

Operating Environment Requirements (OER)

Versions 6.2.3 & 6.2.4, November 2018

Contents

Document Change Log	2
Overview	3
Executive Summary	3
Cb Response Architecture and Sizing	4
Cb Response Deployment Dimensions	4
Estimating Endpoint Activity.....	5
Determining Desired Performance.....	6
Determining Data Retention.....	6
Disk Storage Requirements	7
Disk Space Requirements for Non-Data Drives	7
Hard Disk Performance	8
Determining Maximum Endpoints per Server	8
Deployment Planning and Sizing Summary.....	9
Sizing the Installation	9
Sizing the Server	9
Sizing the Cluster.....	9
Multiple Cluster Environments	10
Virtual Server Deployments	10
Virtual Deployment Considerations	10
Adjusting Server Requirements for Other Variables	11
Cb Response Server Disk Configuration	11
Overview	11
Disk Configuration	11
Appendix A: FAQs	12
1. What Sensor OSEs does Cb Response support?	12
2. What Cb Response console/server operating systems are supported?.....	12
3. What is the sensor impact on the endpoint?	12
4. How much network bandwidth does Cb Response require?.....	12
5. I have remote locations and/or users who travel. How does Cb Response work for those users? Do I need to consider impact on server sizing and configuration?	13
6. If I do not have sufficient server resources, what is the impact?	14
7. Is there a way to avoid storing data for high volume processes?.....	14
8. What is the recommended number of days to retain the data and why?	14
9. What are the concerns if I want to store more than the recommended number of days' worth of data?.....	15
10. Are virtual servers supported?	15
11. Do the sensors support VDI environments?.....	15
12. How does Cb Response support disaster recovery (DR), high availability (HA) & backups?	15
13. I have fewer than 18,750 endpoints. Can I get by with fewer server resources/hardware?	15

Copyright ©2011-2018 Carbon Black, Inc. All rights reserved. This product may be covered under one or more patents pending. Cb Response is a registered trademark of Carbon Black, Inc. in the United States and other countries. Any other trademarks and product names used herein may be the trademarks of their respective owners.

Carbon Black.

14. Do Server requirements change depending on the sensor type (Windows, Mac, or Linux)?	16
15. Can agent-based software (antivirus, performance monitoring, backup utilities) be installed on Cb Response server?	16
16. How do I contact Carbon Black Support?	16

Document Change Log

Date	Revision
10/19/2016	Draft
12/20/2016	Updated endpoint limits to 18,750 per node
04/05/2017	Remove Draft notation
05/10/2017	GA 6.1 release revision
05/25/2017	Minor adjustments to non-data disk size requirements for consistency
11/07/2018	Updated supported Linux versions for the server. Reformatted the document. Other minor editing.

Overview

This document provides insight into the performance and scalability considerations in deploying Carbon Black (Cb) Response and provides the following:

- An overview of how Cb Response operates and the underlying need to provide the resources required for a successful implementation.
- Descriptions of the dimensions of performance and scale on Cb Response and recommendations for server and cluster sizing.
- An overview covering the most common questions and answers about the performance and scale of Cb Response within an enterprise.
- Recommended hardware for reference sizes of installations, enabling a cost estimate for the hardware to meet the scalability, and performance and storage needs of customers at specific install sizes.

In addition to this document, the following additional resources exist to assist to help properly size a Cb Response installation:

- **Cb Response Sizing Estimator:** Providing interactive sizing estimates, the Cb Response sizing calculator can be used in conjunction with a Cb Response sales engineer to understand the rough sizing requirements based on key user inputs.
- **Professional Services:** Cb Response customers may also engage customer support or professional services for additional assistance in scoping, sizing, and tuning their installations.

Executive Summary

Continuous recording and analysis of the activity on endpoints is required to detect and respond to today's complex threat landscape. This collection and analysis of endpoint information allows organizations to more quickly detect, respond to, and properly remediate incidents.

