

# Symantec™ Endpoint Protection 14 Sizing and Scalability Best Practices White Paper

# Symantec Endpoint Protection Sizing and Scalability Best Practices

Product version: 14.x

Documentation version: 4

This document was last updated on: April 24, 2019

## Legal Notice

Copyright © 2019 Symantec Corporation.

All rights reserved.

Symantec, the Symantec Logo, the Checkmark Logo and LiveUpdate are trademarks or registered trademarks of Symantec Corporation or its affiliates in the U.S. and other countries. Other names may be trademarks of their respective owners.

# Best practices for sizing and scalability

This document includes the following topics:

- [About sizing and scalability for Symantec Endpoint Protection](#)
- [The challenge of sizing security protection in the enterprise](#)
- [System design and planning](#)
- [Site design](#)
- [Determining client-to-server ratios](#)
- [Management server and database sizing](#)
- [Reducing log sizes](#)
- [Windows client installation packages](#)
- [Windows client installation package and content update sizes](#)
- [Calculating total disk space requirements](#)

## About sizing and scalability for Symantec Endpoint Protection

Symantec Endpoint Protection provides best-of-breed endpoint security to enterprises of all sizes. Symantec recommends that you consider multiple variables to determine your sizing and deployment needs. Careful consideration can help to create optimum protection and serviceability. This white paper provides recommendations in the following areas:

- Single- and multiple-site environments

See “[Site design](#)” on page 6.

- Client-to-server ratios  
See “[Determining client-to-server ratios](#)” on page 8.
- Database sizing  
See “[Symantec Endpoint Protection Manager hardware and software recommendations](#)” on page 14.  
See “[Choosing a database based on client-to-server ratio](#)” on page 15.
- Log-keeping and sizes  
See “[Reducing log sizes](#)” on page 17.  
See “[Calculating log sizes](#)” on page 18.

The architecture, designs, and recommendations that are provided in this guide are based on metrics from internal testing of the product. These tests are performed in an isolated environment. Implementations in production environments may result in some performance metrics that vary from the testing scenarios. These variations can alter the recommended sizing and architecture.

This guide references possible changes and modifications to Symantec Endpoint Protection capability, functions, metrics, and features. These changes are subject to continuous evaluation and should not be considered as firm commitments by Symantec.

---

**Note:** The data in this white paper reflects the scenarios that were tested using Symantec Endpoint Protection build 14.0.1899, unless otherwise noted.

---

For more information about how to configure the options in this guide, see the [Symantec Endpoint Protection Installation and Administration Guide](#).

## The challenge of sizing security protection in the enterprise

Successful Symantec Endpoint Protection configurations and deployments depend on the following factors:

- The Symantec Endpoint Protection technologies to be deployed
- Whether different security policies are needed for users in different locations
- Whether different administrative groups and policies are needed for desktops, servers, laptops, users, and departments
- The number of geographic locations within the company
- The frequency at which content updates are applied

- Whether client patches should be automatically deployed
- The desired method of content distribution
- Whether a high availability infrastructure is present or desired
- Log retention times
- The frequency of requests for log or reporting data older than one week, one month, and one year
- How often metrics need to be gathered
- Who and where are people who need access to the data
- Requirements to tie into an existing third-party tool or authentication scheme

Knowing how to evaluate these variables is crucial to establishing an effective, efficient, and sustainable endpoint protection solution.

## System design and planning

Effective and efficient endpoint security requires a balance of protection technologies, a manageable infrastructure, and adequate forensic data to properly monitor network security activities. Before you deploy Symantec Endpoint Protection, you need to make several decisions about the best ways to configure the components for your particular environment.

### Architectural components

Symantec Endpoint Protection contains the following main architectural components that work together to protect your company from security threats:

- Symantec Endpoint Protection Manager – The management server that is used to configure clients, reports, and alerts.
- Symantec Endpoint Protection database – The embedded database or Microsoft SQL Server database that stores all configuration, updates, and reporting information.
- Symantec Endpoint Protection Manager console – A lightweight user interface that you use to access Symantec Endpoint Protection Manager. The Symantec Endpoint Protection Manager console is used to manage and view deployment activity, configurations, updates, and client reports.
- Symantec Endpoint Protection client – Software that is deployed to the Windows, Mac, and Linux computers in your network. The client monitors your security policies and automates your policy compliance.

