



Carbon Black Response 5.2

Operating Environment Requirements

Version 5.2
October 2016

Carbon Black, Inc.
1100 Winter Street, Waltham, MA 02451 USA
Tel: 617.393.7400 Fax: 617.393.7499
E-mail: support@carbonblack.com
Web: <http://www.carbonblack.com>

Copyright © 2004-2016 Carbon Black, Inc. All rights reserved. This product may be covered under one or more patents pending. Carbon Black is a 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.

Document Change Log

Date	Revision

Contents

- Overview..... 5
- Executive Summary 6
- Cb Response Architecture and Sizing 7
- Cb Response Deployment Dimensions..... 7
- Estimating Endpoint Activity..... 8
- Determining Desired Performance 9
- Determining Data Retention..... 9
- Disk Storage Requirements..... 10
 - Disk Space Requirements for Non-Data Drives 11
 - Hard Disk Performance..... 11
- Determining Maximum Endpoints per Server 11
- Deployment Planning and Sizing Summary 12
 - Sizing the Installation..... 12
 - Sizing the Server 12
 - Sizing the Cluster 12
 - Multiple Cluster Environments..... 13
 - Virtual Server Deployments 13
 - Virtual Deployment Considerations..... 14
 - Adjusting Server Requirements for Other Variables 14
- Cb Response Server Disk Configuration..... 14
 - Overview 14
 - Disk Configuration 14
- Appendix A: FAQs..... 15
 - 1. What Sensor OSes does Cb Response support? (endpoints + servers) 15
 - 2. What Cb Response console/server operating systems are supported? 15
 - 3. What is the sensor impact on the endpoint? 15
 - 4. How much network bandwidth does Cb Response require? 16
 - 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? 16
 - 6. If I do not have sufficient server resources, what is the impact?..... 17
 - 7. Is there a way to avoid storing data for high volume processes? 17

Operating Environment Requirements

8. What is the recommended number of days to retain the data and why? 17

9. What are the concerns if I want to store more than the recommended number of days' worth of data? 18

10. Are virtual servers supported? 18

11. Do the sensors support VDI environments? 18

12. How does Cb Response support disaster recovery (DR), high availability (HA) & backups? 18

13. I have fewer than 6,250 endpoints. Can I get by with a fewer server resources/hardware? 19

14. Do Server requirements change depending on the sensor type (Windows, Mac, or Linux)? 19

15. How do I contact Carbon Black Support? 20

16. Can agent-based software (antivirus, performance monitoring, backup utilities) be installed on Cb Response server? 20

Appendix B: Sample Hardware Specifications 20

Appendix C: Sample Server Configuration 21

 Dell R320 Server 21

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 1Tb, 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 6,250 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 50,000 sensors per cluster.
- The number of sensors supported and the duration of sensor data stored are driven mostly by the number 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 more CPU and disk IO capabilities.
 - 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 the per-process container. This is referred to as a process document.
 - The process document count is a key driver of performance and storage capacity.

- **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 few 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 standard activity estimate and those that require to hit the desired days of data stored should use the conservative estimates:

- **Standard Activity Estimate:** This estimate assumes an average sensor and process activity rate. This estimate is more accurate for Windows only environments, and should be selected by the majority of customers for their hardware estimates.
- **High Activity Estimate:** This estimate assumes a slightly above average sensor activity rate. This estimate will be more accurate for the full enterprise deployment, including a distribution of Linux and Mac sensors.
- **Non-Standard Activity Estimate:** This estimate assumes a higher-than-average sensor activity level. This estimate represents approximately 5-10% of all Cb Response installs and should not be considered typical.

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

Installation Estimate Type	Processes / Sensor /Day	Document Size
Standard Activity Estimate	1200	5.5 kb
High Activity Estimate	2400	7.0 kb
Non-Standard Activity Estimate	3600	9.0 kb

Note: 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 the table above 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 amount of RAM available, the performance of the disk subsystem, and the amount of hard drive storage.
- **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

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) as a single data structure. This is the “process document”.

