



Contents

Document Change Log	3
Overview	3
Executive Summary	4
Carbon Black EDR Architecture and Sizing	4
Carbon Black EDR Deployment Dimensions	5
Factors Impacting Performance and Retention	6
Estimating Endpoint Activity	6
Estimating Endpoint Activity Size on Disk	6
Determining Desired Operational Environment	7
Disk Space Requirements for Non-Data Drives	9
Hard Disk Performance	9
Cluster Sizing	9
Multiple Cluster Environments	10
Virtual Server Deployments	11
Virtual Deployment Considerations	11
Carbon Black EDR Server Disk Configuration	12
Appendix A: FAQs	12
1. What endpoint operating systems do Carbon Black EDR sensors support?	12
2. What Carbon Black EDR console/server operating systems are supported?	12
3. What processors are supported?	13
4. Is there any difference in CPU core count if vCPUs are used?	13
5. What is the sensor impact on the endpoint?	13
6. Can I use a spinning disk if my deployment size or retention needs are small?	13
7. How much network bandwidth does Carbon Black Response require?	13
8. I have remote locations and/or users who travel. How does Carbon Black EDR work for those users? Do I need to consider the impact on server sizing and configuration?	14
9. What is the impact if I do not have sufficient server resources?	15
10. Is there a way to avoid storing data for high volume processes?	15
11. What are the concerns if I want to store more than the recommended number of days' worth of data?	16
12. Are virtual servers supported?	16
13. Do sensors support VDI environments?	16

14.	How does Carbon Black EDR support disaster recovery (DR), high availability (HA) and backups?	16
15.	How do I contact Carbon Black Support?	16
16.	Can agent-based software (antivirus, performance monitoring, backup utilities) be installed on a Carbon Black EDR server?	17

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.
01/30/2019	Updated branding.
03/22/2019	GA 6.3 release revision. Updated supported operating systems. Other minor editing.
04/01/2020	GA 7.1 release revision. Major documentation update.
04/07/2020	Updated link to supported endpoint operating systems.
09/25/2020	GA 7.3 release revision. Updated branding. Updated supported operating systems. Added processor requirement.

Overview

This document provides insight into performance and scalability considerations in deploying VMware Carbon Black EDR. It specifically provides the following information:

- An overview of how Carbon Black EDR operates, and the resources that are required for a successful implementation.
- Dimensions of performance and scale and recommendations for server and cluster sizing.
- The most common questions and answers about the performance and scale of a Carbon Black EDR installation.
- Recommended hardware for referenced sizes of installations, which enables a hardware cost estimate to meet the scalability, performance, and storage needs for specific install sizes.

The following additional resources can help properly size a Carbon Black EDR installation:

- **Carbon Black EDR Sizing Estimator:** By providing interactive sizing estimates, the Carbon Black EDR sizing calculator is used together with a Carbon Black sales engineer to understand rough sizing requirements based on key user inputs.
- **Professional Services:** Carbon Black customers can engage Customer Support or Professional Services for additional assistance in scoping, sizing, and tuning their installations.

Executive Summary

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

Carbon Black EDR is a continuous real-time endpoint monitoring, collection, processing, and analytics solution that manages very large amounts of data and demands a unique hardware infrastructure. Carbon Black EDR is a big data solution and is similar to netflow or data aggregation products in function and processing demands. Carbon Black EDR is unlike a typical database-driven web app. It is not uncommon for database sizes to grow beyond 10TB; Carbon Black EDR often processes and analyzes billions of data points of information per day.

Insufficient or inappropriate hardware configurations account for the majority of performance-related issues that Carbon Black EDR customers encounter. A properly-configured system ensures that Carbon Black EDR delivers the highest-possible user experience. Therefore, we require conformance to our server sizing guidance. Carbon Black Support and Professional Services teams cannot assist with performance-related issues until the deployment conforms to the recommendations in this guide.

This document is 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 will guide you to the necessary hardware and storage configurations to provide a great experience, and put your organization in the best posture possible. We look forward to working with you to design an infrastructure that meets your specific needs.

