



Carbon Black Server Sizing Guide

(Operating Environment Requirements)

Carbon Black v5.1.0.150914

10 October 2015

Bit9, Inc.
1100 Winter Street, Waltham, MA 02451 USA
Tel: 617.393.7400 Fax: 617.393.749
E-mail: support@bit9.com
Web: <http://www.bit9.com>

Copyright ©2011–2015 Bit9, Inc. All rights reserved. This product may be covered under one or more patents pending. Bit9 and Carbon Black are registered trademarks of Bit9, Inc. in the United States and other countries. Any other trademarks and product names used herein may be the trademarks of their respective owners.

Table of Contents

[Welcome](#)

[Executive Summary](#)

[Change Log](#)

[Carbon Black Architecture and Sizing](#)

[Overview](#)

[Dimensions of a Carbon Black deployment](#)

[Estimating Endpoint Activity](#)

[Determining Desired Performance](#)

[Data Retention](#)

[Translating document count to storage requirements:](#)

[Determining maximum endpoints per server](#)

[Deployment planning](#)

[Sizing Carbon Black Installation:](#)

[Sizing the server cluster](#)

[Multiple Cluster Environments](#)

[Virtual Server Deployments](#)

[Adjusting for other variables](#)

[Appendix A: FAQ](#)

[1. What Sensor OSES does Carbon Black support? \(endpoints + servers\)](#)

[2. What Carbon Black console/server OSES are supported?](#)

[3. What is the sensor impact on the endpoint?](#)

[4. How much network bandwidth does Carbon Black require?](#)

[5. I have remote locations and/or users who travel -- how does Carbon Black work for those users, Do I need to consider impact on server sizing and configuration?](#)

[6. If I do not have sufficient server resources, what is the impact?](#)

[7. Is there a way to not store data for high volume processes?](#)

[8. What is the recommended number of days to retain the data and why?](#)

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

[10. Are virtual servers supported?](#)

[11. Do the sensors support VDI environments?](#)

[12. How does Carbon Black support disaster recovery \(DR\), high availability \(HA\) & backups?](#)

[13. I have fewer than 5,000 endpoints. Can I get by with fewer server resources/hardware?](#)

[14. Do Server requirements change depending on the sensor type \(Windows, Mac or Linux\)](#)

[15. How do I contact Bit9 Support?](#)

[Appendix B - Data Retention Calculations](#)

[Appendix C - Sample Hardware Specifications](#)

Welcome

This document provides insight into the performance and scalability considerations in deploying Carbon Black. It contains the following sections:

1. **Executive Summary:** Providing an overview of how Carbon Black operates and the underlying need to provide the resources required for a successful implementation.
2. **Carbon Black Architecture and Sizing:** Containing descriptions of the dimensions of performance and scale on Carbon Black and recommendations for server and cluster sizing
3. **Carbon Black Architecture and Sizing FAQ:** An overview covering the most common questions and answers about the performance and scale of Carbon Black within an enterprise
4. **Sample Hardware Configurations:** Listing recommended hardware for reference sizes of installations, enabling a cost estimate for the hardware to meet the scalability, performance and storage needs of customers at specific install sizes.

In addition to this document, additional resources exist to assist in properly sizing a Carbon Black installation, including:

1. **Carbon Black Sizing Estimator:** Providing interactive sizing estimates, the Carbon Black sizing calculator can be used in conjunction with a Bit9 + Carbon Black sales engineer to understand the rough sizing requirements based upon key user inputs.
2. **Professional Services:** Bit9 + Carbon Black 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 organization to more quickly detect, respond and properly remediate incidents.

Carbon Black is a continuous real-time endpoint monitoring, collection, processing, and analytics solution managing very large amounts of data, demanding a somewhat unique hardware infrastructure. In short, Carbon Black is a big data solution and is similar to netflow or data aggregation products in terms of function and processing demands. Carbon Black is unlike a typical database-driven web app. Standard databases sizes grow beyond 1Tb and Carbon Black 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 Carbon Black customers. A properly-configured system ensures Carbon Black delivers the highest-possible user experience. It is for these reasons that we require conformance to the Server Sizing Guide.

This server sizing guide has been carefully designed with our customers' success as the top priority. We make a concerted effort to ensure the Server Sizing 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.

The Server Sizing Guide 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 the Carbon Black Server Sizing Guide and we look forward to working with you to design an infrastructure that meets your specific needs.

Change Log

2015.09.25 Document Revisions:

Page	Content
1	Updated Document date to 25 September 2015 Updated Release Build to 150914
5	Added Change Log
9	Changed table 3 to reflect updated requirements Updated "Disk space requirements for the non-data drives" section to reflect current recommendations
11	Updated numbers to current specification and included space requirements for non-data drives
13	Updated Linux OS support information for Sensors and Server
17	Note added on clarification around running servers with less than 32GB RAM

Carbon Black Architecture and Sizing

Overview

Carbon Black 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 Carbon Black console. The centralized server infrastructure, can be one server or multiple servers in a cluster for larger deployments.