The available system RAM determines the number of process documents that can be efficiently stored and indexed by a single Cb Response server and therefore plays a significant role in determining the number of concurrently active sensors supported or days of event data retained.

The following chart provides conservative data retention numbers for a recommended server deployment with 128GB of RAM.

Table 2: Data Retention (Days) for Activity Estimates - 128 GB RAM Server

Activity Estimate/Sensors Per Server	Standard Activity Estimate	High Activity Estimate	Non-Standard Activity Estimate
1500:	120	54	31
3000:	60	27	16
4000:	45	20	12
6250:	29	12	8

Notes:

- Once maximum storage capacity is reached, 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 the table above will change. Selecting the **High** data suppression level results in a reduction to these numbers. Selecting **None** results in an increase to these numbers.

Other RAM configurations are provided here:

Table 3: Retention (days) for 96 GB of RAM

Activity Estimate/Sensors Per Server	Standard Activity Estimate	High Activity Estimate	Non-Standard Activity Estimate
1500:	90	40	24
3000:	45	20	12
4000:	34	15	9
6250:	22	10	6

Table 4: Retention (days) for 64 GB of RAM

Activity Estimate/Sensors Per Server	Standard Activity Estimate	High Activity Estimate	Non-Standard Activity Estimate
500:	180	81	47
1500:	60	27	16
3000:	30	13	8
4000:	23	10	6

Table 5: Retention (days) for 32 GB of RAM

Activity Estimate/Sensors Per Server	Standard Activity Estimate	High Activity Estimate	Non-Standard Activity Estimate
500:	90	40	24
1500:	30	13	8
3000:	15	7	4

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 at least 30% 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 retention rate shown in Table 2:

Table 6: Recommended Hard Disk Capacity per Server at 6250 sensors/server

Activity Estimate	Standard Activity Estimate	High Activity Estimate	Non-Standard Activity Estimate
Data Partition Size	2107 GB	2324 GB	2542 GB

** The values in Table 1 do not include disk space requirements for non-data drives.

Notes:

- Required disk space can be reduced by reducing days of retention, number of active sensors, or type and number of events collected per process.
- 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.

Disk Space Requirements for Non-Data Drives

We recommend 75GB for (OS, /tmp, /etc) plus 70% of total RAM space for logging/diagnostics. For maximum stability, follow this example partition scheme:

- 75GB for “/” partition
- 70% of total RAM for “/var/log/cb”

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. The non-data space 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 driven by system RAM and disk storage, slower disk IO can have adverse effects on search performance. When this occurs, the number of storable process documents is reduced, and thus the overall data retention is reduced. We recommend using high-performance 15k spinning drives (RAID 10 minimum) or SSD drives for Cb Response storage.

Determining Maximum Endpoints per Server

A typical Cb Response server has an ingest rate of about 300-400 process documents/second, corresponding to a maximum ingest of 34.5M process documents a day. This rate is sufficient to support a typical enterprise deployment of 6250 endpoints per server for any of the three estimated activity loads. However, if endpoints generate more processes per day, then the number of supported endpoints per server may be reduced:

Notes:

- In addition to server RAM, ingest rate also 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.
- If endpoints are generating more than 6000 processes per host per day, the Cb Response Professional services team can help you deploy techniques to suppress known-good processes.

