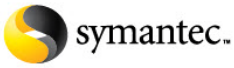

Symantec CCS Vulnerability Manager 10 Administrator's Guide

Revision history

Revision date	Description
June 15, 2010	Created document.
August 16, 2010	Added instructions for enabling FIPS mode, offline activations and updates.
September 13, 2010	Corrected a step in FIPS configuration instructions; added information about how to configure data warehousing.
September 22, 2010	Added instructions for verifying that FIPS mode is enabled; added section on managing updates
October 25, 2010	Updated instructions for activating, modifying, or renewing licenses.
December 13, 2010	Added instructions for SSH public key authentication.
December 20, 2010	Added instructions for using Asset Filter search and creating dynamic asset groups. Also added instructions for using new asset search features when creating static asset groups and reports.
March 16, 2011	Added instructions for migrating the database, enabling check correlation, including organization information in site configuration, managing assets according to host type, and performing new maintenance tasks.
March 31, 2011	Added a note to the database migration verification section.
April 18, 2011	Updated instructions for configuring Web spidering and migrating the database.
July 11, 2011	Added information about Scan Engine pooling, expanded permissions, and using the command console.
July 25, 2011	Corrected directory information for pairing the Security Console with Scan Engines.
September 19, 2011	Updated information about Dynamic Scan Pooling and FIPS mode configuration.
November 15, 2011	Added information about vAsset discovery, dynamic site management, new Real Risk and TemporalPlus risk strategies, and the Advanced Policy Engine.
December 5, 2011	Added note about how vAsset discovery currently finds assets in vSphere deployments only. Corrected some formatting issues.
January 23, 2012	Added information about the platform-independent backup option.
March 21, 2012	Added information about search filters for virtual assets, logging changes, and configuration options for Kerberos encryption.
June 6, 2012	Removed information about deprecated logging configuration page.
August 8, 2012	Added information about PostgreSQL database tuning; updated required JAR files for offline updates; added troubleshooting guidance for session timeout issues.

Contents

About this guide

This guide helps you to ensure that Symantec CCS Vulnerability Manager works effectively and consistently in support of your organization's security objectives. It provides instruction for doing key administrative tasks:

- configuring host systems for maximum performance
- planning a deployment, including determining how to distribute Scan Engines
- managing user accounts, roles, and permissions
- maintenance and troubleshooting

A note about documented features

The following features are not available in Symantec CCS Vulnerability Manager:

- Policy scanning
- Advanced Policy Engine
- FDCC
- USGCB
- Custom Policy scanning
- Policy Editor

Who should read this guide

You should read this guide if you fit one or more of the following descriptions:

- It is your responsibility to plan your organization's Symantec CCS Vulnerability Manager deployment.
- You have been assigned the Global Administrator role, which makes you responsible for maintenance, troubleshooting, and user management.

Other documents and Help

Click **Help** on any page of the Security Console Web interface to find information quickly. You will also find the following documents useful. You can download them from the *Support* page in *Help*.

User's guide

The user's guide helps you to gather and distribute information about your network assets and vulnerabilities using the application. It covers the following activities:

- logging onto the Security Console and familiarizing yourself with the Web interface
- managing vAsset discovery
- setting up sites and scans
- running scans manually
- viewing asset and vulnerability data
- creating remediation tickets

- using preset and custom report templates
- using report formats
- reading and interpreting report data
- configuring scan templates
- configuring other settings that affect scans and reports

API guide

The API guide helps you to automate some Symantec CCS Vulnerability Manager features and to integrate its functionality with your internal systems.

Document conventions

Words in bold are names of hypertext links and controls.

Words in italics are document titles, chapter titles, and names of Web interface pages.

1. Steps of procedures are indented and are numbered.

Items in `Courier font` are commands, command examples, and directory paths.

Items in **`bold Courier font`** are commands you enter.

Variables in command examples are enclosed in box brackets.
Example: `[installer_file_name]`

Options in commands are separated by pipes.
Example: `$ /etc/init.d/[daemon_name] start|stop|restart`

Keyboard commands are bold and are enclosed in arrow brackets.
Example: Press and hold **<Ctrl + Delete>**

NOTES, **TIPS**, and **WARNINGS** appear in the margin.