Carbon Black EDR Architecture and Sizing

Carbon Black EDR 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 EDR console. The centralized server infrastructure can be one server or multiple servers in a cluster.

- Carbon Black EDR can support up to 18,750 sensors and/or up to 10.5 TB of process event data 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 (provided that event data volume per minion server remains under 10.5 TB).
- The number of sensors supported and the duration of stored sensor data are primarily driven by the number and the process activity launched by each endpoint. Endpoints vary widely in the volume of processes generated depending on the operating system (OS) and the software that is running on the endpoints. These are the most important factors that drive scale.
- The Carbon Black EDR 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.

Carbon Black EDR Deployment Dimensions

Different data dimensions impact performance and scalability. This section provides an overview of the most critical components that drive scale of a Carbon Black EDR server installation.

- **Incoming data rates:** Each sensor sends a stream of data to the server; that data requires indexing and storage. The incoming data volume that is generated for an installation is impacted by three measures:
 - The number of concurrently active sensors: the server must process event data from each endpoint. More sensors increases CPU and disk IO requirements.
 - The number of processes per sensor per day: Planning requires estimates for the typical rate of processes per endpoint per day, but there can be wide variation between installations. Operating systems vary in the number of processes that are generated per sensor.
 - The average activity per process: Similar to the count of processes per endpoint per day, the activity of those processes (for example, file modifications, registry changes, network connections, child processes, etc.) can vary between endpoints. The amount of storage each process document takes is dependent on the activity of the process.

- **Notes:**
 - Carbon Black EDR uses a storage model that is based on per-process storage where the activities of a given process are stored within a set of per-process containers. These are referred to as *process documents*; each process document represents a snapshot of the process activity in a period of time (typically five minutes).
 - Each process can have one or more associated process documents. New process documents are created each time the sensor sends data to the server. Therefore, longer-living processes (such as `svchost.exe`) can have many associated process documents. The average process-to-document ratio is between two and three.
 - The Carbon Black EDR datastore collates multiple process documents into logical data volumes called *shards*. Multiple shards are created (based on size or number of days) to achieve the desired retention period while maintaining optimum search performance.
 - The process document count is a key driver of performance and storage capacity.

- **Threat intelligence feeds:** Depending on the number of enabled feeds, the server monitors for activity related to a number of unique indicators. Most organizations do not notice any impact; however, some organizations might monitor a very large number of indicators, and therefore require additional resources.
- **Watchlists:** For each configured watchlist, the server runs a search every ten minutes. Performance is impacted by the number and complexity of these searches. Most organizations will not notice an impact, but some organizations might monitor many watchlists or process complex watchlists; this will impact performance (for example, lengthy watchlists or searches that contain wildcards).

Factors Impacting Performance and Retention

The count and size-on-disk of process documents are the key factors that drive performance and storage requirements of a Carbon Black EDR deployment. This section reviews factors that impact endpoint activity levels, and provides guidance around typical ranges across current Carbon Black EDR customers.

Estimating Endpoint Activity

Endpoint activity varies significantly across different deployment environments and OS platforms.

The following factors contribute to endpoint activity levels:

- **The endpoint OS:** Carbon Black EDR tracks and reports endpoint activity on a per-process execution basis. In most cases, Microsoft Windows creates and dismantles fewer processes than *nix-based operating systems such as macOS and Linux. Therefore, Microsoft Windows endpoints result in lower endpoint activity levels.
- **Endpoint type:** In most cases, an endpoint that is deployed as a server results in higher endpoint activity levels than a general purpose workstation.
- **Endpoint use case:** A build machine results in higher endpoint activity levels (for example, file modifications and created binaries). A DNS server results in higher endpoint activity levels (for example, created network connections.)