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 7: 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"> RAM Disk Size Disk Performance

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 6,250 **Windows** endpoints, and for servers within a cluster, we recommend the following hardware configuration:

Table 8: Server Sizing

Endpoints	Cluster Configuration
< 6250	Single server; 8 core @ 2.5 GHz, 128 GB RAM, at least 15k disks in RAID 10 or SSDs (all hard drives), 2.1 TB storage for data retention and 165 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 6,250 endpoints, a cluster (more than one server) is required. In general, we recommend:

- One server per 6,250 endpoints; each cluster up to 50,000 endpoints.
- One additional dedicated ‘master’ server for installations over 12,500 endpoints.
- More than 50,000 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 9: Cluster Configuration by Endpoint Count

Endpoints	Cluster Configuration
< 6,250	Single server (R320 as specified below)
6,250-12,500	Two servers, both indexers
12,500-50000	One server per 6,250 endpoints , plus one dedicated “master” node
> 50,000	Multiple clusters

If the installation has more than 12,500 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

For installations with more than 6,250 endpoints, a cluster (more than one server) is required. In general, we recommend:

Sites with more than 50,000 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 include:

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).
 - Disk and SAN considerations:
 - The /var partition should be as fast as possible and the preference is 15K disk or SSD.
- If you choose to use a SAN, we recommend placing /var 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 assume the only variable is the number of endpoints. Two other key variables exist: endpoint activity and data retention period. To explore the sizing impacts 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. 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.

- **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 section covers the most common questions and answers to Cb Response deployments.

1. What Sensor OSES does Cb Response support? (endpoints + servers)

- Windows: XP SP3 - 10 / Server 2003 - 2012R2, x86 and x64
 - Windows embedded OSES are individually evaluated.
- Mac: OS X 10.6 through 10.11.6, x64 on Intel
- Linux: RHEL & CentOS 6.4-6.8, 7.0-7.2 x64 – standard kernel versions (2.6.32-358.el6, 2.6.32-431.el6, 2.6.32-504.el6, 2.6.32-573.el6, 2.6.32-642.el6, 3.10.0-123.el7, 3.10.0-229.el7, 3.10.0-327.el7) and the standard minor/maintenance releases are supported. Non-RHEL/CentOS distributions or modified RHEL/CentOS environments (those built on the RHEL platform) are not supported.

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

- Linux
 - CentOS 6.4-6.8, (64-bit)
 - Red Hat Enterprise Linux (RHEL) 6.4-6.8 (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:
 - CPU: ~1% CPU usage. This can vary depending on system activity.
 - 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.
- Network: The Cb Response sensor generates 8-12MB of traffic per day depending on the sensor activity level. Traffic can be throttled by sensor group, per-day, per-hour.

4. How much network bandwidth does Cb Response require?

- Per endpoint:
 - 0.8-1.0 kilobits per second (kbps) per host (conservative average)
 - 8-12 megabytes (MB) per host, per day (conservative average)
- 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.

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 is 2GB on the endpoint (enough for multiple months) of data. The default is configurable on by sensor group.
 - Once the local data storage limit is reached, the sensor stops storing new log messages.
- 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 amount of allocated memory 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 significant 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 real 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 6,250 endpoints. Can I get by with a fewer server resources/hardware?

Yes. It is possible to use fewer resources for smaller deployments. Some general guidelines for Windows endpoints are as follows:

Table 10: Server Sizing

Windows Endpoints	Cluster Configuration
< 500	4 core @ 2.5GHz, 32GB RAM, 500GB storage + 97GB for non-data drives (OS, /tmp/, etc.)
501-2500	4 core @ 2.5GHz, 64GB RAM, 1TB + 120GB for non-data drives (OS, /tmp/, etc.)
2500+	8 core @ 2.5GHz, 96GB RAM, 2.1TB storage + 142GB for non-data drives (OS, /tmp/, etc.)
3500-6250	8 core @ 2.5GHz, 128GB RAM, 2.1TB storage + 165GB for non-data drives (OS, /tmp/, etc.)

Note: On the < 500 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). Although RAM might be less, the return on investment is high with the recommended 32GB of RAM. These guidelines assume typical endpoint data volumes and 25-30 days of data retention.

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. 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 (877.BIT9.098)
- Fax: 617.393.7499
- Hours: 8 a.m. to 8 p.m. EST

16. 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.

Appendix B: Sample Hardware Specifications

Below is a manifest of two systems: a Dell R320 and a Dell Blade Server (for use in a cluster).

