

Symantec™ Endpoint Protection 12.1 Sizing and Scalability Best Practices White Paper

Symantec Endpoint Protection Sizing and Scalability Best Practices White Paper

The software described in this book is furnished under a license agreement and may be used only in accordance with the terms of the agreement.

Product version: 12.1

Documentation version: 2

Legal Notice

Copyright © 2013 Symantec Corporation. All rights reserved.

Symantec, the Symantec Logo, the Checkmark Logo, and Altiris, LiveUpdate, Norton, Norton 360, and TruScan 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.

This Symantec product may contain third party software for which Symantec is required to provide attribution to the third party ("Third Party Programs"). Some of the Third Party Programs are available under open source or free software licenses. The License Agreement accompanying the Software does not alter any rights or obligations you may have under those open source or free software licenses. Please see the Third Party Legal Notice Appendix to this Documentation or TPIP ReadMe File accompanying this Symantec product for more information on the Third Party Programs.

The product described in this document is distributed under licenses restricting its use, copying, distribution, and decompilation/reverse engineering. No part of this document may be reproduced in any form by any means without prior written authorization of Symantec Corporation and its licensors, if any.

THE DOCUMENTATION IS PROVIDED "AS IS" AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE HELD TO BE LEGALLY INVALID. SYMANTEC CORPORATION SHALL NOT BE LIABLE FOR INCIDENTAL OR CONSEQUENTIAL DAMAGES IN CONNECTION WITH THE FURNISHING, PERFORMANCE, OR USE OF THIS DOCUMENTATION. THE INFORMATION CONTAINED IN THIS DOCUMENTATION IS SUBJECT TO CHANGE WITHOUT NOTICE.

The Licensed Software and Documentation are deemed to be commercial computer software as defined in FAR 12.212 and subject to restricted rights as defined in FAR Section 52.227-19 "Commercial Computer Software - Restricted Rights" and DFARS 227.7202, "Rights in Commercial Computer Software or Commercial Computer Software Documentation", as applicable, and any successor regulations. Any use, modification, reproduction release, performance, display or disclosure of the Licensed Software and Documentation by the U.S. Government shall be solely in accordance with the terms of this Agreement.

Symantec Corporation
350 Ellis Street
Mountain View, CA 94043
<http://www.symantec.com>

Technical Support

Symantec Technical Support maintains support centers globally. Technical Support's primary role is to respond to specific queries about product features and functionality. The Technical Support group also creates content for our online Knowledge Base. The Technical Support group works collaboratively with the other functional areas within Symantec to answer your questions in a timely fashion. For example, the Technical Support group works with Product Engineering and Symantec Security Response to provide alerting services and virus definition updates.

Symantec's support offerings include the following:

- A range of support options that give you the flexibility to select the right amount of service for any size organization
- Telephone and/or Web-based support that provides rapid response and up-to-the-minute information
- Upgrade assurance that delivers software upgrades
- Global support purchased on a regional business hours or 24 hours a day, 7 days a week basis
- Premium service offerings that include Account Management Services

For information about Symantec's support offerings, you can visit our website at the following URL:

www.symantec.com/business/support/

All support services will be delivered in accordance with your support agreement and the then-current enterprise technical support policy.

Contacting Technical Support

Customers with a current support agreement may access Technical Support information at the following URL:

www.symantec.com/business/support/

Before contacting Technical Support, make sure you have satisfied the system requirements that are listed in your product documentation. Also, you should be at the computer on which the problem occurred, in case it is necessary to replicate the problem.

When you contact Technical Support, please have the following information available:

- Product release level

- Hardware information
- Available memory, disk space, and NIC information
- Operating system
- Version and patch level
- Network topology
- Router, gateway, and IP address information
- Problem description:
 - Error messages and log files
 - Troubleshooting that was performed before contacting Symantec
 - Recent software configuration changes and network changes

Licensing and registration

If your Symantec product requires registration or a license key, access our technical support Web page at the following URL:

www.symantec.com/business/support/

Customer service

Customer service information is available at the following URL:

www.symantec.com/business/support/

Customer Service is available to assist with non-technical questions, such as the following types of issues:

- Questions regarding product licensing or serialization
- Product registration updates, such as address or name changes
- General product information (features, language availability, local dealers)
- Latest information about product updates and upgrades
- Information about upgrade assurance and support contracts
- Information about the Symantec Buying Programs
- Advice about Symantec's technical support options
- Nontechnical presales questions
- Issues that are related to CD-ROMs, DVDs, or manuals

Support agreement resources

If you want to contact Symantec regarding an existing support agreement, please contact the support agreement administration team for your region as follows:

Asia-Pacific and Japan	customercare_apac@symantec.com
Europe, Middle-East, and Africa	semea@symantec.com
North America and Latin America	supportsolutions@symantec.com

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](#)
- [Application learning and its impact to the database](#)
- [Management server and database sizing](#)
- [Preventing and correcting false positive detections with Symantec Endpoint Protection](#)

About sizing and scalability for Symantec Endpoint Protection

Symantec Endpoint Protection provides best-of-breed endpoint security to enterprises of all sizes. Several decisions factor into correctly sizing and deploying the Symantec Endpoint Protection environment for optimum protection and serviceability. This white paper provides the following information:

- Detailed recommendations for single- and multiple-site environments
- Client-to-server ratios

- Database sizing recommendations
- Log-keeping and maintenance

The architecture, designs, and recommendations 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 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 ongoing evaluation and should not be considered as firm commitments by Symantec.