[Optional components for Symantec Endpoint Protection](#)

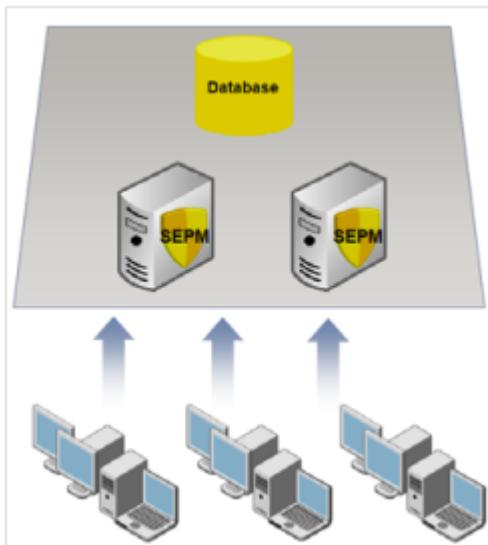
# Site design

A Symantec Endpoint Protection site design begins with the choice of the site architecture. At the highest level, designs are divided between single-site designs and multiple-site designs.

## Single-site design

An organization with one data center can generally use a single-site design with the following attributes:

- Two instances of Symantec Endpoint Protection Manager, for redundancy and load balancing
- Database clustering, to support high availability



## Multi-site design

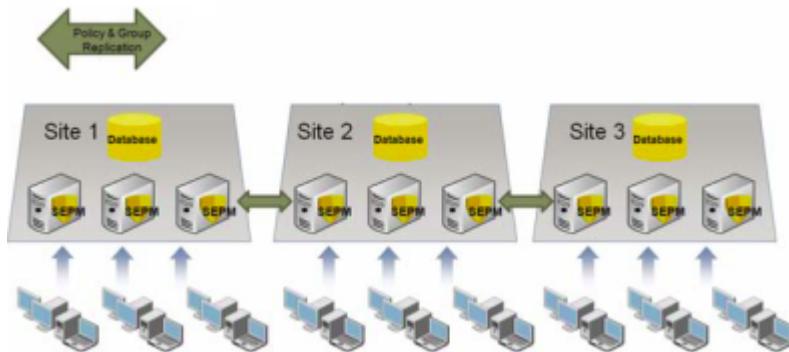
An organization that has more than one data center or that has many large locations should use a multiple-site design. Consider the following three primary designs for a multiple-site environment:

- Distributed
- Central logging
- High availability

## Distributed

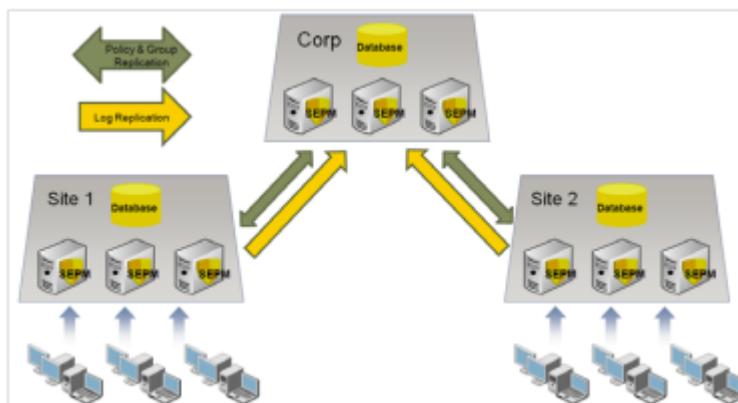
The distributed design is recommended when immediate access to remote site data is not critical. This design has the following attributes:

- Each site performs bidirectional replication of groups and policies.
- Logs and content are not replicated by default.
- To view the site reports, administrators use the console to connect to the Symantec Endpoint Protection Manager server at each remote site.



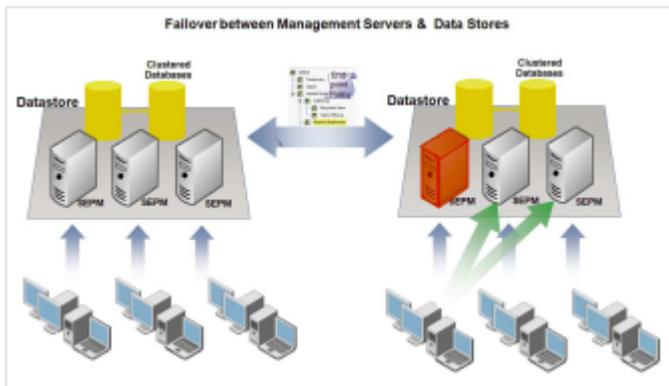
## Central logging

The central logging design is recommended when centralized reporting is required. The principal feature of this design is log forwarding to a centralized repository. In the following example, the corporate headquarters site is the central repository for the logs that are forwarded from corporate sites 1 and 2.



## High availability

The high availability design takes advantage of multiple Symantec Endpoint Protection Manager installations and multiple clustered databases to provide redundancy, failover, and disaster recovery. Several options are available to optimize performance, failover, and recovery. For example, the high availability design can be configured to have client computers automatically switch to an alternate Symantec Endpoint Protection Manager if the primary server becomes unavailable.