Estimating endpoint activity can be challenging. Incorporating known factors (OS breakdown, server versus workstations) into the sizing process results in a better experience. The following estimates can help determine the required server specifications:

Table 1 Endpoint activity level percentiles

	Windows	macOS	Linux
Median	7,800	12,000	59,750
75-Percentile	10,750	18,750	125,000
90-Percentile	16,000	25,500	195,750
99-Percentile	34,750	82,750	819,250

Table 1 shows endpoint activity level percentiles for process documents per endpoint per day, generated by different OS platforms across a Carbon Black EDR customer base (minimum of 100 endpoints).

Most endpoints have activity levels within the median range for each OS type. Servers and endpoints that are used for high performance computing, simulations, or build machines can fall within the 75-percentile to 90-percentile range. Special cases might encounter higher endpoint activity levels, but it is unlikely that all endpoints will be above the 90-percentile range.

Estimating Endpoint Activity Size on Disk

A Carbon Black EDR server stores endpoint process activity on disk as searchable and retrievable process documents. The size of these process documents varies across environments and OS platforms.

The following factors contribute to on-disk size of process documents:

- **Endpoint OS.** Microsoft Windows processes are in general longer-lived compared to *nix based operating systems such as macOS and Linux. Therefore, such processes record more events per process, resulting in larger size on disk. Short-lived macOS and Linux processes (for example, `ps`, `cat`, `ls`) result in much smaller size on disk per process.
- **Endpoint type.** In most cases, a server results in larger process document size on disk because servers run long-lived services (for example, daemons).
- **Endpoint use case.** An endpoint that runs applications that modify many files or registry entries results in higher process document sizes on disk.

Estimating process document size on disk can be challenging. Incorporating known factors (OS breakdown, server versus workstations, etc.) into the sizing process results in a better experience. The following estimates can help gauge the required server specifications:

Table 2 On-disk size percentiles

	Process Document Size (Bytes)
Median	3,600
75-Percentile	4,750
90-Percentile	6,250
99-Percentile	13,800

Table 2 shows process document on-disk size percentiles across a Carbon Black EDR customer base (minimum of 100 endpoints).

Determining Desired Operational Environment

A Carbon Black EDR server stores each instance of a process execution and all event data with which it is associated (for example, module loads, registry or file modifications, and network connections) in process documents. Process documents from multiple sensors are stored in database structures known as *shards*. To provide optimum storage and search performance, shards are periodically rotated when they reach a certain disk size.

For best performance, the recommended maximum size of a single shard is 500 GB, with up to 12 searchable shards per server for a total process event retention of 6 TB. However, for increased retention, shard size and count can be increased to 750 GB and 14 respectively with additional hardware. The latter configuration can extend total event retention of a single Carbon Black EDR server to 10.5 TB. It is important to note that, while more and bigger shards result in increased retention, fewer and smaller shards result in an improved search performance. You can determine the desired configuration (and retention) for your needs as follows.

Compute an estimate of the event data that the server will store (given your retention needs). Use the data provided in [Table 1](#) and [Table 2](#) to calculate this value:

Equation 1 Calculating Total Event Data Volume

$$\text{Total Event Data Volume} = (\text{Number of Endpoints} * \text{Number of Process Documents per Endpoint per Day} * \text{Size of Process Document on Disk}) * \text{Retention Days}$$

If the estimated event data volume is less than 10.5 TB, see [Table 3](#) to select the appropriate server size. You should choose the smallest shard count and size combination that meets your total event data. A single Carbon Black EDR server is sufficient to support your deployment. If the estimated value is larger than 10.5 TB, see [Cluster Sizing](#) to determine the clustered deployment that is necessary to support your environment.

If OS or endpoint breakdown is known or can be estimated, this information can be incorporated into the previous calculation to improve planning accuracy. For example, a deployment includes 1,000 Windows workstations (50-percentile) and 100 Linux servers (75-percentile). A reasonable estimate would be:

$$\text{Required Event Data Storage} = (1000 * 7800 + 100 * 125000) * 3.6\text{KB} * 30 \text{ days} \sim 2.1\text{TB}$$

Because this value is less than 10.5 TB, a single server can accommodate this deployment. The next step is to determine the appropriate server size based on the desired search performance.