NOTES contain information that:

- enhances a description or a procedure.
- provides additional details that only apply in certain cases.

TIPS provide hints, best practices, or techniques for completing a task.

WARNINGS provide information about how to avoid potential loss of data or damage to data or a loss of system integrity.

Throughout this document, Symantec CCS Vulnerability Manager is referred to as *the application*.

For technical support

- Contact your Symantec account representative, or click the **Support** link on the Security Console Web interface.

Configuring maximum performance in an enterprise environment

This chapter provides system configuration tips and best practices to help ensure optimal performance of Symantec CCS Vulnerability Manager in an enterprise-scale deployment. The emphasis is on the system that hosts the Security Console. Some considerations are also included for Scan Engines.

Even if you are configuring the application for a smaller environment, you may still find some of this information helpful, particularly the sections maintaining and tuning the database, Scan Engine scaling, and disaster recovery considerations.

Configuring and tuning the Security Console host

The Security Console is the base of operations in a deployment. It manages Scan Engines and creates a repository of information about each scan, each discovered asset, and each discovered vulnerability in its database. With each ensuing scan, the Security Console updates the repository while maintaining all historical data about scans, assets, and vulnerabilities. The Security Console includes the server of the Web-based interface for configuring and operating the application, managing sites and scans, generating reports, and administering users.

The Security Console is designed to meet the scaling demands of an enterprise-level deployment. One Security Console can handle hundreds of Scan Engines, thousands of users, and any number of reports as long as it is running on sufficient hardware resources and is configured correctly.

Scan volume drives resource requirements

In an enterprise environment, the Security Console's most resource-intensive activities are processing, storing, and displaying scan data.

To determine resource sizing requirements, consider these important factors:

- The number of IP addresses that the application will scan: Every target generates a certain amount of data for the Security Console to store in its database. More targets mean more data.
- The frequency with which it will scan those assets: Scanning daily produces seven times more data than scanning weekly.
- The depth of scanning. A Web scan typically requires more time and resources than a network scan.
- The amount of detailed, historical scan data that it will retain over time: To the extent that scan data is retained in the database, this factor acts as a multiplier of the other two. Each retained set of scan data about a given target builds up storage overhead, especially with frequent scans.

Selecting a Security Console host for an enterprise deployment

The following hardware configuration is recommended to host the Security Console in an enterprise-level deployment. The definition of "enterprise-level" can vary.

Experience with past deployments indicates that 25,000 IP addresses or more, scanned with any reasonable frequency, warrants this recommended configuration:

- **vendor:** preferably IBM or Hewlett-Packard (These products are lab tested for performance)
- **processor:** 2x Intel quad-core Xeon 55xx "Nehalem" CPUs (2 sockets, 8 cores, and 16 threads total)
- **RAM:** 48-96 GB with error-correction code (ECC) memory; some 2-socket LGA1366 motherboards can support up to 144GB, with 8GB DDR3 modules
- **storage:** 8-12 x 7200RPM SATA/SAS hard drives, either 3.5" or 2.5" (if the chassis can only support that many drives in this form factor); total capacity should be 1+TB
- **network interface card (NIC):** 2 x 1GbE (one for scans, and one for redundancy or for a private-management subnet)

Examples of products that meet these specifications include the following:

- HP ProLiant DL380 G6
- IBM System x3650 M2Y

Your IT department or data center operations team may have preferred vendors. Or, your organization may build "white box" servers from commodity parts.

Linux expertise is essential

If your requirements dictate that you use a Linux-based host, consider the level of expertise in your organization for maintaining a Linux server.

Note that Red Hat Enterprise Linux 5.4 and 5.5 64-bit are the supported versions.

Configuring host memory for optimal performance

As indicated in the preceding section, 48 to 96 GB of RAM is recommended for enterprise deployments of the Security Console. Generally speaking, the Security Console should have 8 GB of RAM for performing its various functions. The remaining amount of RAM should be sufficient to cache the entire database in memory and so minimize disk access for read-only SQL queries.