## Determining client-to-server ratios

Deploying Symantec Endpoint Protection with the proper client-to-server ratio is crucial to providing a high-performance endpoint security environment. The significant parameters that affect the client-to-server ratio are client-server communication, desired update speeds, and the security technologies that are deployed in the network environment.

### Client-server communication

Symantec Endpoint Protection clients and management servers exchange status information, content data, and policy information. Clients initiate this communication with Symantec Endpoint Protection Manager depending on the release:

- If you install version 14 or later for the first time, the clients communicate with management server using HTTPS and the TLS 1.2 protocol. Clients on Windows versions that do not support TLS natively cannot communicate with a 12.1.6 or later management server using HTTPS.

[Enable HTTPS client-server communications](#)

- If you upgrade from 12.1.x to 14 or later, the clients communicate using the protocol that was previously in use, either HTTP (TCP port 8014) or HTTPS (TCP port 443). In the event of a conflict, you can change the ports.

The frequency of communication is a polling interval called the heartbeat. The heartbeat and other communication configurations can affect speed and performance on your network.

You should set the heartbeat interval to correspond to other communication requirements. You can set a smaller interval if less data is exchanged between the client and the management server.

## Communication modes

You configure the Symantec Endpoint Protection clients to communicate with Symantec Endpoint Protection Manager using either pull mode or push mode. Each communication mode has advantages and the disadvantages that you need to assess for each environment.

For best performance, keep the database close to the management server and use pull mode. Symantec does not recommend push mode for large networks.

**Table 1-1** Pull mode versus push mode

Communication mode	Description
Pull mode	<p>The client connects to the management server according to the heartbeat, or polling interval. This procedure repeats indefinitely. The clients obtain any policy changes at the next check-in.</p> <p>The number of clients that pull mode can support depends on the following conditions:</p> <ul style="list-style-type: none"><li>■ Server performance</li><li>■ Network bandwidth for client communication</li><li>■ Server communication</li><li>■ Heartbeat frequency</li></ul> <p>In general, the less frequent the heartbeat, the more clients a management server can support. You are not limited to a maximum number of clients that can connect to a particular management server.</p> <p>Three management servers can support 55,000 clients, with an average of 18,000 clients per Symantec Endpoint Protection Manager (14.x).</p> <p>Use pull mode for remote sites.</p>

**Table 1-1** Pull mode versus push mode (*continued*)

Communication mode	Description
Push mode	<p>The client establishes a persistent TCP connection to the management server. The clients check in to the management server as soon as there are any policy changes and obtain the updated policies immediately.</p> <p>Push mode is more resource intensive than pull mode because of the persistent TCP connection.</p> <p>In push mode communication:</p> <ul style="list-style-type: none"> <li>■ The server notifies the client whenever the server changes status.</li> <li>■ Logs are sent from the client to Symantec Endpoint Protection Manager at the heartbeat interval.</li> </ul> <p>The maximum ratio of concurrently connected clients to management servers is 500:1 if Symantec Endpoint Protection Manager uses the HTTP protocol and 250:1 if Symantec Endpoint Protection Manager uses the HTTPS protocol.</p>

In 12.1.5 and later, by default Windows clients upload critical events immediately to the server. This setting bypasses the heartbeat interval. Critical events include any risk found (except cookies) and any intrusion event. To find this option, click **Clients > Policies > Communications Settings**.

[Configuring push mode or pull mode to update client policies and content](#)

## Recommended heartbeat intervals

The following performance measurement for the heartbeat interval was performed in a controlled environment.

**Table 1-2** Heartbeat interval

Number of clients	Number of management servers	SQL Server CPU	Shortest recommended heartbeat interval
55,000	3	Six cores (6 CPU)	60 minutes

This heartbeat interval does not include performance overhead such as site-to-site database replication or reporting activity. Other factors can adversely affect performance numbers, such as lower hardware specifications, available network bandwidth, or network congestion. These factors may require you to increase the heartbeat intervals to achieve the desired performance in your environment. The test data that is provided in this document is based on performance in a physical environment. Because of the nature of resource allocation in a virtual environment, add 25-30% more time to your calculation for the heartbeat interval setting.

This environment is configured as a single site that runs multiple management servers with the following specifications.

Symantec Endpoint Protection Manager:

- CPU: Intel Xeon Processor E5420 (2.5 GHz, Quad-Core)
- Physical memory (RAM): 8 GB
- Operating system: Microsoft Windows Server 2012 R2 Standard 64-bit

Microsoft SQL Server:

- CPU: Six-Core AMD Opteron Processor Model 2435 (2.6 GHz)
- Physical memory (RAM): 64 GB
- Operating system: Microsoft Windows Server 2012 R2 Standard 64-bit
- Database version: SQL Server 2012 SP2

## Calculating content distribution time

A key metric for provisioning an environment is the time it takes to distribute content updates to an organization. Content updates vary in size and frequency depending upon the type of content and their availability. Content updates can include virus definitions, intrusion prevention signatures, and Symantec Endpoint Protection engines updates, among other content types.