Table 3 Server sizing chart based on event data volume

Event Data Volume	Shard Size	Number of Shards	Disk Space for Event Data*	Memory	CPU	Example Configuration
< 500GB	50GB	10	1TB	16 GB	2 cores	2.3 GHz Intel Xeon E5-2686 v4 Processor or similar. SSD disks with 16000 IOPS (250 MiB/s throughput)
< 2TB	200GB	10	3TB	32 GB	4 cores	
< 3.6TB	300GB	12	5TB	64 GB	8 cores	
< 6TB	500GB	12	8TB	128 GB	16 cores	
< 10.5TB	750GB	14**	13TB	256 GB	32 cores	

*Disk size is the space that is required to store Carbon Black EDR event datastore, including necessary overhead for database optimization and binary storage. Additional disk space is required for storing non-event-data OS files.

**Complex queries take more time to execute, resulting in slower console response times.

After the server sizing is completed, configure the Carbon Black EDR application with the following settings in `cb.conf`:

- `SolrTimePartitioningMinutes=8640` (for 6 days)
- `SolrTimePartitioningActivePartitions= [Number of Shards]`
- `SolrTimePartitioningMaxSizeMB= [Shard Size in MB]`

- #MaxEventStoreDays (comment out, partition size and count determines days)
- #MaxEventStoreSizeInMB (comment out, partition size and count determines size)
- MaxEventStoreSizeInPercent=90 (a safeguard to avoid running out of disk)

See the *VMware Carbon Black EDR Server Configuration (cb.conf) Guide* for instructions on how to configure these parameters. See also [Cb Enterprise Response - Managing Retention](#).

Disk Space Requirements for Non-Data Drives

In addition to 10 GB for the root drive that hosts the OS, we recommend reserving 75 GB for /tmp and disk space that equals to 70% of total RAM for logging/diagnostics and process memory dumps (if needed).

The following is an example partition scheme for a machine that has 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

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

Note: Additional disk space can be optionally added to enable greater retention of logging and diagnostic data.

Hard Disk Performance

Process document retention is primarily limited by total available disk storage. System performance is determined by disk IOPS and throughput. We recommend using SSD drives on RAID 5 for all Carbon Black EDR deployments with 16,000 IOPS (250 MiB/s throughput) or equivalent performance characteristics.

We do not support network-attached storage (NAS) because NAS is often slower, displays larger latencies with a wider deviation in average latency, and is a single point of failure. We also do not support NFS for similar reasons.

Cluster Sizing

If the **Total Event Data Volume** calculated from [Equation 1](#) exceeds the capacity of a server's event data volume ([Table 3](#) Column 1), a single server will not meet the requirements. Carbon Black EDR supports a clustered configuration to allow horizontal scaling for larger deployments. A clustered deployment is a set of servers (minions) that work together with a head node (master) for horizontal scaling. Each minion and master of a clustered deployment must conform to the same specifications of a single server deployment.

The number of minions that are required to support event data volume can be determined by dividing the value calculated in [Equation 1](#) by the supported event data volume of your server

hardware (Table 3, Column 1) and rounding up to the nearest integer. If the result is larger than two, an additional eventless master is needed.

Note: A special use case exists if the result is *two*. A single master and single minion cluster configuration can be used where both the master and minion store event data (the master must be provisioned according to the minion hardware requirements).

Table 4: Cluster size selection based on core size estimation

Total Data Volume Required	Cluster Configuration
<= 6TB	Single server
6TB - 12TB	Two servers, both indexers
12TB - 48TB	One server per 6TB of event data, plus one dedicated master node
> 48TB	Multiple clusters

Table 4 assumes that 6 TB server event storage size from Table 3 is used for each minion. However, if another server event storage size is chosen, corresponding data volume size (Table 3, column 1) should replace the value of 6 TB in calculating the necessary number of cluster nodes.

For example, if the deployment is for 25,000 Windows workstations and 3,500 Windows servers, an estimated data volume for 25 days of retention is:

$$(25000 * 7800 + 3500 * 10750) * 3.6KB * 25 \text{ days} \sim 19.5TB$$

Because this value is larger than 6 TB, the required number of minions in the clustered deployment must be at least four (plus one head node). Alternatively, a two-minion and one-master node cluster can be configured, with each using the 10.5 TB data configuration.