To project how much the database will grow, note its current size based on current target count and scan frequency. Then, factor in the projected maximum numbers for these parameters.See *Scan volume drives resource requirements* on page 7. Depending on these variables, a production database can grow at a rate of 5 to 10 GB per week.

Storing the entire database in RAM is a best practice that optimizes Security Console performance. It minimizes the need to query data stored on disk, which is much slower than RAM. Another endorsement for Linux is that its multi-level buffer caching system efficiently manages data that is being cached in RAM, swapping in and out blocks of data to optimize how the operating system fulfills requests from the application.

More RAM means more speed. However, the cost of RAM starts to become prohibitive above 72 to 128 GB per node, depending on the type of memory. If the database seems on track to grow beyond that size, consider trimming old or unnecessary data.

The definition of "unnecessary" depends on your organization. It may be subject, for example, to compliance requirements that data be retained for a certain amount of time. Or data retention for a certain time period may be a matter of corporate policy. Otherwise, you can set retention criteria based simply on whether a given set of scan data is still useful for any reason.

Setting up an optimal RAID array

It should also be noted that the application cannot completely avoid querying data on disk. So, configuring a performance-friendly RAID array is important, especially given the fact that disk requirements can range up to 1TB.

Symantec recommends arranging multiple disks in a configuration of striped mirrors, also known as a RAID 1+0 or RAID 10 array, for better random disk I/O performance without sacrifice to redundancy. Symantec CCS Vulnerability Manager and PostgreSQL should be installed on this high-performing RAID 1+0 array. The PostgreSQL transaction log should be on independent disks, preferably a 2-drive mirror array (RAID 1). The operating system, which should generate very little disk I/O, may share this 2-drive mirror with the PostgreSQL transaction log.

A good purchasing approach will favor more desks over expensive disks. 8 to 12 disks are recommended. The application, the operating system, and PostgreSQL should each run on its own partition.

Maintaining the database

Given the amount of data that an enterprise deployment will generate, regularly scheduled backups are important. Nightly backups are recommended. During a database backup, Symantec CCS Vulnerability Manager goes into a maintenance mode and cannot run scans. Planning a deployment involves coordinating backup periods with scan windows. The time needed for backing up the database depends on the amount of data.

A backup saves the following items:

- the database
- configuration files (nsc.xml, nse.xml, userdb.xml, and consoles.xml)
- licenses
- keystores
- report images
- custom report templates
- custom scan templates
- generated reports
- scan logs

It is recommended that you perform the following database maintenance routines on a weekly basis:

- Clean up the database to remove leftover data that is associated with deleted objects, such as sites, assets, or users.
- Compress database tables to free up unused table space.
- Rebuild database indexes that may have become fragmented or corrupted over time.

Another maintenance task can be used to regenerate scan statistics so that the most recent statistics appear in the Security Console Web interface.

Additionally, a database optimization feature applies optional performance improvements, such as vulnerability data loading faster in the Security Console Web interface. It is recommended that you run this feature before running a backup.

For information on performing database backups and maintenance, see *Performing a backup and restore* on page *53*.

PostgreSQL also has an autovacuum feature that works in the background performing several necessary database maintenance chores. It is enabled by default and should remain so.

Tuning PostgreSQL 9

The application uses a PostgreSQL database server to store scan results and user data. The database server comes with a basic configuration to work with a wide range of installations. You can tune the configuration for maximum performance depending on your business needs and hardware specifications.

The tuning process involves modifying a PostgreSQL configuration file. The following section provides instructions for modifying the file (see *Modifying postgresql.conf* on page *11*) and recommendations for tuned settings (see *Tuned PostgreSQL settings* on page *12*). If you increase the shared_buffers setting, an additional step is required. See *Increasing the shmmax kernel parameter* on page *14*.

The tuning recommendations incorporate extensive performance testing on systems that represent typical midrange and enterprise configurations. The testing included common operations such as scanning, report generation, and administrative actions in typical load scenarios. The recommendations address both enterprise and midrange installations.

Which configuration should you choose?