Cb Response is a continuous real-time endpoint monitoring, collection, processing, and analytics solution that manages very large amounts of data and demands a somewhat unique hardware infrastructure. In short, Cb Response is a big data solution and is similar to netflow or data aggregation products in terms of function and processing demands. Cb Response is unlike a typical database-driven web app. Standard databases sizes grow beyond 10Tb, and Cb Response often processes and analyzes billions of data points of information a day.

Insufficient or inappropriate hardware configurations account for the majority of performance-related issues encountered by Cb Response customers. A properly-configured system ensures Cb Response delivers the highest-possible user experience. It is for these reasons that we require conformance to the Server Sizing Guide.

This document has been carefully designed with our customers' success as the top priority. We make a concerted effort to ensure that this guide receives the appropriate attention by all stakeholders, including those from IT, SecOps, database management, and datacenter teams. Alignment across stakeholders helps ensure an on-time deployment and minimal time-to-value.

This document is designed to guide you to the necessary hardware and storage configurations to provide a great experience and put your organization in the best posture possible. We thank you for taking the time to review and understand this guide, and we look forward to working with you to design an infrastructure that meets your specific needs.

Cb Response Architecture and Sizing

Cb Response consists of two main components: sensors, which reside on and monitor the endpoints, and the centralized server infrastructure, which stores the sensor data and serves the Cb Response console. The centralized server infrastructure, can be one server or multiple servers in a cluster for larger deployments.

- Cb Response can support up to 18,750 sensors per server.
- Up to eight-servers, plus one head node, can be grouped in a cluster under a single user console to support up to 150,000 sensors per cluster.
- The number of sensors supported and the duration of sensor data stored are mostly driven by the number and activity of processes launched by each endpoint. Endpoints vary widely in the volume of processes generated, based on the software running on the endpoints. While there are other factors that drive scale, this is the most important factor.
- The Cb Response data store is Apache Solr for events with a Postgres management database.
- Proper sizing of server infrastructure to support a high-performance installation for each installation is critical to successful implementations.

Cb Response Deployment Dimensions

There are a number of different dimensions of data that can impact performance and scalability of a Cb Response server install. This section provides an overview of the most critical components that drive scale.

- **Incoming data rates:** Each sensor sends a stream of data to the server; that data requires indexing and storage. The incoming data rate generated for an installation is impacted by three measures:
 - The number of concurrently active sensors: The server must process the event data from each endpoint. More sensors means higher CPU and disk IO requirements.
 - The number of processes per sensor per day: Planning requires estimates for the typical rate of processes per host per day, but there can be wide variation between installations. Operating systems may also vary widely in the number of processes per generated sensor.
 - The average activity per process: Similar to the count of processes per host per day, the activity of those processes (such as file modifications, registry changes, network connections, and created child processes - among others) can vary between hosts. The amount of storage each process document takes up is dependent on the activity of the process.
- **Notes:**
 - Cb Response uses a storage model based on per-process storage where the activities of a given process are stored within set of per-process containers. These are referred to as a *process documents* and each represents snapshot of the process activity in period of time (typically 5 minutes).
 - Each process can have one or more process documents associated with it. New process documents are created each time a sensor sends more data to the server. Therefore, longer-living processes (such as svchost.exe) can have many documents associated with them. Average observed process to document ratio is between 2 and 3.
 - The process document count is a key driver of performance and storage capacity.

Carbon Black.

- **Threat intelligence feeds:** Depending on the number of feeds enabled, your server will be monitoring for activity related to a number of unique indicators. While most organizations won't notice the impact, some organizations may wish to monitor a very large number of indicators, requiring additional resources.
- **Watchlists:** For each watchlist configured, your server will run a search every ten minutes. Performance can be impacted by the number and complexity of those searches. Most organizations won't notice the impact, but some organizations may want to monitor many watchlists or process complex watchlists, and this will impact performance (e.g. lengthy watchlists or those containing wildcards).

Estimating Endpoint Activity

Because the activity of the sensors can vary greatly across different software ecosystems, it can be challenging to estimate the size of data generated for the server. This document provides three estimates that will help gauge the required server specifications below.