Note: When using a dedicated eventless master node, disk that is allocated to event data can be reduced to 500 GB. In eventless master nodes, this storage amount provides the necessary space to store threat intelligence feeds, alerts, and binary metadata information that is shared across the cluster. CPU and memory allocations should match minion nodes. Additional disk space is required for storing non-event-data OS files.

Multiple Cluster Environments

For installations that exceed a cluster size of eight servers (plus one master), multiple clusters are required. Carbon Black EDR enables enterprise-wide management of multiple clusters via four primary subsystems:

- **Custom Carbon Black 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 other cluster configuration details. Scripts for common tasks are available from our support staff.
- **Enterprise Messaging Bus:** Can subscribe to event streams. Carbon Black EDR provides an add-on Event Forwarder to forward event data to third party SIEMs or other integration points. For best performance, we recommend that Event Forwarder be configured to run on a separate server. The recommended Event Forwarder configuration is as follows:

- Server Operating System:
 - CentOS 6.7-6.10 (64-bit)
 - CentOS 7.3-7.8 (64-bit)
 - The server can be either physical or virtualized.
- Choose the same CPU/RAM level from [Table 3](#) to match your Carbon Black EDR server specification
- 4TB Enterprise Grade SSD for store-and-forward.
- **Syslog:** Combines and forwards (alerting) information from multiple clusters into a central location (such as a SIEM).
- **Unified View Server:** (See the *VMware Carbon Black EDR Unified View User Guide*.) Provides unified login and capabilities across multiple clusters. A base configuration requires the following:
 - Server Operating System (see [What Carbon Black EDR console/server operating systems are supported?](#))
 - The server can be either physical or virtualized.
 - For < 10 Carbon Black EDR servers:
 - 8 GB of RAM
 - 4 CPU cores and 500 GB Storage
 - For 10 to 100 Carbon Black EDR servers:
 - Minimum 16 GB of RAM
 - Minimum 8 CPU cores. - 1 TB storage
 - Storage is required for OS files and logging only. Unified View is not I/O intensive.

Multi-cluster solutions mitigate bandwidth where endpoints are geographically dispersed. Network bandwidth loads are constrained to local, higher-speed links in local area networks; 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 required resources are available to Carbon Black EDR. 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 must be dedicated and reserved to mitigate these risks.

Virtual Deployment Considerations

- Dedicated and Reserved CPU, RAM, and IOPs resources provided to the Carbon Black EDR virtual machine per this document
- Hardware specifications and recommendations persist (# processor, # RAM, disk allocation).
- Disk and SAN considerations:
 - Data partition (defaults to `/var/cb/data` directory) should be as fast as possible. SSD Disk I/O must meet requirements listed in [Table 3](#).
- If you use a SAN, we require placing the data partition (defaults to `/var/cb/data` directory) on a dedicated LUN.

Carbon Black EDR Server Disk Configuration

Because Carbon Black EDR is a very I/O intensive application, high performance disks and RAID configurations are required. This section provides the recommended disk configuration for each type of disk partition.

- The data disk partition stores the primary data that requires high performance equipment. Carbon Black EDR places its high-volume data in 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 that require 2 TB of storage space or more, at least five solid-state SAS drives in a RAID5 configuration are required.
- For partitions requiring 5 TB or more of storage space, we require two equally sized volumes to alternate in storing and optimizing event partitions. See the *VMware Carbon Black EDR Server/Cluster Management Guide*.
- We require that the disks are qualified using the Carbon Black EDR Qualifier tool. Please work with Carbon Black staff to qualify the disk and determine that sufficient throughput is available for your deployment size.
- Non-data disk partitions do not require the same I/O and space requirements as the data disks. Therefore, 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 common questions and answers to Carbon Black EDR deployments.