The midrange configuration will yield optimal performance for deployments on systems with 8 GB of RAM or more. The enterprise configuration will yield optimal performance for deployments on systems with at least 64 GB of RAM. For systems with more than 8 GB, you can use the midrange settings and also set the effective_cache_size to one half of available RAM. See *Tuned PostgreSQL settings* on page *12*.

Applying the enterprise configuration to systems that don't meet the enterprise system requirements will likely degrade performance rather than improving it. Be sure to select the appropriate configuration for your needs.

Modifying postgresql.conf

Tuning PostgreSQL involves editing configuration settings in the postgresql.conf file.

1. Back up the postgresql.conf file before making any changes.
2. Locate the postgresql.conf file at [product_installation-directory]/nsc/nxpgsql/nxpdata/postgresql.conf.
3. Open the file in a text editing program.
4. Consult the tuning recommendations in the following table, and change desired parameters.
5. Save and close the file.
6. Restart the Security Console to allow the changes to take effect.

NOTES: Lines in the postgresql.conf file that begin with the pound sign (#) are comments and for informational purpose only. They do not affect configuration. For multiple listings of the same setting, the instance that appears last is the one that affects the configuration.

Tuned PostgreSQL settings

The following table lists PostgreSQL configuration parameters, their descriptions, default settings, and their recommended "tuned" settings. The table continues on the following page.

- The *Recommended midrange settings* are intended to work with 64-bit hardware running on 8GB of RAM.
- The *Recommended enterprise business settings* are intended to work in a higher-scan-capacity environment in which the application is installed on high-end hardware with 72 GB of RAM. See the *Selecting a Security Console host for an enterprise deployment* on page *8*.

Parameter	Description	Default value	Recommended midrange settings	Recommended enterprise settings
shared_ buffers	This is the amount of memory that is dedicated to PostgreSQL for caching data in RAM. PostgreSQL sets the default when initializing the database based on the hardware capacity available, which may not be optimal for the application. Enterprise configurations will benefit from a much larger setting for shared_buffers. Midrange configurations should retain the default that PostgreSQL allocates on first installation. **NOTE:** Increasing the default value may prevent the database from starting due to kernel limitations. To ensure that PostgreSQL starts, see *Increasing the shmmax kernel parameter* on page *14*	This value is set on PostgreSQL startup based on operating system settings.	24 MB	1950 MB
max_connections	This is the maximum number of concurrent connections to the database server. Increase this value if you anticipate a significant rise in the number of users and concurrent scans. Note that increasing this value requires approximately 400 bytes of shared memory per connection slot.	100	200	300
work_mem	This is the amount of memory that internal sort operations and hash tables use before switching to temporary disk files.	1 MB	32 MB	32 MB
checkpoint_segments	PostgreSQL writes new transactions to the database in files known as write ahead log (WAL) segments, which are 16 MB in size. These entries trigger checkpoints, or points in the transaction log sequence at which all data files have been updated to reflect the content of the log. The checkpoint_segments setting is the maximum distance between automatic checkpoints. At the default setting of 3, checkpoints can be can be resource intensive, producing 48 MB (16 MB multiplied by 3) and potentially causing performance bottlenecks. Increasing the setting value can mitigate this problem.	3	3	32
effective_cache_size	This setting reflects assumptions about the effective portion of disk cache that is available for a single query. It is factored into estimates of the cost of using an index. A higher value makes an index scan more likely. A lower value makes sequential scans more likely.	128 MB	4 GB (For configurations with more than 16 GB of RAM, use half of the available RAM as the setting.)	32 GB