The time that is required to perform a content distribution update in a best-case scenario can be calculated with the following formula:

$(\text{Concurrent connections} \times \text{Average content size}) \div \text{Available bandwidth (in Kbps)} = \text{Content distribution time (in seconds)}$

Table 1-3 uses this formula, where the average content size value is 200 KB. In addition, the results assume that the entire network is dedicated to the update. In other words, the content download consumes all available bandwidth, and all clients download content concurrently.

**Table 1-3** Content distribution time with a 200-KB update

Bandwidth	Number of clients	Time
T1 (1.54 Mbps, or 192.5 Kbps)	5,000	90 minutes
	15,000	4.5 hours
10 Mbps (1250 Kbps)	5,000	14 minutes
	15,000	40 minutes
100 Mbps	5,000	80 seconds
	15,000	4 minutes

**Table 1-3** Content distribution time with a 200-KB update (*continued*)

Bandwidth	Number of clients	Time
1 Gbps	5,000	8 seconds
	15,000	24 seconds

Network utilization and protocol overhead also affect latency.

To decrease the time that is required to distribute content distribution updates, you can choose any of the following options:

- Distribute the client load across multiple management servers.
- Deploy Group Update Providers.
- Use alternative methods to distribute the content such as LiveUpdate servers or third-party distribution tools.

See [“Windows client installation package and content update sizes”](#) on page 24.

See [“About Group Update Providers”](#) on page 13.

---

**Note:** The data in this section was tested in version 12.1.5.

---

## Bandwidth control for client communication

In 12.1.5 and later, you can also adjust bandwidth for client communication by using the `mod_bw` Apache module. This module is now provided as part of the management server. `Mod_bw` can limit the number of clients that download content simultaneously from the management server. It can also specify the maximum network bandwidth that the server consumes to deliver content to clients.

Bandwidth control is disabled by default, and should be used with caution. It is useful, however, in the following scenarios:

- You have many roaming clients that connect intermittently to your network. These clients are more likely to require full content downloads or to send larger logs. Both these activities require more network bandwidth.
- You have multiple locations that connect to a single management server and updates clog your WAN network.
- Any scenario where you need to control bandwidth precisely to optimize network performance.

Sample settings are provided in the `bw.conf` configuration file for the following cases:

- Full content downloads

- Content delta downloads
- Client upgrade packages

To minimize network load, you can also randomize content downloads.

The configuration file for `mod_bw` can be found in the following location on the management server:

*Symantec Endpoint Protection Manager installation path*/apache/conf/bw/bw.conf

For more information on configuring the Apache `mod_bw` module, see:

[Symantec Endpoint Protection Bandwidth Control for Client Communication](#)

## About Group Update Providers

The Group Update Provider (GUP) provides updates to clients belonging to a group, and to any subgroups that are configured to inherit group policies. If you are concerned about multiple updates occurring across a given connection, consider deploying a GUP. A single GUP can support up to 10,000 clients. The throughput of the hardware for the designated GUP system dictates the performance of the GUP. Typically, a client that is configured as a GUP requires between 500 MB and 1 GB of disk space to store content updates. The number varies depending upon the age of the clients connecting for updates, and the size of the delta that is used to update them.

If you need to deploy updates to more than 10,000 clients, consider an alternative update method, such as:

- Multiple GUPs
- Additional installations of Symantec Endpoint Protection Manager
- An internal LiveUpdate server

[Configuring Group Update Providers](#)

[Choose a distribution method to update content on clients](#)

## About LiveUpdate servers

You may want to install LiveUpdate Administrator on an internal server if your environment meets the following criteria:

- You provide content updates to more than 10,000 clients
- You cannot deploy an additional Symantec Endpoint Protection Manager
- You cannot deploy multiple GUPs

LiveUpdate servers add very little overhead to an existing management server, and so have negligible effect on server resources.

To keep all clients up-to-date with the latest available content, configure one external LiveUpdate 2.x server with multiple distribution centers to remote sites throughout the network.

For more information about setting up an internal LiveUpdate server, see the *Symantec LiveUpdate Administrator User's Guide*.

## Management server and database sizing

Several factors influence the size of the Symantec Endpoint Protection Manager database and the storage space that is required on a management server. These factors include the following variables:

- Database maintenance settings  
See [“Optimizing database performance”](#) on page 15.
- Database backup requirements  
See [“Calculating backup storage requirements”](#) on page 16.
- Log size and expiration timeframes  
See [“Reducing log sizes”](#) on page 17.  
See [“Calculating log sizes”](#) on page 18.
- Content update sizes  
See [“Windows client installation package and content update sizes”](#) on page 24.
- Client installation package sizes

## Symantec Endpoint Protection Manager hardware and software recommendations

Hardware and software requirements vary depending on the number of clients that Symantec Endpoint Protection Manager manages.