1. What endpoint operating systems do Carbon Black EDR sensors support?

For the current list of supported operating systems for Carbon Black EDR sensors, see [Carbon Black EDR Supported Versions Grid](#).

Distributions that are not specified in this document, or modified versions of those distribution environments, are not supported.

2. What Carbon Black EDR console/server operating systems are supported?

For best performance, Carbon Black recommends running the latest supported software versions. For default CentOS/RHEL installations, the product supports the following for the Carbon Black EDR server and Unified View:

- CentOS 6.7-6.10 (64-bit)
- CentOS 7.3-7.8 (64-bit)
- CentOS 8.1-8.2 (64-bit)
- Red Hat Enterprise Linux (RHEL) 6.7-6.10 (64-bit)
- Red Hat Enterprise Linux (RHEL) 7.3-7.8 (64-bit)
- Red Hat Enterprise Linux (RHEL) 8.1-8.2 (64-bit)

However, if you are pinning dependencies to a specific OS version, the product only supports the following for the Carbon Black EDR server and Unified View:

- CentOS 6.7-6.10 (64-bit)
- CentOS 7.5-7.8 (64-bit)
- CentOS 8.2 (64-bit)
- Red Hat Enterprise Linux (RHEL) 6.7-6.10 (64-bit)
- Red Hat Enterprise Linux (RHEL) 7.5-7.8 (64-bit)
- Red Hat Enterprise Linux (RHEL) 8.2 (64-bit)

Note: For the Carbon Black EDR server and Unified View, versions 7.3, 7.4, and 8.1 (64-bit) of CentOS/ RHEL are not supported if you are pinning dependencies. In general, VMWare Carbon Black discourages pinning dependencies as per the [recommendation from CentOS/RHEL](#).

Installation and testing are performed on default installations using the minimal distribution and the distribution's official package repositories. Customized Linux installations must be individually evaluated and are ineligible for support.

3. What processors are supported?

Carbon Black EDR supports Intel or AMD x86 processors. Carbon Black EDR does not support ARM processors.

4. Is there any difference in CPU core count if vCPUs are used?

CPU Cores refers to CPU threads. You can allocate 1 vCPU = 1 physical CPU core.

5. What is the sensor impact on the endpoint?

The Carbon Black EDR sensor is designed to have no performance impact. Endpoint activity levels might impact actual values. Typical ranges for the impact of the Carbon Black EDR sensor are as follows:

- CPU – < 5% CPU usage, depending on system activity
- Memory – 12-50 MB RAM
- Disk Storage – The sensor regularly sends data to the server, requiring minimal storage on the endpoint (500 KB to 3 MB). If the sensor cannot communicate with the server, data queues up to an adjustable threshold (2 GB by default, expected 30-60 days activity on a normal system). The data is synced when server communications are reestablished.

6. Can I use a spinning disk if my deployment size or retention needs are small?

For installations that have less than 1 TB of data, you can reduce the recommended disk configuration to four SSDs, or use 6 Gb/s SAS 15K RPM spinning drives.

7. How much network bandwidth does Carbon Black Response require?

It is difficult to predict the actual network traffic that Carbon Black EDR requires. Network bandwidth depends on many factors, including sensor activity and the number of unique binary files that are uploaded to the server. Apply the following estimates:

- Per endpoint:
 - 1-4 kilobits per second (kbps) per host

- 10-40 megabytes (MB) per host per day
- Table 5 shows server-side expected average network traffic based on sensor activity estimates (not including network usage for uploading unique binaries to the server):

Table 5: Average expected server network incoming traffic in MB/s for event data

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 limits bandwidth from a group of sensors. Throttling is often used on low-bandwidth sites or sites that are bandwidth-constrained at certain times of 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. This override instructs the host to ignore any configured throttles and immediately send all data.
 - 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 limited the collection of certain type of events per-process on a per-sensor-group basis. For more about sensor groups, see the *VMware Carbon Black EDR User Guide*.
- Maximum sensor check-in rate can be configured through `SensorCheckingDelayRate` in `cb.conf`.
 - The default value is 100, which corresponds to a maximum 100 check-ins/second/server node. Reducing this value reduces check-in network traffic, but also reduces how often sensors send statistics and retrieve configuration changes.