In this example, we will provide three example activity rates used in making estimations on server performance. The majority of customers will fall into the medium activity estimate and those that require to hit the desired days of data stored should use the conservative estimates:

- **Low Activity Estimate:** This estimate assumes lower sensor and process activity rate. This estimate is more accurate for Windows only environments with non-developer systems.
- **Medium Activity Estimate:** This estimate assumes average sensor activity rate. This estimate will be more accurate for the full enterprise deployment, including a smaller subset of Linux and Mac's sensors.
- **High Activity Estimate:** This estimate assumes a higher-than-average sensor activity level. This estimate represents small percentage all Cb Response installs and should not be considered typical unless deployment consists of large percentage of Linux and Mac sensors.

Table 1: Sample Processes per Sensor per Day and Storage Size Estimates

Installation Estimate Type	Processes / Sensor / Day	Storage Size / Sensor / Day (at 10 KB process size)
Low Activity Estimate	1,600	16 MB
Medium Activity Estimate	3,200	32 MB
High Activity Estimate	6,400	64 MB

Notes:

1. Storage size /day assumes that individual process size (in Solr) is 10 KB. This is average value we observe on customers today, but can differ in some circumstances. Storage size per day will be affected by this difference
2. If you change the **Data Suppression Level** in the **Create Group/Edit Group Settings** dialogs (discussed in the *Carbon Black Response User Guide*) from the default of **Medium** to **High** or **None**, the estimates in *Table 1* will change. Selecting the **High** data suppression level results in a reduction to these numbers. Selecting **None** results in an increase to these numbers.

Determining Desired Performance

Several desired performance parameters must be taken into account when choosing a set of server specifications. These areas include:

- **Data Retention:** This determines how long the Cb Response server will store the information generated by each sensor. The longer the data is stored, the more resources must be available on the server. This is primarily driven by the performance of the disk subsystem, and the amount of hard drive storage. Available RAM and CPU determine performance of searches when done across entire retention history
- **Sensors per Server:** Several factors determine the scalability of the server in terms of the number of sensors it can handle. While some of this is determined by the activity rate of the sensor, increasing the CPU and hard drive IO performance can help increase the number of sensors a server can handle.

Determining Data Retention

As mentioned before, Cb Response server stores each instance of a process execution and all event data associated with it (e.g. module loads, registry or file modification, and network connections) in a set of “process documents”.

The following chart provides conservative storage requirements for each day of storage given number of sensors.

Table 2: Daily storage requirements for different activity estimates

Sensors Per Server	Low Activity Estimate	Medium Activity Estimate	High Activity Estimate
1,000	16 GB	32 GB	64 GB
3,125	50 GB	100 GB	200 GB
6,250	100 GB	200 GB	400 GB
18,750	300 GB	600 GB	1.2 TB

Notes:

- Once maximum storage capacity is reached (90% of storage disk size), Cb Response purges the oldest data.
- If you change the **Data Suppression Level** in the **Create Group/Edit Group Settings** dialogs (discussed in the *Carbon Black Response User Guide*) from the default of **Medium** to **High** or **None**, the estimates in Table 2 will change. Selecting the **High** data suppression level results in a reduction to these numbers. Selecting **None** results in an increase to these numbers.

Disk Storage Requirements

Cb Response server requires free space for regular operation of the data-store in addition to the size of the retained data. We recommend maintaining 25% - 30% extra free space for performance. Based on the estimated activity levels in *Table 2*, the following provides information on the hard disk storage requirements to maintain the 30 day retention rate for 18,750 endpoint server:

*Table 3: Recommended Hard Disk Capacity per Server at 18,750 sensors/server***

Activity Estimate	Low Activity Estimate	Medium Activity Estimate	High Activity Estimate
Data Partition Size	12 TB	24 TB	48 TB
**These values do not include disk space requirements for non-data drives			

The next table shows expected disk writes for given activity estimates. This is important when determining required write performance of the disk. Note that disk writes can vary based on the nature of the activity reported by sensor, and the values provided here are only estimates.