Symantec makes the following recommendations for Symantec Endpoint Protection Manager hardware:

Symantec Endpoint Protection Manager that serves fewer than 10,000 clients:

- 8 GB RAM minimum
- Single processor

Symantec Endpoint Protection Manager that serves between 10,000 and 18,000 clients:

- 8 GB RAM minimum
- Dual processor

Symantec Endpoint Protection Manager that serves more than 18,000 clients:

- 64 GB RAM minimum
- Quad-core processor

For better performance, Symantec recommends using the following software:

- Windows Server 2008 or 2012 R2 x64-bit  
You can also obtain additional optimization for disk I/O performance on the Symantec Endpoint Protection Manager.
- SQL Server 2012 or later  
High-throughput hard drives with 10,000 rpm or higher drive speed.
- SAN environment with a management product such as Veritas Storage Foundation

## Choosing a database based on client-to-server ratio

For the following client-to-server ratios, Symantec recommends the following databases:

- For installations with a client-to-server ratio of  $\leq 5,000:1$ , using the default log settings, Symantec recommends the embedded database.
- For installations with a client-to-server ratio  $> 5,000:1$ , or that require large log sizes, Symantec recommends a separate SQL Server database. For SQL Server database sizing, see the database vendor's sizing tools or guides.

## Recommended database backup plans

Database backups create a copy of the database. In the event that data corruption or hardware failure occurs, you can revert to a previous copy of a backup to restore lost data.

Symantec recommends the following backup plans:

- SQL Server database only: Use the SQL Server Enterprise Manager to set up a maintenance plan that includes automatic backups.
- Embedded database or SQL Server database: Use the Symantec Endpoint Protection Manager console to perform on-demand backups and also schedule automatic backups. You can also use the Symantec Database Backup and Restore utility, which is automatically installed when you install Symantec Endpoint Protection Manager.

Create backups regularly and store them on a separate disk drive, preferably at a secure off-site facility.

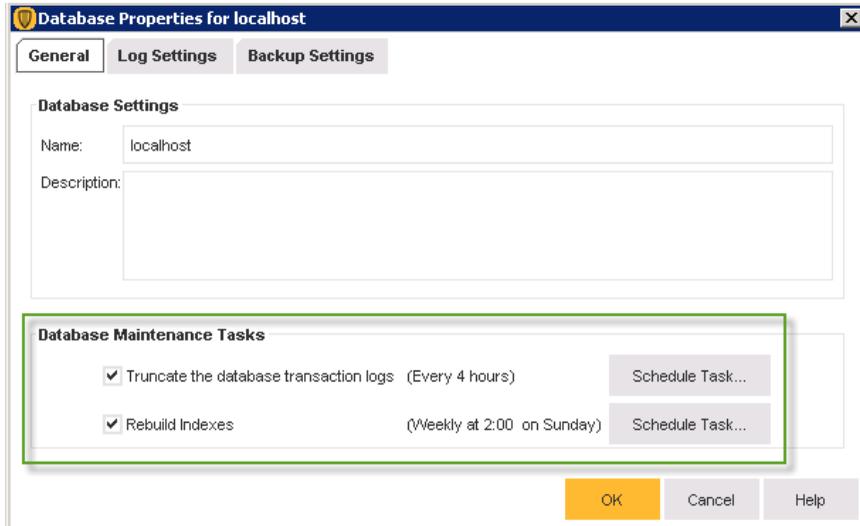
[Backing up the database and logs](#)

## Optimizing database performance

For either the embedded database or the SQL Server database, Symantec recommends that you periodically defragment and reindex the database. Reindexing can improve performance

and ensure an optimal database structure, particularly in large environments. Reindexing and defragmenting should be included as part of the regular database maintenance plan.

You enable these settings on the **Admin > Servers > localhost > Edit Database Properties > General** tab.



[Creating a maintenance plan in SQL Server 2012 or later to optimize database performance](#)

## Calculating backup storage requirements

The size and number of retained backups affect the required disk space on the Symantec Endpoint Protection Manager server.

For a new installation of Symantec Endpoint Protection Manager, the backup size is approximately 75% of the database size, which you multiply by the number of copies to retain:

3-GB database x 0.75 x 3 copies = 6.75 GB of disk space

---

**Note:** The data in this section was checked in version 12.1.5.

---

## Application learning and its effect on the database

Application learning lets Symantec Endpoint Protection clients report information and statistics about the executables that are run on them. This information is provided to Symantec Endpoint Protection Manager and aggregated into its database. The purpose of this information is to build a list of known applications in an environment. This list helps to create application-based

firewall rules and Host Integrity rules. It can also be used as a reference for developing Application Control rules and exceptions.