Note: Due to the number of processes that are generated on those endpoints, macOS and Linux sensors can drive higher bandwidth utilization.

8. I have remote locations and/or users who travel. How does Carbon Black EDR work for those users? Do I need to consider the impact on server sizing and configuration?

- If endpoints at remote locations (for example, outside the organization’s network) can reach the Carbon Black EDR server, all operations are identical to when the endpoint is in the network.

- When they are not connected to the server, Carbon Black EDR sensors queue data on the endpoint (up to a configurable threshold) until the server is reconnected.
 - Default storage on the endpoint is 2 GB or 2% of total disk storage (whichever is reached first). This should be sufficient for multiple months of data. The default is configurable by sensor group.
 - After the local data storage limit is reached, the sensor stops storing new log messages.
- The Carbon Black EDR server can be deployed in the DMZ or directly on the Internet.
- For installations in a DMZ or with direct Internet access, it is best practice to configure Carbon Black EDR to restrict access to the management interface (the console) to a separate, internal network interface.
- This behavior does not impact server sizing.

9. What is the impact if I do not have sufficient server resources?

If there are insufficient resources, the server will throttle the sensors in their sending of data. If the CPU, RAM, or storage resources or performance are insufficient for the stored process documents, the search performance of Carbon Black EDR is impacted and the environment will be ineligible for support until the environment is brought into alignment with the resource requirements that are outlined in this document.

If the server throttles sensor uploads (or an external throttle, such as network level QoS or traffic shaping is implemented), the data queues up at the endpoints until one of two things occurs: the server can handle the load, or the sensors meet the local threshold for how much data to queue. Endpoints continue to operate without any noticeable performance impact.

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

You can deploy event filters on the server 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 to enable event filtering.

You can reduce the data that is recorded per-process on a per-sensor-group basis. In the sensor group's configuration page, one can limit the collection of the following:

- 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 reduces data volumes, but at a loss of forensic visibility. 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. Engage with Sales Engineering or Professional Services to determine the best approach.

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

Days stored increase the required data storage and can impact the number of servers that are required to process the stored data. A performance degradation occurs when the number of stored process documents exceeds the recommended number of days stored. Use the Sizing Estimator together with Sales Engineering or Professional Services for tailored guidance. Forwarding select or all events to a separate SIEM solution for longer retention is a recommended best practice.

12. Are virtual servers supported?

Yes. See [Virtual Server Deployments](#).

13. Do sensors support VDI environments?

Yes. Sensors running on a Virtual Desktop Infrastructure (VDI) are supported for both persistent and non-persistent VDI setups. To maintain continuity through non-persistent sessions, Carbon Black EDR has developed logic in the sensor to make sure that the VDI session maintains the sensor ID. This persistence ensures that each sensor is depicted one time only in the console. For VDI, Carbon Black EDR limits the disk writes for both persistent and non-persistent sessions to optimize for zero or thin sessions. See the *VMware Carbon Black EDR Integration Guide*.

14. How does Carbon Black EDR support disaster recovery (DR), high availability (HA) and backups?

- Carbon Black EDR relies on existing system administration procedures for backup and recovery. See the *VMware Carbon Black EDR Server/Cluster Management Guide*.
 - Best practices when using virtual infrastructure include taking snapshots and backing up using your enterprise's IT management procedures.
 - Best practices on physical hardware use RAID to help against hard drive failures; maintain cold spares.
- 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 when the connection is re-established. In the Carbon Black EDR console, you can configure the data that is stored on the endpoint.
- For backup, send critical Carbon Black EDR events to SIEMs as a best practice.
- You can backup and archive the Carbon Black EDR event data, but this requires custom scripting and can be large.

15. How do I contact Carbon Black Support?

Carbon Black Technical Support offers several channels for resolving support questions:

- [Customer Support Portal](#) in the User Exchange
- 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 a Carbon Black EDR server?

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