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 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 Symantec Endpoint Protection 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
- Frequently gathered metrics
- Who and where people are that need access to the data
- Whether multiple administrative groups exist within the organization (such as groups for IT, security, desktops, or servers)

- 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

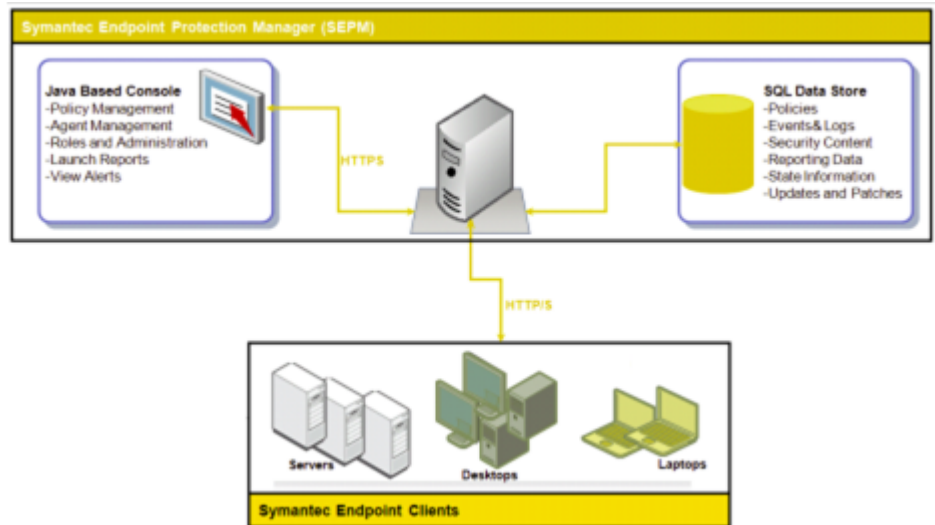
Effective and efficient endpoint security requires a balance of protection technologies, a manageable infrastructure, and adequate forensic data to properly monitor network security activities. Before the solution is deployed, several decisions need to be made about the best ways to configure the Symantec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

Architecture

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

- Symantec Endpoint Protection Manager (“management server”) – The management server that is used to configure clients, reports, and alerts.
- Symantec Endpoint Protection database (“database”) – The Microsoft SQL Server database or embedded database that stores all configuration, updates, and reporting information.
- Symantec Endpoint Protection Manager console (“console”) – A lightweight user interface that is used to access the Symantec Endpoint Protection Manager. The Symantec Endpoint Protection Manager console is used to manage and view deployment activity, configurations, updates, and Symantec Endpoint Protection client reports.
- Symantec Endpoint Protection client (“client”) – Software that is deployed to networked computers. The client is used to monitor policies and automate policy compliance activities.

Figure 1-1 Basic Symantec Endpoint Protection architecture




Site design

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

Single-site design

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

- Two Symantec Endpoint Protection Managers (for redundancy and load balancing)
- Database clustering (to support high availability)



Multiple-site design


An organization with more than one datacenter or with multiple large physical locations should use a multiple-site design. There are 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 bi-directional 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 at each remote site.



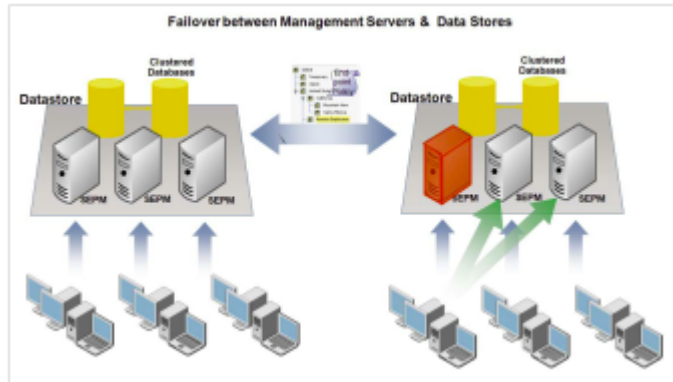
Central logging

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 logs forwarded from corporate sites 1 and 2.



High availability

The high availability (HA) 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 instance, the high availability design can be configured to have client computers automatically switch to an alternate Symantec Endpoint Protection Manager should the primary server become 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. Chief among the parameters that affect the client-to-server ratio are client-server communication, desired update speeds, and the security technologies deployed in the network environment.