If the database is left to run indefinitely, it can grow considerably and eventually slow processing or cause other database problems. For this reason, Symantec recommends that you keep application learning turned off for the management servers that use the embedded database.

To turn off application learning, click **Clients > Policies > Communications**, and uncheck **Learn applications that run on the client computers**.

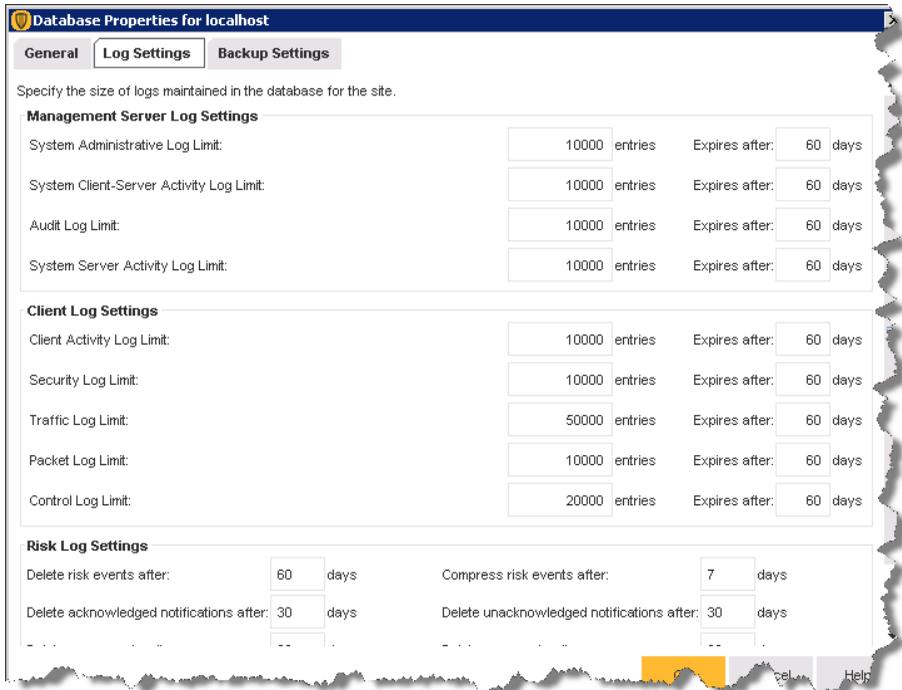
[Best Practices Guide to Application Learning in Symantec Endpoint Protection Manager](#)

## Reducing log sizes

You can configure database log options for the data that are stored in the database. The log options help you to manage the size of your database by specifying compression settings and how long to keep data.

Scheduled deletion of events ensures log entries are deleted regularly to prevent your database from growing too large. Event compression consolidates multiple "risk-found" events into a single security event. Over time, and especially during a security event, event compression can help keep the database size within manageable limits.

To configure these settings, click **Admin > Servers > localhost > Edit Database Properties > Log Settings** tab.



Specifying the log size and how long to keep log entries in the database

Specifying client log size and which logs to upload to the management server

## Calculating log sizes

You can configure the logs to optimize storage requirements and to comply with the company policies that control retention of logged data. The following parameters are commonly used to control logging activity:

- Maximum number of the entries that are stored in the logs
- Length of time in days to store log entries

Table 1-4 illustrates the key factors affecting log size and storage requirements.

**Table 1-4** Log size and storage requirements

Log	Size per 10,000-log entries (MB)
System Administrative	4
System Client-Server Activity	5

**Table 1-4** Log size and storage requirements (*continued*)

Log	Size per 10,000-log entries (MB)
Audit	10
System Server Activity	6
Client Activity	7
Security	23
Traffic	8
Packet	.5
Control	11

**Table 1-5** Approximate sizes of detected and quarantined virus events

Number of viruses in database	Approximate space (MB)
1,000	.28
5,000	1.4
15,000	4.3
25,000	7.2
45,000	13

The average database requirement for a 26,825-client deployment is roughly 15,000 detected and quarantined virus events every 60 days.

**Table 1-6** Example of log data statistics for a 26,825-client environment

Log	Average events per log
System Administrative	Usually very few
System Client-Server Activity	1 event per client per day
Audit	Usually very few
System Server Activity	240 events per server per day
Client Activity	10 events per endpoint per day
Security	0 events per client per day

**Table 1-6** Example of log data statistics for a 26,825-client environment (*continued*)

Log	Average events per log
Traffic	21 events per client per day
Packet	Can be large depending on policies
Control	Can be large depending on policies
Viruses	573 events per month for 1,000 clients

The log metric data varies for each customer.

---

**Note:** The data in this section was tested in version 12.1.5.

---

## Windows client installation packages

Symantec Endpoint Protection version 14 and later provides real-time protection for Windows clients by accessing definitions from the cloud (Intelligent Threat Cloud Service).

