /tmp/.cb_tmp/cb-server.key /etc/cb/certs/cb-server.key
Cmnd_Alias COPY_CLIENT_CA_KEY = /usr/bin/rsync --remove-source-files
--verbose /tmp/.cb_tmp/cb-client-ca.key /etc/cb/certs/cb-client-
ca.key
Cmnd_Alias COPY_CB_REPO = /usr/bin/rsync --remove-source-files --
verbose /tmp/.cb_tmp/CarbonBlack.repo /etc/yum.repos.d/
CarbonBlack.repo
Cmnd_Alias COPY_CLUSTER_CONF = /usr/bin/rsync --remove-source-files -
-verbose /tmp/.cb_tmp/cluster.conf /etc/cb/cluster.conf
Cmnd_Alias COPY_ERLANG_COOKIE = /usr/bin/rsync --remove-source-files
--verbose /tmp/.cb_tmp/erlang.cookie /var/cb/.erlang.cookie
Cmnd_Alias COPY_SERVER_LIC = /usr/bin/rsync --remove-source-files --
verbose /tmp/.cb_tmp/server.lic /etc/cb/server.lic
Cmnd_Alias COPY_SERVER_TOKEN = /usr/bin/rsync --remove-source-files -
-verbose /tmp/.cb_tmp/server.token /etc/cb/server.token
Cmnd_Alias CBCHECK_FIREWALL = /usr/share/cb/cbcheck firewall --apply
Cmnd_Alias CB_ENTERPRISE = /etc/init.d/cb-enterprise
Cmnd_Alias CAT_VERSION = /bin/cat /usr/share/cb/VERSION
Cmnd_Alias CBUPGRADE = /usr/share/cb/cbupgrade --non-interactive
Cmnd_Alias CBUPGRADE_CHECK = /usr/share/cb/cbupgrade --check

my user ALL=(ALL) NOPASSWD: HOSTNAME, CB_INIT, YUM_INSTALL_CB,
YUM_INSTALL_RSYNC, MKDIR_ETC_CB, MKDIR_ETC_CB_CERTS,
COPY_ALLIANCE_CERT, COPY_SERVER_CERT, COPY_CLIENT_CA_CERT,
COPY_ALLIANCE_KEY, COPY_SERVER_KEY, COPY_CLIENT_CA_KEY, COPY_CB_REPO,
COPY_CLUSTER_CONF, COPY_ERLANG_COOKIE, COPY_SERVER_LIC,
COPY_SERVER_TOKEN, CBCHECK_FIREWALL, CB_ENTERPRISE, CAT_VERSION,
CBUPGRADE, CBUPGRADE_CHECK

```

See <https://community.carbonblack.com/docs/DOC-5692> for a version of these entries formatted so that you can copy and paste them into your environment.

If any of the required permissions are not configured, the `cbcluster` command will prompt for the missing permissions during initial validation.

## Defining Users

### Note

To avoid potential conflicts, do not select “cb” for a username.

There are two ways to invoke `cbcluster` and configure minion nodes as a user other than root:

- If this is a new minion which has not been added to the cluster, the new user can be configured when running `add-node`, by including the new `add-node` option:

```
--user=<arg>
```

For example:

```
$ /usr/share/cb/cbcluster add-node --hostname <my_host> --user  
<my_user>
```

This will configure `cbcluster` to perform all remote actions for `<my_host>` as `<my_user>`. The `/etc/cb/cluster.conf` file will be updated to reflect the new configuration with the key-value pair `"User=<my_user>"`.

- If this is an existing minion, the `/etc/cb/cluster.conf` file can be modified directly. Open the file in an editor, and for each minion node, add the key-value pair `"User=<my_user>"` where `<my_user>` is the actual username you want `cbcluster` to use when connecting to that minion.