Client-server communication

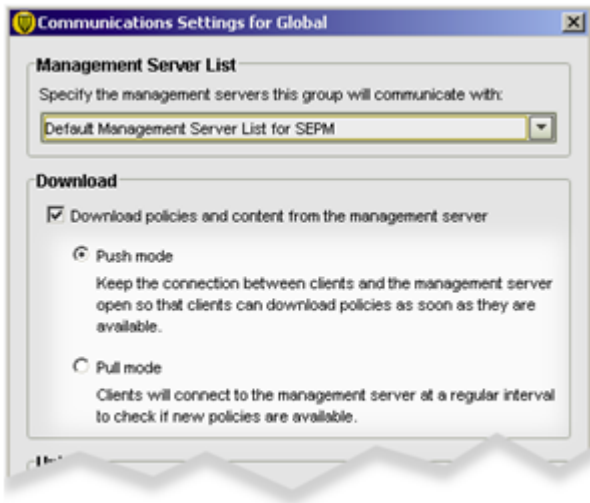
Symantec Endpoint Protection clients and management servers exchange status information and content data. Clients initiate this communication with the Symantec Endpoint Protection Manager from an ephemeral port to the management server on TCP port 8014 (or 443 if using SSL). In the event of a conflict, this port is configurable. The frequency of communication depends on the heartbeat (also called a "polling interval") and communication configuration.

The size of the client heartbeat depends on how much data is exchanged between the client and the management server. If client logging is not enabled and there are no new policies or content to download from the management server, the size of the client heartbeat is between 3 KB and 5 KB. If the maximum level of client logging is enabled and all the client protection technologies are enabled, the size of the client heartbeat is between 200 KB and 300 KB. This number does not include packet-level firewall logs, which is not recommended in a production environment.

Communication modes

Symantec Endpoint Protection clients can be configured to communicate with the Symantec Endpoint Manager using either push mode or pull mode. For best

performance, keep the database close to the management server and use pull mode.



Each communication mode has advantages and disadvantages that need to be assessed for each environment.

Table 1-1 Push mode versus pull mode

Communication mode	Description
Push mode	<p>The client establishes a persistent TCP connection to the management server. If a client cannot connect to the management server, it retries periodically, depending on the heartbeat frequency. The following conditions apply to push mode communication:</p> <ul style="list-style-type: none">■ The server notifies the client whenever the server changes status.■ Logs are sent from the client to the Symantec Endpoint Protection Manager at the heartbeat interval.■ Push mode is more resource intensive than pull mode because of the persistent TCP connection. <p>In push mode, the theoretical maximum ratio of clients to management servers is 50,000:1. However, Symantec generally recommends a maximum ratio of 5000:1 for push mode.</p>

Table 1-1 Push mode versus pull mode (*continued*)

Communication mode	Description
Pull mode	<p>The client connects to the management server according to the heartbeat frequency. This procedure repeats indefinitely. The number of clients that can be supported in pull mode depend on the following conditions:</p> <ul style="list-style-type: none">■ Server performance■ Network bandwidth used for clients■ Server communication■ Heartbeat frequency <p>In general, the less frequent the heartbeat, the more clients a management server can support. There is no maximum number of clients that can connect to a particular management server.</p>

Heartbeat interval

The performance figures provided in the following tables are based on testing performed in a controlled environment using servers with the following specifications and configured as a single site. All times are measured in minutes.

Symantec Endpoint Protection Manager:

- CPU: Intel Core2 Duo E6600, 2.40 GHz
- Physical memory (RAM): 4 GB
- Operating system: Microsoft Windows Server 2008, 64-bit

Microsoft SQL Server:

- CPU: Intel Xeon E5420 2.5 GHz (for CPU core performance)
- Physical memory (RAM): 64 GB
- Operating system: Microsoft Windows Server 2008, 64-bit
- Database version: Microsoft SQL Server 2008

Table 1-2 5,000 clients

SQL Server CPU	1 Symantec Endpoint Protection Manager	2 Symantec Endpoint Protection Managers
Single core	50	30
Dual core	20	15
Quad core	15	10

Table 1-2 5,000 clients (*continued*)

SQL Server CPU	1 Symantec Endpoint Protection Manager	2 Symantec Endpoint Protection Managers
2x quad core	10	10

Table 1-3 15,000 clients

SQL Server CPU	1 Symantec Endpoint Protection Manager	2 Symantec Endpoint Protection Managers	3 Symantec Endpoint Protection Managers
Dual core	50	35	25
Quad core	20	15	10
2x quad core	20	10	10

Table 1-4 25,000 clients

SQL Server CPU	1 Symantec Endpoint Protection Manager	2 Symantec Endpoint Protection Managers	3 Symantec Endpoint Protection Managers	4 Symantec Endpoint Protection Managers	5 Symantec Endpoint Protection Managers
Dual core	85	55	45	35	30
Quad core	30	20	20	15	10
2x quad core	30	20	15	10	10

Table 1-5 50,000 clients

SQL Server CPU	2 Symantec Endpoint Protection Managers	3 Symantec Endpoint Protection Managers	4 Symantec Endpoint Protection Managers	5 Symantec Endpoint Protection Managers
Dual core	80	65	50	40
Quad core	30	25	20	15
2x quad core	25	20	15	10

Table 1-6 100,000 clients

SQL Server CPU	2 Symantec Endpoint Protection Managers	3 Symantec Endpoint Protection Managers	4 Symantec Endpoint Protection Managers	5 Symantec Endpoint Protection Managers
Quad core	50	40	35	30
2x quad core	40	35	30	20

The heartbeat intervals listed in the tables do not include the performance overhead introduced by actions such as site-to-site database replication or reporting activity. Other factors such as lower hardware specifications, available network bandwidth, or network congestion can adversely affect performance numbers and require you to increase the heartbeat intervals to achieve the desired performance in your environment. Test data provided in this document is based on performance in a physical environment. Due to the nature of resource allocation in a virtual environment, you should add 25-30% more time to your calculation for the heartbeat interval setting.