- Carbon Black can support up to **5,000 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 **40,000 sensors per cluster**.
- The number of sensors supported, and the duration of sensor data stored, is 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 Carbon Black 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.

Dimensions of a Carbon Black Deployment

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

1. **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:
 - a. **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.
 - b. **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 sensor generated.
 - c. **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 child process created (among others), can vary between hosts. The amount of storage each process document takes up is dependent on the activity of the process.
- Notes:**
 - Carbon Black 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.
2. **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.
3. **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 wish to monitor many watchlists or process complex watchlists, impacting 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 Carbon Black 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	1,500	5.5 kb
High Activity Estimate	3,000	7.0 kb
Non-Standard Activity Estimate	4,500	9.0 kb

Determining Desired Performance

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

1. **Data Retention:** This determines how long the Carbon Black server will store the information generated by each sensor. The longer the data is stored, the more resources that 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.
2. **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.

Data Retention

Carbon Black server stores each instance of a process execution and all event data associated with it (e.g. module loads, registry or file modification, network connections) as a single data structure, the “process document”. Available system RAM determines the number of process documents that can be efficiently stored and indexed by a single Carbon Black 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. Appendix A includes information on other RAM configurations.

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:	110	46	27
3000:	53	23	14
4000:	40	18	10
5000:	31	14	9

Note: Once maximum storage capacity is reached, Carbon Black purges the oldest data.

Disk Storage Requirements

Carbon Black server requires free space for regular operation of the data-store in addition to the size of the retained data, therefore it is recommended that at least *30% free space* is maintained for performance (included in below figures). Based on the estimated activity levels above, the following table provides information on the hard disk storage requirements to maintain the retention rate shown above:

Table 3. **Recommended Hard Disk Capacity per Server at 5000 sensors/server**

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

*Note: Values in Table 3 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 the non-data drives

We recommend 75GB for (OS, /tmp, /etc) plus 70% of total RAM space for logging/diagnostics. For maximum stability, below is an 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 etc. 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 overall data retention is reduced. We recommended that high performance 15k spinning drives RAID 10 minimum or SSD drives recommended be used for Carbon Black storage.

Determining Maximum Endpoints per Server

A typical Carbon Black 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 5000 endpoints per server for any of the three estimated activity loads. However, if endpoints generate more processes per day, then 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, Bit9 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 to assure proper performance as shown in the following table:

Table 4. **Key Factors in Deployment Planning and Sizing**

Area	Key Factors
Endpoints	<ul style="list-style-type: none"> • Number of processes generated: driven by number of endpoints, the OS type and the activity levels of endpoints • Size of process documents: Driven by the activity of the endpoints
Server(s)	<ul style="list-style-type: none"> • RAM • Disk Size • Disk Performance

Deployment Planning

Sizing Carbon Black Installation

There are two phases when sizing a Carbon Black installation:

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

Sizing the Server

For up to 5,000 Windows endpoints, and for servers within a cluster, we recommend the following hardware configuration:

Endpoints	Cluster Configuration
< 5000	Single server; 8 core @ 2.5GHz, 128 GB RAM, at least 15k disks in RAID 10 or SSDs (all hard drives), 2.1 TB storage for data retention and 160GB for non-data drives (OS, /tmp/, etc.)

Note: Please see the FAQ for details if reduction in the server sizing is preferred for smaller deployments.

Sizing the Server Cluster

For installations with more than 5,000 endpoints, a cluster (more than one server) is required.

In general, we recommend:

- One server per 5,000 endpoints; each cluster up to 40,000 endpoints.
- One additional dedicated 'master' server for installations over 10,000 endpoints.
- More than 40,000 endpoints requires multiple Carbon Black clusters.
- All nodes within the server cluster must be on the same network segment and the same 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 5. **Cluster Configuration by Endpoint Count**

Endpoints	Cluster Configuration
< 5,000	Single server
5,000-10,000	Two servers, both indexers
10,001-40,000	One server per 5K nodes, plus a dedicated 'master' node
> 40,000	Multiple clusters

If installation has more than 10,000 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 40,000 endpoints require multiple clusters. Carbon Black enables enterprise-wide management of multiple clusters via four primary subsystems:

- **Custom Alliance Feeds:** In-house threat intelligence can be syndicated over the network to all clusters from a single location
- **REST API:** The REST API provide 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:** Ability to subscribe to event streams
- **Syslog:** 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 Carbon Black. 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 Carbon Black 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, it's best to place /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. There are two other key variables: endpoint activity and data retention period. To explore sizing impacts of adjusting those values, review the output of the Carbon Black sizing calculator with Bit9 + Carbon Black sales engineering to estimate the number of servers and approximate storage required.

Appendix A: FAQ

This FAQ section covers the most common questions and answers to Carbon Black deployments.