Table 4: Disk write MB/s for different activity estimates

Sensors Per Server	Low Activity Estimate	Medium Activity Estimate	High Activity Estimate
1,000	1.6 MB/s	3.2 MB/s	6.4 MB/s
3,125	5 MB/s	10 MB/s	20 MB/s
6,250	10 MB/s	20 MB/s	40 MB/s
18,750	30 MB/s	60 MB/s	120 MB/s

Notes:

- Disk usage could be reduced in case some portion of sensors is not always active or using ingress filters,
- If insufficient disk space is provided, system will automatically delete oldest data, effectively reducing total days of retention,
- There are additional storage requirements for non-process related activity, such as management configuration storage and the storage of binaries, which is assumed to be <150 Gigs of storage and are included in the figures above. System will automatically delete binaries that have been sent to Alliance, but if the binary sharing is turned off, binaries will be kept and consume storage permanently.

Disk Space Requirements for Non-Data Drives

In addition to 10 GB for root drive hosting OS, we recommend reserving 75GB for /tmp and space equal to 70% of total RAM for logging/diagnostics and process memory dumps (if needed). Here is an example partition scheme for machine with 128 GB of RAM:

- 10 GB free for root / drive - OS files and installed applications
- 75 GB free for “/tmp” directory - diagnostic files
- 90 GB for “/var/log/cb” (equal to 70% of total RAM on the server) - logs and memory dumps

Carbon Black.

This partitioning scheme will ensure that the OS will remain responsive and not run out of free space due to memory dumps, logging, diagnostic data, and so on. The partition scheme is only a recommendation. If non-data space is on a single volume, it should equal 75GB + 70% RAM for total allocated space.

Note: Additional disk space can be added at customer discretion (not required) to enable greater retention of logging and diagnostic data.

Hard Disk Performance

While document retention is primarily limited by total available disk storage, system throughput is determined by disk performance (IOPS and throughput). We recommend using high-performance 15k spinning drives (RAID 10) for smaller deployments (up to few thousand endpoints) and SSD drives for larger deployments. See [Disk Configuration](#) for more details.

Determining Maximum Endpoints per Server

A typical Cb Response server node has sustained ingest rate of about **1,600 process documents/second**, corresponding to a maximum ingest of approx 140M process documents a day, or approx 60M sensor processes per day. This rate is sufficient to support a typical enterprise deployment of 18,750 endpoints per server node for low and medium activity estimate loads. If endpoints generate different number than 3,200 processes per day on average, then the number of supported endpoints per server node needs to be adjusted correspondingly:

Notes:

- Ingest rate primarily depends on the read/write rates of the disk subsystem. Network-based disk subsystems or low RPM spinning disks may negatively impact server performance. Solid State Drives (SSDs) are recommended for best performance and will result in above mentioned 1,600 process documents/second.
- If endpoints are generating too many processes per host per day on average, the Cb Response Professional services team can help you deploy techniques to suppress known-good processes.
- 1,600 process documents/second per node capacity applies to environments consisting of 6.x sensors. In case there are older sensors (5.x) connected to the server, this number will be reduced by 33% due to extra processing overhead when translating older sensor data and associated disk activity. Therefore, for medium activity load, server node with old (5.x) sensors will support up to 12,500 sensors
- Number of sensors per node will not exceed 18,750, and number of sensors per cluster will not exceed 150,000 regardless of the activity level of sensors

Table 5: Calculating max number of sensors per server node

Activity Estimate/Sensors Per Server Node	Low Activity Estimate (1,600 processes/day)	Medium Activity Estimate (3,200 processes/day)	High Activity Estimate (6,400 processes/day)
Max 6.x sensors per server node	18,750	18,750	9,375
Max 5.x sensors per server node	18,750	12,500	6,250

Deployment Planning and Sizing Summary

In summary, the key factors in planning your deployment are driven by endpoint activity for the input to the requirements and key factors within the server infrastructure. This assures proper performance as shown in the following table:

Table 6: Key factors in deployment planning and sizing

Areas	Key Factor
Endpoints	<ul style="list-style-type: none"> • Number of processes generated: driven by number of endpoints, OS type, and activity levels of endpoints • Size of process documents: Driven by the activity of the endpoints
Servers	<ul style="list-style-type: none"> • CPU • RAM • Disk Performance • Disk Size

Sizing the Installation

There are two phases when sizing a Cb Response installation:

- Size the server or server cluster
- Determine best hardware for the server(s)

Sizing the Server

For up to 18,750 **Windows** endpoints, and for servers within a cluster, we recommend the following hardware configuration. Note that this assumes **low sensor activity** of 1,600 processes per day per sensor. For higher sensor activity estimates, storage will need to be increased accordingly.

*Table 7: Server sizing***

Endpoints	Cluster Configuration
< 18,750	Single server; 16 core @ 2.5 GHz, 128 GB RAM, at least 15k disks in RAID 10 or SSDs (all hard drives), 12 TB storage for 30 days of data retention and 175 GB for non-data drives (OS, /tmp/, etc.)
** See the FAQs for details if reduction in the server sizing is preferred for smaller deployments.	

Sizing the Cluster

For installations with **more than 18,750 endpoints**, a cluster (more than one server) is required. In general, we recommend:

- One server per 18,750 endpoints; each cluster up to 150,000 endpoints.
- One additional dedicated 'master' server for installations over 37,500 endpoints.
- More than 150,000 endpoints requires multiple Cb Response clusters.
- All nodes within the server cluster must be on the same network segment and switch, including the head end.
- Clustered servers can reside in a single chassis for efficiencies associated with hosting multiple servers (i.e., space, cooling, power).

Table 8: Cluster Configuration by Endpoint Count

Endpoints	Cluster Configuration
< 18,750	Single server
18,750-37,500	Two servers, both indexers
37,500 - 150,000	One server per 18,750 endpoints , plus one dedicated “master” node
> 150,000	Multiple clusters

If the installation has **more than 37,500 endpoints**, we recommend that that dedicated master node be configured with the following configuration:

- 8 core @ 2.5 GHz, 96 GB RAM, with at least 1TB of storage and 143GB for non-data drives (OS, /tmp/, etc.) using 15k disks or SSDs.

Multiple Cluster Environments

Sites with **more than 150,000 endpoints** require multiple clusters. Cb Response enables enterprise-wide management of multiple clusters via four primary subsystems:

- Custom Cb Threat Intel Feeds: In-house threat intelligence can be syndicated over the network to all clusters from a single location.
- REST API: Provides simple REST endpoints for searching and managing the other cluster configuration details. Scripts for common tasks are available from our support staff.
- Enterprise Messaging Bus: Provides the ability to subscribe to event streams.
- Syslog: Provides the ability to combine and forward (alerting) information from multiple clusters into a central location (such as a SIEM).

Multi-cluster solutions also mitigate bandwidth concerns where endpoints are geographically dispersed. Network bandwidth loads are constrained to local, higher-speed links found within local area networks, and only API calls, alliance communications, and queries are sent over the more constrained wide area network.

Virtual Server Deployments

Virtual deployments are supported, as long as the hardware specifications are met and the resources required are available to Cb Response. Virtual environments create economies of scale by time-sharing hardware resources between multiple machines. This can create performance bottlenecks, particularly in the disk subsystem, if those resources are oversubscribed. Hardware resources should be dedicated or reserved to mitigate these risks.

Virtual Deployment Considerations

- Dedicated CPU, RAM, and IOPs resources to be given to the Cb Response virtual machine.
- Hardware specs/recommendations remain the same as covered above (# processor, # RAM, disk allocation).
 - o Disk and SAN considerations:
 - o Data partition (defaults to /var/cb/data directory) should be as fast as possible and the preference is 15K disk or SSD.

Carbon Black.

- If you choose to use a SAN, we recommend placing data partition (defaults to /var/cb/data directory) on its own dedicated LUN.
- Be aware of resource throttles that may restrict performance. Throttling within a virtual machine is documented by VMware here.