Heartbeat sizing example

The following table assumes the Symantec Endpoint Protection Manager specifications shown in the previous tables.

Table 1-7

Number of total clients	Number of Symantec Endpoint Protection Managers	SQL Server CPU	Shortest recommended heartbeat interval
50,000	2 (25,000 clients per management server)	Dual core (2 CPU)	80 minutes
50,000	2 (25,000 clients per management server)	Quad core (4 CPU)	30 minutes

Calculating content distribution time

A key metric for provisioning a Symantec Endpoint Protection environment is the time it takes to distribute content updates to an organization. Content updates can include virus definitions and Intrusion Prevention signatures, among other content types.

Content updates vary in size and frequency depending upon content types and content update availability. The time required to perform a content distribution update in a best-case scenario can be calculated with the following formula:

*Concurrent connections X Average content size ÷ Available bandwidth = Content distribution time**

Where *Average content size* = 250-500 KB

*Latency is also affected by network utilization and protocol overhead.

To decrease the time that is required to distribute content distribution updates, do one or more of the following tasks:

- 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 [“Size of client installation packages and content updates”](#) on page 25.

About Group Update Providers

The Group Update Provider (GUP) provides updates to clients belonging to a group, and any subgroups that are configured to inherit group policies as set on the Clients tab. If you are concerned about multiple updates occurring across a given connection, then administrators should consider deploying a GUP. The GUP can support up to 10,000 clients. The performance of the GUP is dictated by the throughput of the hardware for the system that is designated as a GUP. Typically, a client 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 created to update them.

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

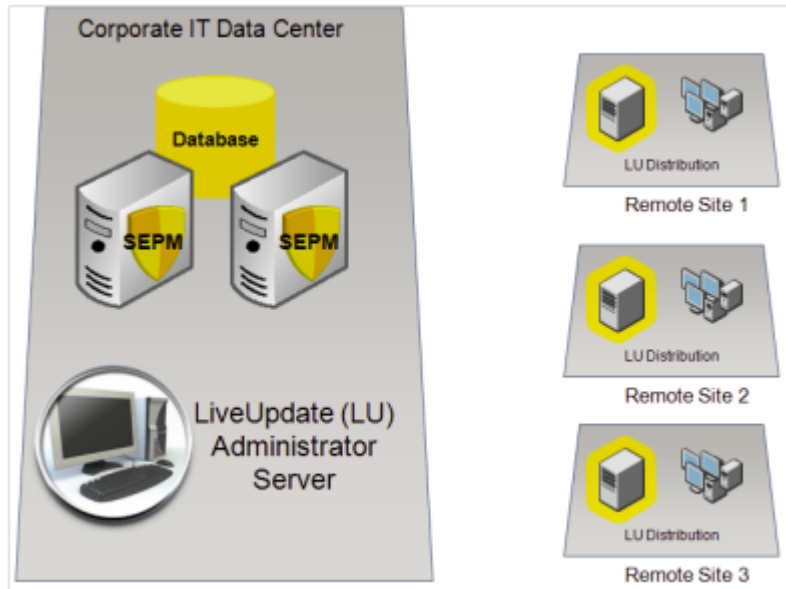
- Additional GUPs
- Additional Symantec Endpoint Protection Managers
- A Symantec LiveUpdate server

For more information about GUPs, see: [Configuring Group Update Providers](#)

About LiveUpdate servers

Environments that provide content updates to more than 10,000 clients, but cannot provide an additional Symantec Endpoint Protection Manager or additional GUPs can use a Symantec LiveUpdate server.

The following graphic depicts an architecture design that uses a LiveUpdate server to provide content updates.



This design features a LiveUpdate server at the Corporate Site (HQ) which redistributes content to LiveUpdate servers at each remote site. LiveUpdate servers act as simple re-distribution points and use HTTP, FTP, or network shares to distribute content updates. LiveUpdate servers add very little overhead to an existing server, and so have negligible impact on server resources.

For a complete list of hardware and software requirements for the Symantec Endpoint Protection components, see:

[Release Notes and System Requirements for all versions of Symantec Endpoint Protection and Symantec Network Access Control](#)

Application learning and its impact to the database

Application learning allows Symantec Endpoint Protection clients to report information and statistics about the executables that are run on them. This information is provided to the Symantec Endpoint Protection Manager and aggregated into the Symantec Endpoint Protection Manager database. The purpose of this information is to build a list of known applications in an environment to create application-based firewall rules and Host Integrity rules and can be used as a reference for developing Application Control rules and exceptions.