A default hardware recommendation is included in the following table. It assumes the reference hardware is suitable, default sensor activity levels for Windows sensors, and a 25-30 day retention period.

Table 11: Sample Hardware Specs

Endpoints	Hardware
< 6,250	1x Dell R320 as specified below
6,250 – 12,500	2x Dell R320 as specified below
12,500 – 18,750	4x Dell R320 as specified below
18,750 – 50,000	1x Dell R320 dedicated master and + up to 8 Dell R320s

Appendix C: Sample Server Configuration

Dell R320 Server

Note: This is a configuration for a high-quality server for installations upwards of 6,250 sensors and roughly 30 days' worth of storage. Servers can be priced on the Dell server configuration Web Site. For smaller installations, or where there is server cost sensitivity, replacing the SSD drives with 15k RPM drives decreases the cost significantly.

- Power Edge R320
 - Power Edge R320, Intel(r) Xenon(r) E-24XX v2 Processors
- Chassis Configuration
 - 2.5" Chassis with up to 8 Hot Plug Drives [318-2039]
- Processor
 - Intel(r) Xenon(r) E5-2430 v2 2.50 GHz, 15m Cache, 7.2GT/s QPI, Turbo 6C, 80W, Max Mem 1600MHz [317-9826][338-BDZU]
- Memory Configuration Type
 - Performance Optimized [331-4428]
- Memory DIMM Type and Speed
 - 1333MT/s RDIMMs
- Memory Capacity (4)
 - 32GB RDIMM, 1333 MT/s, Low Volt, Quad Rank, x4 Data Width [319-1812]
- Operating System
 - No Operating System [420-6320]
- OS Media Kits
 - No Media Required [421-5736]
- RAID Configuration
 - RAID 1+RAID 10 for H710/H310 (2 SAS + 46 SSD HDDs in pairs)
- RAID Controller: PERC H710 Integrated RAID Controller, 512MB NV Cache [342-3529]
- Hard Drives
 - 300GB 10K RPM SAS 6Gbps 2.5in Hotplug (Quantity, 2)
 - Read Intensive MLC 3Gbps 2.5in
 - Hot-plug Drive Limited Warranty [400-ABBQ][469-1615]
 - 800GB Solid State Drive SATA (Quantity, 4)
 - Read Intensive MLC 3Gbps
 - 2.5in Hot-plug Drive Limited Warranty [400-ABBU][469-1615]

Operating Environment Requirements

- Embedded System Management
 - iDRAC7 Enterprise [421-5340][421-6085]
- Add-in Network Adapter
 - On-Board Broadcom 5720
 - Dual Port 1Gb LOM [420-4715]
- Power Supply
 - Dual, Hot-plug, Redundant Power Supply, 350W [331-7022][331-2027]
- Power Cords
 - NEMA 5-15P to C13 Wall Plug (Quantity, 2)
 - 125 Volt, 15 AMP, 10 Feet (3m), Power Cord[310-8509]
- Rack Rails
 - ReadyRails(tm) Sliding Rails with Cable Management Arm [331-4765]
- Bezel
 - Bezel [318-1431]
- Internal Optical Drive
 - No Internal Optical Drive for 8HD Chassis [318-1392]
- System Documentation
 - Electronic System Documentation and OpenManage DVD Kit for R320 [331-6962]
- Shipping
 - PowerEdge R320 Shipping [331-6952]
- Training Services
 - DES - No Dell PowerEdge Server Training Requested [973-2185]
- Hardware Support Services
 - 4 Year ProSupport and Mission Critical 4HR 7x24 Onsite Pack
 - [938-3564][938-3644][939-6767][939-6857][939-6997][989-3439][995-8481]
- Remote Consulting Service
 - Deployment Consulting 1Yr. 1 Case Remote Consulting Service [996-3179]
- Proactive Maintenance
 - No Installation [900-9997]
- Proactive Maintenance
 - Maintenance Declined [926-2979]