Adjusting Server Requirements for Other Variables

The above server requirements are based only on the number of endpoints. Two other key variables exist: endpoint activity and data retention period. To explore the impact of adjusting those values, review the output of the Cb Response sizing calculator with Cb Response sales engineering to estimate the number of servers and approximate storage required.

Cb Response Server Disk Configuration

Overview

Cb Response is a very I/O intensive application. As a result, high performance disks and RAID configurations are required. This section provides the recommended disk configuration for each type of disk partition.

Disk Configuration

- **Data Disk** – The data disk partition, as mentioned above, will store the primary data that requires the heavy requirements. Note that Cb Response places its high-volume data to the /var/cb/data directory by default. Therefore, the data volume should be mounted to this directory if you plan to keep the default configuration.
- For partitions requiring 2TB of storage space or more, the following specifications are required:
 - **Option 1, SSD** – At least five solid-state SAS drives in a RAID5 configuration.
 - **Option 2, Spinning Disk** – At least eight 6Gb/s SAS 15K RPM drives in a RAID1+0 and at least 1GB of battery-backed write cache.

As the endpoints, activity levels and subsequent disk space requirements reduce, other options can apply. For installations with less than 1TB of data, you can reduce option 1 to four SSDs or option 2 to six spinning disks.

Additional notes:

- Disk needs to be qualified using the **CbR Qualifier tool**. Work with Carbon Black staff to qualify your disk to ensure enough throughput for your deployment size.
- In case a single disk array doesn't have sufficient throughput, you might need to provide two equally sized volumes that will alternate in storing and optimizing event partitions
- **Non-Data Disk** – The non-data disks partitions do not require the same I/O and space requirements as the data disks. As a result, the non-data partitions should have at least one RAID1 partition across two spinning or two SSD disks.

Appendix A: FAQs

This FAQ covers the most common questions and answers to Cb Response deployments.

1. What Sensor OSES does Cb Response support?

- Please refer to the latest sensor OS release support information here: <https://community.carbonblack.com/docs/DOC-7991>

2. What Cb Response console/server operating systems are supported?

- Linux
 - o Red Hat Enterprise Linux (RHEL) 6.7 - 6.10 (64-bit)
 - o Red Hat Enterprise Linux (RHEL) 7.3 - 7.5 (64-bit)
 - o CentOS 6.7-6.10 (64-bit)
 - o CentOS 7.3-7.5 (64-bit)

Installation and testing is done on default installs using the 'minimal' distribution and the distribution's official package repositories. Customized Linux installations must be individually evaluated.

3. What is the sensor impact on the endpoint?

- The Cb Response sensor is designed to have no performance impact. The expected impact of Cb Response on a the endpoint are:
 - o CPU: ~1% CPU usage. This can vary depending on system activity.
 - o Memory: 12-20MB RAM
- Disk Storage: This is minimal. When the sensor can communicate with the server, the sensor stores data on the endpoint and regularly sends the data to the server. If the sensor cannot communicate with the server, data will queue up to an adjustable threshold (2GB by default, expected 30-60 days activity on a normal system). The data will be synced upon re-establishment of server communications.

4. How much network bandwidth does Cb Response require?

It is difficult to predict actual network traffic that Cb response requires. It depends on many factors, including activity estimate of the sensor and number of unique binary files seen and uploaded to the server. However, here are some estimates:

- Per endpoint:
 - o 1-4 kilobits per second (kbps) per host
 - o 10-40 megabytes (MB) per host, per day
- The following table shows server-side expected average network traffic based on sensor activity estimates:

Table 9: Average expected server network incoming traffic in MB/s (megabytes per second)

Activity Estimate/Sensors Per Server	Low Activity Estimate	Medium Activity Estimate	High Activity Estimate
1,000	0.12 MB/s	0.24 MB/s	0.5 MB/s
3,125	0.4 MB/s	0.75 MB/s	1.5 MB/s
6,250	0.75 MB/s	1.5 MB/s	3 MB/s
18,750	2.25 MB/s	4.5 MB/s	9 MB/s