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

For more information, see:

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

Management server and database sizing

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

- Database maintenance settings
- Log size and expiration timeframes
- Content update sizes
- Client installation package sizes
- Backup information requirements

See [“Database maintenance settings”](#) on page 20.

See [“Logging options and sizes”](#) on page 21.

See [“Size of client installation packages and content updates”](#) on page 25.

See [“Database backups”](#) on page 24.

Database maintenance settings

You can configure database maintenance options for the data that are stored in the database. Database maintenance 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.

Database Properties for localhost

General Log Settings Backup Settings

Specify the size of logs maintained in the database for the site.

Management Server Log Settings

System Administrative Log Limit: 10000 entries Expires after: 60 days

System Client-Server Activity Log Limit: 10000 entries Expires after: 60 days

Audit Log Limit: 10000 entries Expires after: 60 days

System Server Activity Log Limit: 10000 entries Expires after: 60 days

Client Log Settings

Client Activity Log Limit: 10000 entries Expires after: 60 days

Security Log Limit: 10000 entries Expires after: 60 days

Traffic Log Limit: 50000 entries Expires after: 60 days

Packet Log Limit: 10000 entries Expires after: 60 days

Control Log Limit: 20000 entries Expires after: 60 days

Risk Log Settings

Delete risk events after: 60 days

Compress risk events after: 7 days Delete compressed events after: 7 days

Delete acknowledged notifications after: 30 days Delete unacknowledged notifications after: 30 days

Delete scan events after: 30 days Delete commands after: 30 days

☒ Delete unused virus definitions ☒ Delete EICAR events

Logging options and sizes

Logging options are configured by the administrator to optimize storage requirements and comply with company policies that control retention of logged data. The following parameters are commonly used to control logging activity:

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

The following examples illustrate the key factors affecting log size and storage requirements:

Table 1-8 Log size and storage requirements

Log	Size per 10,000 log entries (MB)
System Administrative	10
System Client-Server Activity	9
System Enforcer	6
Audit	6
System Server Activity	66

Table 1-8 Log size and storage requirements (*continued*)

Log	Size per 10,000 log entries (MB)
Client Activity	45
Security	45
Traffic	45
Packet	45
Control	45
Enforcer Client	16
Enforcer Server	14
Enforcer Traffic	9

Table 1-9 Approximate detected/quarantined virus event sizes

Number of viruses in database	Approximate space (MB)
1,000	0.8
5,000	4.3
15,000	12.9
25,000	21.6
50,000	43.2

The average database requirement for a 17,000-client deployment is roughly 15,000 detected and quarantined virus events every 60 days.

Table 1-10 Example of log data statistics for a 17,000-client environment*

Log	Size per 10,000 log entries (MB)
System Administrative	10 events per day per admin
System Client-Server Activity	9 events per day per machine
Audit	Usually very small
System Server Activity	650 events per server per day
Client Activity	120 events per machine per day

Table 1-10 Example of log data statistics for a 17,000-client environment*
(continued)

Log	Size per 10,000 log entries (MB)
Security	1 event per day per machine
Traffic	2400 events per machine per day
Packet	Could be extremely large depending on policies
Control	Could be extremely large depending on policies
Viruses	250 per month per 1000 clients

*Log metric data varies from customer to customer

Symantec Endpoint Protection Manager hardware recommendations

Hardware requirements vary depending on the number of clients served by the Symantec Endpoint Protection Manager. Symantec makes the following recommendations for Symantec Endpoint Protection Manager hardware:

Symantec Endpoint Protection Managers serving less than 10,000 clients:

- 2 GB RAM minimum
- Single processor

Symantec Endpoint Protection Managers serving more than 10,000 clients:

- 4 GB RAM minimum
- Dual processor

Database recommendations

For installations with a client-to-server ratio of 5,000 clients or less, using the default log settings, Symantec recommends using the embedded database. For installations with a client-to-server ratio greater than 5,000 clients, or using a higher level of log settings, Symantec recommends using a separate Microsoft SQL Server database. For SQL database sizing, Symantec recommends using the database vendor's recommended sizing tools or guides.

Database performance enhancement recommendations

For added performance, Symantec recommends using the following options:

- Microsoft Windows Server 2003 or 2008 64-bit or later
- Microsoft SQL Server 2005 64-bit or later
High-throughput hard drives with 10,000 RPM or higher drive speed.
- Use a SAN environment with a management product such as Symantec Storage Foundation

Additional optimization can be obtained for disk I/O performance on the Symantec Endpoint Protection Manager.

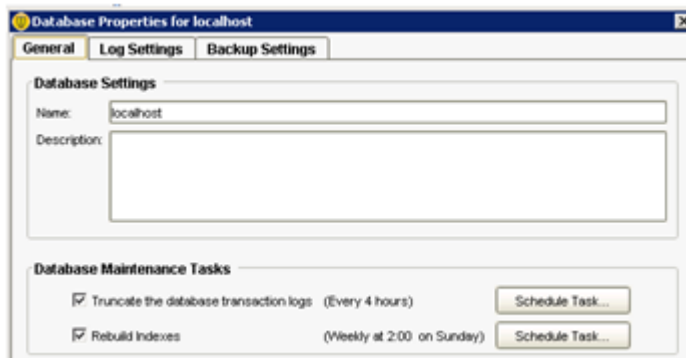
Database backups

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. A database backup is created using the Symantec Endpoint Protection Manager console or by using the Symantec Database Backup and Restore utility. The Symantec Database Backup and Restore utility is automatically installed during the installation of Symantec Endpoint Protection Manager.