Symantec collects information about files from its global community of millions of users and its Global Intelligence Network. The collected information is available to Symantec products in the cloud through Symantec Insight. Symantec Insight provides a file reputation database and the latest virus and spyware definitions. Symantec products leverage Insight to protect client computers from new, targeted, and mutating threats. The data is sometimes referred to as being in the cloud since it does not reside on the client computer. Symantec Endpoint Protection must request or query Insight for information. The queries are called reputation lookups, cloud lookups, or Insight lookups.

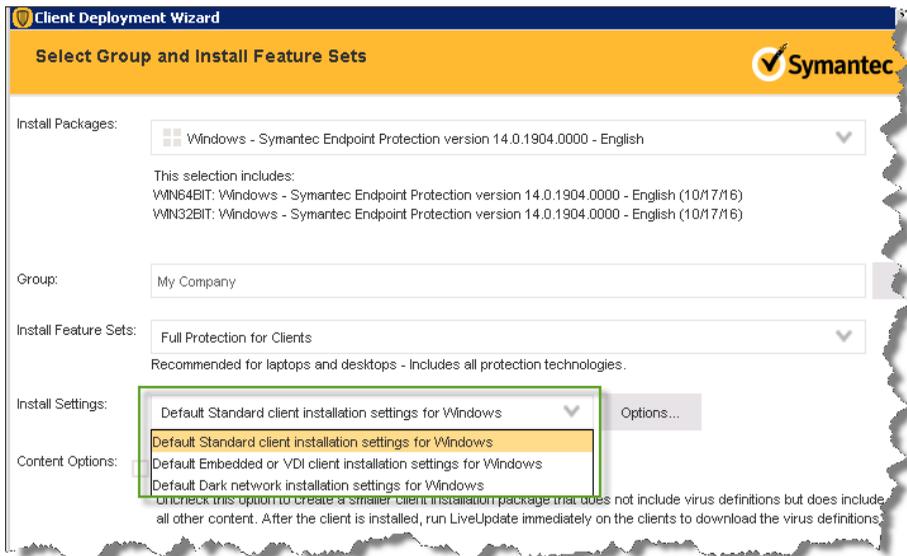
Symantec Endpoint Protection includes the following client installation packages to take advantage of the cloud:

- **Standard client** (cloud-based): Recommended for installations where clients have access to the cloud or the clients are version 12.1.6 and earlier. The standard client is 80% to 90% smaller than a dark network client installation package and includes the most recent virus definitions only. After installation, the client accesses the full set of virus definitions from the cloud.
- **Embedded client or VDI client** (cloud-based): Recommended for clients on embedded devices. The embedded client replaces the reduced-size client that was introduced in version 12.1.6. The embedded client is also 80% to 90% smaller than the standard client and also includes the most recent virus definitions only. After installation, the client accesses the full set of virus definitions from the cloud.

- Dark network client:** Recommended for installations where the client computers are in networks with no access to the cloud. The dark network installs a full set of virus definitions and keeps the definitions locally rather than accessing them from the cloud.
 

The 12.1.6.x standard client is similar to the dark network client, except that the 12.1.6.x client performs reputation lookups using the cloud for Download Insight.

All version 14 and later Windows client installation packages use the same technologies to detect threats.



For more information on the differences between the version 14.x and version 12.1.6.x clients, see:

[Comparison of Symantec Endpoint Protection Windows client type versions](#)

[About Windows client installation packages](#)

[Installing clients with Save Package](#)

## Benefits of cloud-based clients

The cloud-based clients have the following advantages over the dark network client:

- Faster to install**
- Use less disk space on the client computer**
- Maintain a lower false positive rate**

The Intelligent Threat Cloud Service's rapid scan capabilities makes it unnecessary to download all definitions to the client computer to maintain a high level of effectiveness. When the client

requires new definitions, the client downloads or looks up the definitions in the cloud for better performance and speed.

Therefore, Symantec recommends that you install the standard client and not the dark network client. The exception is when the client computers are located in a true dark network with no access to the cloud. Dark network clients must provide a larger set of definitions to improve the odds of getting a conviction in the absence of reputation lookup data for Insight.

The only disadvantage of the standard client or embedded client is that the content updates may take longer to download than the dark network client.

See [“Windows client installation package and content update sizes”](#) on page 24.

## Scan performance and bandwidth for cloud-based clients

Cloud lookups make queries to Symantec Insight for file reputation information and definition checking in the cloud. Therefore, the cloud-based scans take slightly longer, and take slightly more CPU consumption than dark network clients. However, cloud-base clients have a higher rate of malware detection and a lower false positive rate than scans on 12.1.6.x clients.

[Table 1-7](#) displays the test results for scans on a new installation of version 14.x and version 12.1.6.x.