- Throttling can be configured per site via sensor groups, per hour, per day.
 - Throttling will limit bandwidth from a group of endpoint sensors. Often used on low-bandwidth sites or sites that are bandwidth constrained at certain times of the day.
 - The trade-off when throttling is invoked is a delay in data sent back to the central server for analysis against watchlists and the availability of the data in the console.
 - Console users can override the network throttle by enabling “sync” to any individual host to instruct the host to ignore any configured throttles and send all data immediately.
 - Throttles shape the volume of traffic to the server from sensors at particular times. They do not reduce overall traffic. To reduce traffic, you can reduce data collected on the sensor group’s configuration.
- Maximum sensor Checkin rate can be configured through SensorCheckingDelayRate in cb.conf.
 - Default value is 100 and it corresponds to max 100 checkins/second/server node
 - Reducing this value, reduces network traffic due to checkins, but also reduces how often sensors send statistics and retrieve any configuration changes

Note: Due to the number of processes generated on those endpoints, Mac and Linux sensors may drive higher bandwidth utilization.

5. I have remote locations and/or users who travel. How does Cb Response work for those users? Do I need to consider impact on server sizing and configuration?

- If machines at remote locations (for example, outside the corporate network) can reach the Cb Response server, all operations are identical to when the endpoint is within the network.
- While not connected to the server, Cb Response sensors will queue data on the endpoint (up to a configurable threshold) until the server is reachable again.
 - Default storage on the endpoint is 2GB or 2% of total disk storage (whichever is reached first). This should be enough for multiple months of data. The default is configurable by sensor group.
 - Once the local data storage limit is reached, the sensor stops storing new log messages.

Carbon Black.

- Customers can also deploy the Cb Response server in their DMZ or directly on the Internet.
- For installations in a DMZ or with direct internet access, Cb Response can be configured to restrict access to the management interface (i.e., the GUI) to a separate, internal network interface.
- This behavior does not impact server sizing.

6. If I do not have sufficient server resources, what is the impact?

If there are insufficient resources, the server will take action, throttling the sensors in their sending of data. If the CPU resources or storage performance is insufficient for the number of process documents stored, the search performance of Cb Response will suffer.

If the server is throttling sensor uploads, the data will queue up at the endpoints until one of two things occurs: the server can handle the load or the sensors hit their local threshold for how much data to queue. Endpoints will continue to operate for the end user as they usually do without any noticeable performance impact.

7. Is there a way to avoid storing data for high volume processes?

It is possible to deploy event filters on the server that can be used to limit particularly noisy processes. This mechanism should only be used in extreme cases where a large amount of process activity can be attributed to a few processes. We recommend that you engage with professional services regarding the ability to enable event filtering.

It is also possible to reduce the data recorded per-process on a per-sensor-group basis. In the sensor group's configuration page, you can limit the collection of:

- Process information
- File modifications
- Registry modifications
- Binary module (.dll, .sys, .exe) loads
- Network connections
- Binaries
- Binary info
- Process user context
- Non-binary file writes
- Cross process events

Eliminating these events will reduce data volumes, but at a loss of fidelity. In a typical enterprise, the most frequent events are file modifications, registry modifications, and binary module loads. The performance impact of disabling specific event types is highly variable and depends on the endpoint environment. Engaging with sales engineering or Professional Services can help you determine the best path to take.

8. What is the recommended number of days to retain the data and why?

25-35 days: The number of days for which you retain the data is a balance between the value of the data over time and the cost of data storage. 25-35 days allows for access to key information required for incident response and limits the amount of storage that is required. Additional capacity can be added for customers who prefer more history.

9. What are the concerns if I want to store more than the recommended number of days' worth of data?

Days stored will increase the amount of data storage required and may impact the number of servers required to process the stored data. There is a performance degradation that occurs when the number of process documents stored exceeds the recommended number of days stored. Use the Sizing Estimator in conjunction with sales engineering or Professional Services for tailored guidance.

10. Are virtual servers supported?

Yes. See the considerations for Virtual Server deployments above.