1. What Sensor OSES does Carbon Black support? (endpoints + servers)

- **Windows:** XP SP3 - 10 / Server 2003 - 2012R2, x86 and x64
 - Windows embedded OSES are individually evaluated
- **Mac:** OS X 10.7 through 10.10, x64 on Intel
- **Linux:** RHEL & CentOS 6.4-6.6, 7.0-7.1 x64 – standard kernel versions (2.6.32-358.el6, 2.6.32-431.el6, 2.6.32-504.el6, 2.6.32-573.el6, 3.10.0.el7.x86_64) 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. For RHEL & CentOS 6.7 support, please contact Bit9 + Carbon Black Technical Support team to get access to the supported sensor.

2. What Carbon Black console/server OSES are supported?

- **Linux**
 - CentOS 6.4-6.7, (64-bit)
 - Red Hat Enterprise Linux (RHEL) 6.4-6.7 (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 Carbon Black sensor is designed to have no performance impact. The expected impact of Carbon Black on a the endpoint are:
 - **CPU:** ~1% CPU usage. This can vary depending on system activity.
 - **Memory:** 12-20MB RAM
 - **Disk Storage:** 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 reestablishment of server communications.

- **Network:** The Carbon Black sensor generates between 8 and 12MB of traffic per day depending on sensor activity level; traffic can be throttled by sensor group, per-day, per-hour.

4. How much network bandwidth does Carbon Black require?

- Per endpoint:
 - 100-125 bytes per host, per second (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, it does not reduce overall traffic. To reduce traffic, you can reduce data collected on the sensor group’s configuration.
- Note: Mac and Linux sensors, due to the number of processes generated on those endpoints, may drive higher bandwidth utilization.

5. I have remote locations and/or users who travel – how does Carbon Black work for those users? Do I need to consider impact on server sizing and configuration?

- If machines at remote locations (i.e., outside the corporate network) can reach the Carbon Black server, all operations are identical to when the endpoint is within the network.
- While not connected to the server, Carbon Black 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 stop storing new log messages.
- Customers can also deploy the Carbon Black server in their DMZ or directly on the Internet.

- For installations in a DMZ or with direct internet access, Carbon Black 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 Carbon Black will suffer.

If the server is throttling sensor uploads, the data will queue up at the endpoints until the server is either able to handle the load or when 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 not store 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. Please 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 in determining the best path forward.

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 Carbon Black has developed logic in the sensor to ensure that the VDI session maintains the sensor ID, thus ensuring that each sensor is only depicted once in the console. For VDI, Carbon Black has limited the amounts of disk writes for both persistent and non-persistent sessions to ensure that sessions are optimized for zero or thin sessions

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

- Carbon Black relies on existing system administration procedures for backup and recovery.
 - Best practices on virtual infrastructure including taking snapshots and backup 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 Carbon Black console

- For backup, sending critical Carbon Black events to SIEMs \ is a best practice
- It is possible to backup and archive the Carbon Black event data, but it requires custom scripting and can be large.

13. I have fewer than 5,000 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:

Endpoints	Cluster Configuration
< 500	2 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.)
3500-5000	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 disproportionate impact on the system as a whole (e.g. performance/stability impacts) in that although RAM could 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 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 Carbon Black installation.

In addition, On exception in Windows, but more common with Mac and Linux, endpoints running certain applications can be verbose’ as it relates to 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 Bit9 Support?

For your convenience, Bit9 Technical Support offers several channels for resolving support questions:

Technical Support Contact Options
Web: www.bit9.com
E-mail: support@bit9.com
Phone: 877.248.9098 (877.BIT9.098)
Fax: 617.393.7499
Hours: 8 a.m. to 8 p.m. EST

Appendix B - Data Retention Calculations

Retention (days) for 128 GB of RAM

Activity Estimate/ Sensors Per Server:	Standard Activity Estimate	High Activity Estimate	Non-Standard Activity Estimate
1500:	110	46	27
3000:	53	23	14
4000:	40	18	10
5000:	31	14	9

Retention (days) for 96 GB of RAM

Activity Estimate/ Sensors Per Server:	Standard Activity Estimate	High Activity Estimate	Non-Standard Activity Estimate
1500:	81	35	20
3000:	53	20	12
4000:	40	18	10
5000:	24	10	7

Retention (days) for 64 GB of RAM

Activity Estimate/ Sensors Per Server:	Standard Activity Estimate	High Activity Estimate	Non-Standard Activity Estimate
500:	158	69	40
1500:	53	23	14
3000:	27	12	7
4000:	20	9	5

Retention (days) for 32 GB of RAM

Activity Estimate/ Sensors Per Server:	Standard Activity Estimate	High Activity Estimate	Non-Standard Activity Estimate
500:	78	32	20
1500:	27	12	7
3000:	15	6	4

Appendix C - 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 table below. It assumes the reference hardware is suitable, default sensor activity levels for Windows sensors and 25-30 day retention period.

Endpoints	Hardware
< 5000	1x Dell R320 as specified below
5000 - 10000	2x Dell R320s as specified below
10001 - 15000	4x Dell R320s as specified below
15001 - 40000	1x Dell R320 + Dell Blade Chassis w/ required blades

Sample Server Configurations

Dell R320 Server

Note: This is a configuration for a high-quality server for installations upwards of 5,000 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 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]
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]