**Table 1-7** Comparison of scan performance between versions 14.x and 12.1.6.x (time)

Scan type	Parameter	Standard (12.1.6.x)	Standard (14.x)	Embedded	Dark network
On-demand scan	Scan time	137 seconds	189 seconds	190 seconds	107 seconds
	CPU consumption	19%	20.6%	20.4%	18.6%
	Sent packet size	10 KB	231 KB	150 KB	N/A
	Received packet size	32 KB	192 KB	125 KB	N/A
	Number of reputation queries	16	1,635	1,635	N/A

**Table 1-7** Comparison of scan performance between versions 14.x and 12.1.6.x (time)  
*(continued)*

Scan type	Parameter	Standard (12.1.6.x)	Standard (14.x)	Embedded	Dark network
Scheduled scan	Scan time	173 seconds	218 seconds	211 seconds	138 seconds
	CPU consumption	16.3%	19.1%	19.6%	15%
	Sent packet size	10 KB	194 KB	239 KB	N/A
	Received packet size	30 KB	161 KB	201 KB	N/A
	Number of reputation queries	7	1,632	1,641	N/A
Auto-Protect scan	Scan time	1,252 seconds	1,486 seconds	1,453 seconds	935 seconds
	CPU consumption	22.8%	34.2%	34.5%	28%
	Sent packet size	236 KB	149 KB	171 KB	N/A
	Received packet size	465 KB	132 KB	150 KB	N/A
	Number of reputation queries	240	300	299	N/A

The on-demand scan time for dark network clients is faster on Windows 10 and longer on Windows 7. The Auto-Protect scan time for standard or embedded clients is slower on Windows 10 and faster on Windows 7.

These scans were tested in an environment with the following specifications:

- CPU: Intel Core i5-3470 CPU @ 3.20 GHz
- RAM: 2 GB
- Operating system: Windows 10 (64-bit)
- Network connection: 1 Gbps

The tests used a sample data set of 5.2 GB.

[How Symantec Endpoint Protection uses the Intelligent Threat Cloud Service](#)

# Windows client installation package and content update sizes

Client installation packages, product patches, and content updates are also stored in the Symantec Endpoint Protection database and affect the storage requirements. Product patches contain information for client packages and information for each language or locale. Note that patches also create new, full client builds.

[Table 1-8](#) displays the size of the client installation package if the maximum level of client logging and protection technologies are enabled.

**Table 1-8** Windows client installation package size

Client type/Definition type	*Installed with virus definitions?	64-bit package (MB)	32-bit package (MB)
Standard and Embedded (14) CoreDefs-3**	Yes	188	175
	No	93	81
Dark network (14) CoreDefs-1.5	Yes	288	276
	No	93	80
Standard (12.1.6) CoreDefs-1	Yes	335	316
	No	86	70
Reduced (Embedded/VDI) (12.1.6) CoreDefs-3	Yes	182	165
	No	86	70

For these packages, you can set a larger heartbeat. These sizes do not include packet-level firewall logs, which are not recommended in a production environment. If client logging is disabled, and there are no new policies or content to download from the management server, the client installation package is smaller. In this case, you can set a smaller heartbeat.

\* If your network has low bandwidth, install the client package without the virus definitions. As soon as the client connects to the management server, the client receives the full set of virus definitions.

All client installation packages include all features, such as Virus and Spyware, the firewall, the IPS, SONAR, System Lockdown, Application Control, Host Integrity content, and so forth. The difference between the client types are the size of the virus and spyware definitions.

In 12.1.5 and later, content updates require less storage space in the database and on the file system. Instead of storing multiple full revisions, the management server now stores only one full content revision plus incremental deltas. In 12.1.6, full content updates require ~470 MB.

---

**Note:** As of version 14, you can download security patches to clients the same way as other content, using a LiveUpdate server, the management server, or a Group Update Provider. In 12.1.6 and earlier, security patches are only available as part of new release, and only as part of a client deployment package using AutoUpgrade.

---

## Calculating total disk space requirements

The following scenario shows the disk space that Symantec Endpoint Protection requires for 26,466 clients. This example assumes the following metrics:

- An average 15,000 viruses over 60 days
- Retention of 20,000 events for each log
- Retention of five versions of each Symantec Endpoint Protection client (both 32-bit and 64-bit in English and French)
- Retention of seven backups

**Table 1-9** Disk space requirements

Item	Space required (MB)
15,000 detected and quarantined viruses	4.3
20,000 events per log	148
20 client versions	2400
Content updates	470

You must multiply the database size of 3.02 GB x 1.4 to account for the overhead of indexes and other tables in the database.

**The total database size = 4.23 GB**

In addition, to store seven backups on Symantec Endpoint Protection Manager, you would need approximately 29.6 GB of disk space.

---

**Note:** The data in this section was tested in version 12.1.5. This data would also apply to a version 14.x dark network client.

---