11. Do the sensors support VDI environments?

Yes. Sensors running on a Virtual Desktop Infrastructure (VDI) are supported for both persistent and non-persistent VDI setups. To ensure continuity through non-persistent sessions, Cb Response has developed logic in the sensor to ensure that the VDI session maintains the sensor ID. This ensures that each sensor is only depicted once in the console. For VDI, Cb Response has limited the amount of disk writes for both persistent and non-persistent sessions to ensure that sessions are optimized for zero or thin sessions.

12. How does Cb Response support disaster recovery (DR), high availability (HA) & backups?

- Cb Response relies on existing system administration procedures for backup and recovery.
 - Best practices when using virtual infrastructure include taking snapshots and backing up using existing procedures
 - Best practices on physical hardware use RAID to help against hard drive failure and maintain cold spares in the event of other hardware failure.
- Typical Linux tools can (and should) be used to make backups of your certificates, configuration database, and settings to assist in DR and HA.
- Sensors store data locally (if they cannot connect to the server) and transmit the data upon the connection being re-established. The amount of data stored on the endpoint can be configured within the Cb Response console.
- For backup, sending critical Cb Response events to SIEMs is a best practice.
- You can backup and archive the Cb Response event data, but this requires custom scripting and can be large.

13. I have fewer than 18,750 endpoints. Can I get by with fewer server resources/hardware?

Yes. It is possible to use fewer resources for smaller deployments. Some general guidelines for Windows endpoints are as follows. Note that this assumes medium sensor activity of 3,200 processes per day per sensor.

Table 10: Server Sizing

Windows Endpoints	Cluster Configuration
<1,000	2 core @ 2.5GHz, 24GB RAM, 1.1TB storage + 102 GB for non-data drives (OS, /tmp/, etc.)
1,000 - 2,500	4 core @ 2.5GHz, 32GB RAM, 2.7TB storage + 102 GB for non-data drives (OS, /tmp/, etc.)
2,500-5,000	6 core @ 2.5GHz, 64GB RAM, 5.5TB + 130 GB for non-data drives (OS, /tmp/, etc.)
5,000-10,000+	12 core @ 2.5GHz, 96GB RAM, 11TB storage + 152 GB for non-data drives (OS, /tmp/, etc.)
10,000-18,750	16 core @ 2.5GHz, 128GB RAM, 20TB storage + 175GB for non-data drives (OS, /tmp/, etc.)

Notes on Table 10:

1. Data volume sizing in the table above assumes medium activity estimate for the sensor, and upper boundary of the endpoint range. Please use *Table 2* in order to estimate required data volume size more precisely, depending on the sensor activity estimate.
2. RAM specified is minimum required, regardless of activity estimate. Providing more RAM than specified in this table will help performance to some extent since it will allow OS to cache disk data more efficiently, but improvement will be less than linear.
3. On the < 1,000 tier (due to smaller scale and lower total RAM), a small set of sensors operating outside “normal” behavior could have a disproportionate impact on the system as a whole (e.g. performance/stability impacts). For that reason, specified RAM is out of proportion compared to larger tiers.

14. Do Server requirements change depending on the sensor type (Windows, Mac, or Linux)?

Yes. Mac and Linux enterprise machines create a higher number of processes than a Windows machine. This higher volume of processes per endpoint should be taken into consideration while planning and deploying a Cb Response installation.

An exception might occur in Windows but is more common with Mac and Linux. Endpoints running certain applications can be noisy, generating a high volume of processes. These endpoints, and the reason for their high activity, should be discovered as part of a systematic roll out. High-activity endpoints include those that generate more than 10,000 processes per day.

15. Can agent-based software (antivirus, performance monitoring, backup utilities) be installed on Cb Response server?

Yes, configure agent-based software according to industry best practices and Carbon Black best practices.

16. How do I contact Carbon Black Support?

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

- Web: www.carbonblack.com
- E-mail: support@carbonblack.com
- Phone: 877.248.9098
- Fax: 617.393.7499