For more information, see the *Symantec Endpoint Protection and Symantec Network Access Control Installation and Administration Guide*.

Recommended SQL database maintenance

If you use a Microsoft 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.



For more information, see the Symantec Support Knowledge Base article:

Create database maintenance plans in MS SQL Server 2005 using SQL Server Integration Services (SSIS)

Recommended backup plans

Symantec recommends the following backup plans:

- Microsoft SQL database only: Use the Microsoft SQL Server Enterprise Manager to set up a maintenance plan that includes automatic backups.
- Embedded or a Microsoft SQL database: Use the Symantec Endpoint Protection Manager console to perform on-demand backups and also schedule automatic backups.

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

Backup storage requirement calculation

The size and number of retained backups affect the required disk space on the Symantec Endpoint Protection Manager server. The backup size is approximately 75% of the database size multiplied by the number of copies being retained:

4-GB database x 0.75 x 3 *copies* = 9.2 GB of disc space

Size of client installation packages and content updates

Client update packages, patches, and content updates are also stored in the Symantec Endpoint Protection database and affect the storage requirements. Product updates and patches contain information for client packages and information for each language or locale. Note that patches also create new, full client builds. Full client builds are approximately 120 MB.

Full content updates require approximately 470 MB each. These updates contain information for the following technologies:

- Virus definitions
- SONAR definitions
- Intrusion Prevention signatures
- Symantec whitelist
- Centralized reputation settings
- Revocation data
- Submission control thresholds

- Auto-Protect portal list
- Extended file attributes and signatures
- Host Integrity definitions (Symantec Network Access Control)

See [“Calculating content distribution time”](#) on page 17.

Calculating total disk space requirements

The following scenario shows the space required for a Symantec Endpoint Protection implementation with 17,000 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 1-11 Disk space requirements

Item	Space required (MB)
15,000 viruses detected/quarantined	12.9
20,000 events per log	722
20 client versions	2400
20 client versions	300

Total database size = 4.80 GB

*The database size of 3.43 GB must be multiplied by 1.4 to account for the overhead of indexes and other tables in the database.

The space required on the management server to store seven backups is approximately 33.6 GB.

Preventing and correcting false positive detections with Symantec Endpoint Protection

Security technology has always had the potential of identifying a good file as bad. Symantec works hard to balance this risk of a false positive versus the need for aggressive detection. Symantec Endpoint Protection adds many features to reduce

the risk of false positive detections (FPD). Among these are security technologies that require some machine-learning to avoid false positives. As a result, during the initial testing and implementation of Symantec Endpoint Protection, some customers may initially see higher rates of FPD.

However, any disruption to the organization due to FPD can be avoided with some simple precautions. Additionally, improvements to the false positive workflow enable a false positive to be corrected quickly and with minimum disruption.

Symantec Insight technology uses file reputation data to detect and block malicious code from running on protected computers. SONAR uses heuristics as well as file reputation data.

[Table 1-12](#) describes the three ranges of threshold sensitivity associated with Insight and the use cases appropriate for each.

Table 1-12 Use cases for Insight threshold sensitivity

Range	Description
1-3	Appropriate for scenarios or for test environments that cannot tolerate false positive detections or the conviction of good files that still build reputation. At these levels, malware that still builds reputation may evade detection, but the system is very unlikely to convict good files.
4-6	Appropriate for most desktop users running normal software. This range balances the false positive risk and detection to capture most malware with low FPs. Based on Symantec's experience and understanding of the threat landscape, level 5 is the appropriate threshold for the majority of the users. Symantec discourages users from changing the value unless advised by Symantec Technical Support personnel.
7-9	Appropriate for highly secure environments where you want to lock down a server or desktop that does not frequently install new or unproven software. False positive detections occur at this level, but very little malware evades detection.

See [“Preventing false positive detections”](#) on page 27.

Preventing false positive detections

Symantec Endpoint Protection does not detect known good files as malware. You can help make sure that your good files are known as good. Use the following steps to help prevent false positives when after you install Symantec Endpoint Protection.

Table 1-13 Steps to prevent false positives

Step	Action	Description
Step 1	Check the digital signatures	<p>One of the easiest ways to identify that a file is good is to know where it came from and who created it. One of the most important factors to build a positive file reputation is to check its digital signature. Executable files without a digital signature are at risk of being identified as unknown.</p> <ul style="list-style-type: none"> ■ Custom or home-grown applications should be digitally signed with class three digital certificates. ■ Customers should insist that their software vendors digitally sign their applications.
Step 2	Add files to the Symantec whitelist	<p>Symantec has a growing whitelist of over 25 million good files. These files are used to test signatures before the signatures are released to the public. The hash values of these files are also stored in the cloud and are used as a real-time check to avoid FPD on the Symantec Endpoint Protection client. Whitelisting files provides a powerful method to avoid FPD. Customers and vendors can add files to this list.</p> <ul style="list-style-type: none"> ■ Software vendors: Submit executables for inclusion to the Symantec whitelist: https://submit.symantec.com/whitelist/ ■ BCS customers: Submit system images to the Symantec whitelist: https://submit.symantec.com/whitelist/bcs.cgi <p>Note: Do not use these websites to correct a false positive. Instead, use the tools that are available through the whitelisting program to help simplify the submission of files.</p> <p>See “Correcting false positive detections” on page 30.</p>
Step 3	Test the level of false positives by monitoring for potential issues	<p>The initial deployment of Symantec Endpoint Protection during testing should include test computers with representative images of the software you run in your environment, including common third-party applications.</p>
Step 4	Log SONAR high risk heuristic detections and use application learning	<p>For SONAR, set detection action for high risk heuristic detections to Log for a short period of time. Let application learning run for the same period of time. Symantec Endpoint Protection learns the legitimate processes that you run in your network.</p> <p>After the period of time, you should set the detection action back to Quarantine.</p> <p>If you use aggressive mode for low risk heuristic detections, you increase the likelihood of false positive detections.</p>

Table 1-13 Steps to prevent false positives (*continued*)

Step	Action	Description
Step 5	Exclude good files from detection	<p>Symantec Endpoint Protection provides multiple methods to exclude good files from detection. Exclusions can be added from within the Symantec Endpoint Protection Manager console to provide false positive mitigation on the client.</p> <p>See “Excluding good files from detections” on page 32.</p>
Step 6	Send feedback to Symantec through the automated submission of metadata on detections	<p>Each security technology in Symantec Endpoint Protection can collect and send file metadata to Symantec. These submissions are used to identify and mitigate false positives through forensic analysis, heuristic training against collected data sets, and custom and generic whitelisting.</p> <p>Automated submission is enabled by default. You can specify the types of detections for which clients submit information. In the console, click Clients > Policies > External Communications Settings > Submissions tab.</p>

See [“Correcting false positive detections”](#) on page 30.

Adjusting SONAR settings on your client computers

You might want to change the SONAR actions to reduce the rate of false positive detections. You might also want to change the SONAR actions to change the number of detection notifications that appear on your client computers.

Note: The settings for SONAR notifications are also used for TruScan proactive threat scan notifications.

To adjust SONAR settings on your client computers

- 1 In the Virus and Spyware Protection policy, select **SONAR**.
- 2 Make sure that **Enable SONAR** is checked.
- 3 Under **Scan Details**, change the actions for high or low risk heuristic threats.

You can enable aggressive mode for low risk detections. This setting increases SONAR sensitivity to low risk detections. It might increase the false positive detections.
- 4 Optionally change the settings for the notifications that appear on your client computers.

The SONAR settings also control notifications for TruScan proactive threat scans.

- 5
- Under **System Change Events**, change the action for either **DNS change detected** or **Host file change detected**.

Note: The **Prompt** action might result in many notifications on your client computers. Any action other than **Ignore** might result in many log events in the console and email notifications to administrators.

Warning: If you set the action to **Block**, you might block important applications on your client computers.

For example, if you set the action to **Block** for **DNS change detected**, you might block VPN clients. If you set the action to **Block** for **Host file change detected**, you might block your applications that need to access the host file. You can use a DNS or host file change exception to allow a specific application to make DNS or host file changes.

- 6
- Under **Suspicious Behavior Detection**, change the action for high or low risk detections.
- 7
- Click **OK**.

Correcting false positive detections

- During the testing, Symantec wants to know about false positive detections on customer systems for the following reasons:
- To identify the causes of false positive detections.
 - To make adjustments to the detection subsystem that reduces the number of future false positive detections.

Table 1-14 Steps to correct false positives

Step	Action	Description
Step 1	Monitor SONAR detection results	Check the SONAR log to determine which processes are legitimate and which are security risks. See “Monitoring SONAR detection results to check for false positives” on page 31.

Table 1-14 Steps to correct false positives (*continued*)

Step	Action	Description
Step 2	Submit the false positive detection by using a Symantec Web form	<p>For any suspected false positive detections, make submissions to: https://submit.symantec.com/false_positive/</p> <p>Warning: It is critical for resolution of false positives that you include the SHA256 value of the file with the submission.</p> <p>The hash value of a file is presented in a notice on the client. Third-party tools are also available.</p> <p>Once the submission has been processed and Symantec whitelists the file, the quarantine rescan feature automatically restores files detected as heuristic false positives.</p>
Step 3	Exclude good files from detection	<p>Symantec Endpoint Protection Manager provides multiple methods to exclude good files from detection. Exclusions can be added from within the Symantec Endpoint Protection Manager console to provide false positive mitigation on the client.</p> <p>See “Excluding good files from detections” on page 32.</p>

Monitoring SONAR detection results to check for false positives

The client collects and uploads SONAR detection results to the management server. The results are saved in the SONAR log.

To determine which processes are legitimate and which are security risks, look at the following columns in the log:

Event	<p>The event type and the action that the client has taken on the process, such as cleaning it or logging it. Look for the following event types:</p> <ul style="list-style-type: none"> ■ A possible legitimate process is listed as a Potential risk found event. ■ A probable security risk is listed as a Security risk found event.
Application	The process name.
Application type	The type of malware that SONAR or a TruScan proactive threat scan detected.
File/Path	The path name from where the process was launched.

The **Event** column tells you immediately whether a detected process is a security risk or a possible legitimate process. However, a potential risk that is found may

or may not be a legitimate process, and a security risk that is found may or may not be a malicious process. Therefore, you need to look at the **Application type** and **File/Path** columns for more information. For example, you might recognize the application name of a legitimate application that a third-party company has developed.

Legacy clients do not support SONAR. Legacy clients collect similar events from TruScan proactive threat scans, however, and include them in the SONAR log.

To monitor SONAR detection results to check for false positives

- 1 In the console, click **Monitors > Logs**.
- 2 On the Logs tab, in the **Log type** drop-down list, click **SONAR**.
- 3 Select a time from the **Time range** list box closest to when you last changed a scan setting.
- 4 Click **Advanced Settings**.
- 5 In the **Event type** drop-down list, select one of the following log events:
 - To view all detected processes, make sure **All** is selected.
 - To view the processes that have been evaluated as security risks, click **Security risk found**.
 - To view the processes that have been evaluated and logged as potential risks, click **Potential risk found**.
- 6 Click **View Log**.
- 7 After you identify the legitimate applications and the security risks, create an exception for them in an Exceptions policy.

You can create the exception directly from the SONAR Logs pane.

Excluding good files from detections

You can exclude a known-good application from being detected in the following ways:

- A known-good file can appear in the Risk log or the SONAR log as a false positive. You can configure log settings to allow the application to be ignored, which prevents it from appearing in the Risk log or SONAR log.
You can select more than one application, file, URL, or IP address at a time.
- You can add and manage exceptions in the Exceptions policy on the **Client Restrictions** page.
- You can add a trusted Web domain in the Exceptions policy on the **Exceptions** page to exclude it from being detected.

Figure 1-2

In the Symantec Endpoint Protection Manager Monitors page, add exclusions or exceptions for critical files, directories, URLs, and IP addresses

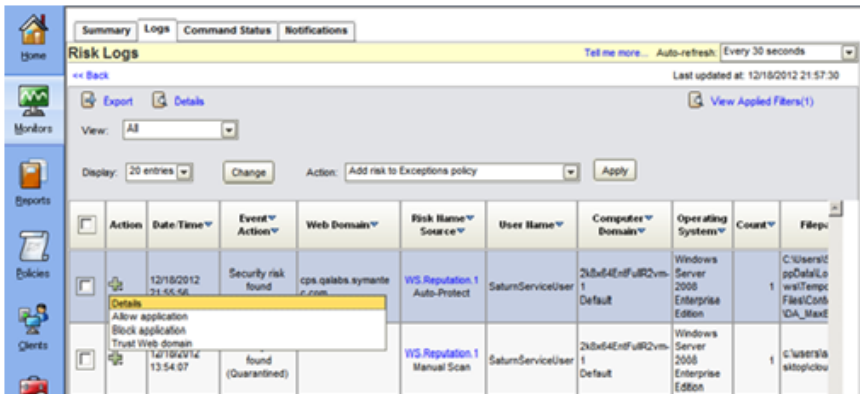
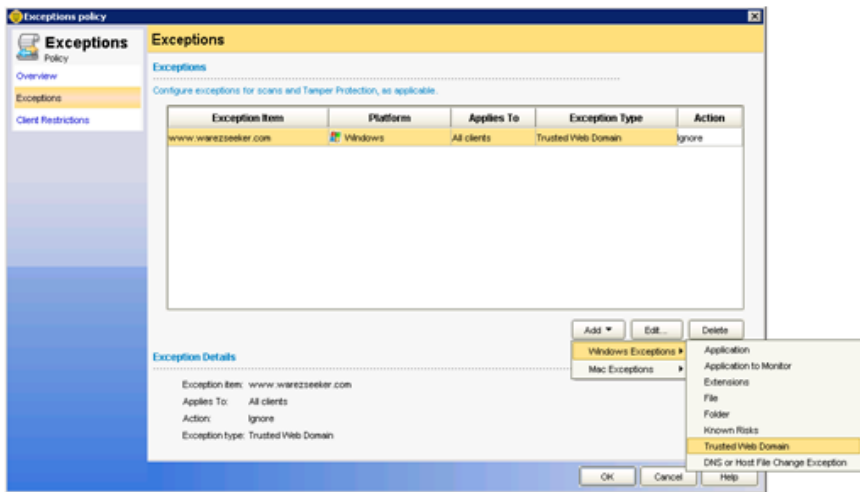


Figure 1-3

Exclude a trusted Web domain from Insight detections



For more information, see the *Symantec Endpoint Protection and Symantec Network Access Control Installation and Administration Guide*.