Parameter	Description	Default value	Recommended midrange settings	Recommended enterprise settings
logging: log_min_error_statement	This setting controls whether or not the SQL statement that causes an error condition will be recorded in the server log. The current SQL statement is included in the log entry for any message of the specified severity or higher. Each value corresponds to one of the following severity levels in ascending order: DEBUG5, DEBUG4, DEBUG3, DEBUG2, DEBUG1, INFO, NOTICE, WARNING, ERROR, LOG, FATAL, and PANIC. The default value is ERROR, which means statements causing errors or more severe events will be logged. Increasing the log level can slow the performance of the application since it requires more data to be logged.	ERROR	ERROR	ERROR
logging: log_min_duration_statement	This setting causes the duration of each completed statement to be logged if the statement ran for at least the specified number of milliseconds. For example: A value of 5000 will cause all queries with an execution time longer than 5000 ms to be logged. The default value of -1 means logging is disabled. To enable logging, change the value to 0. This will increase page response time by approximately 5 percent, so it is recommended that you enable logging only if it is required. For example, if you find a particular page is taking a long time to load, you may need to investigate which queries may be taking a long time to complete.	-1	-1 (Set recommended value to 0 only if required for debugging)	-1 (Set recommended value to 0 only if required for debugging)
wal_buffers	This is the amount of memory used in shared memory for write ahead log (WAL) data. This setting does not affect select/update-only performance in any way. So, for an application in which the select/update ratio is very high, wal_buffers is almost an irrelevant optimization.	64 KB	64 KB	8 MB
maintenance_work_mem	This setting specifies the maximum amount of memory to be used by maintenance operations, such as VACUUM, CREATE INDEX, and ALTER TABLE ADD FOREIGN KEY.	16 MB	16 MB	512 MB

Increasing the shmmax kernel parameter

If you increase the shared_buffers setting as part of tuning PostgreSQL, check the shmmax kernel parameter to make sure that the existing setting for a shared memory segment is greater than the PostgreSQL setting. Increase the parameter if it is less than the PostgreSQL setting. This ensures that the database will start.

1. Determine the maximum size of a shared memory segment:

```
# cat /proc/sys/kernel/shmmax
```

2. Change the default shared memory limit in the proc file system.

```
# echo [new_kernel_size_in_bytes] > /proc/sys/kernel/shmmax
```

It is unnecessary to restart the system.

Alternatively, you can use sysctl(8) to configure the shmax parameters at run-time:

```
# sysctl -w kernel.shmmax=[new_kernel_size_in_bytes]
```

NOTE: If you do not make this change permanent, the setting will not persist after a system restart.

To make the change permanent, add a line to the /etc/sysctl.conf utilities file, which the host system uses during the startup process. Actual command settings may vary from the following example:

```
# echo "kernel.shmmax=[new_kernel_size_in_bytes]" >> /etc/sysctl.conf
```

Scan Engine scaling

To determine how many Scan Engines you need, consider the number of target IP addresses, the required scan frequency, and the number of available scan hours each day.

Empirical lab data indicates that one Scan Engine can completely scan 400 to 500 targets IP addresses in an hour. So, for example, if you have 30,000 live IP addresses and an 8-hour scan window, you need 8 Scan Engines. This calculation assumes no significant bandwidth limitations or scan window restrictions. Your results may vary based on your production environment.

Generally speaking, each added Scan Engine provides an incremental, linear boost to scan performance.

For suggestions on where to place Scan Engines, see *Distribute Scan Engines strategically* on page *24*.

For suggestions on optimizing scans, see the chapter on working with scan templates in the user's guide.

Disaster recovery considerations

As previously mentioned, one Security Console is sufficient for handling all activities at the enterprise level. However, an additional, standby Security Console may be warranted for your organization's disaster recovery plan for critical systems. If a disaster recovery plan goes into effect, this "cold standby" Security Console would require one database-restore routine in order to contain the most current data.

Disaster recovery may not warrant doubling the fleet of Scan Engines in the data center. Instead, a recovery plan could indicate having a number of spares on hand to perform a minimal requirement of scans—for example, on a weekly basis instead of daily—until production conditions return to normal. For example, if your organization has 10 Scan Engines in the data center, an additional 5 may suffice as temporary backup. Having a number of additional Scan Engines is also helpful for handling occasional scan spikes required by events such as monthly Microsoft patch verification.

Using anti-virus software on the server

Anti-virus programs may sometimes impact critical operations that are dependent on network communication, such as downloading updates and scanning. Blocking the latter may cause degraded scan accuracy.

If you are running anti-virus software on your intended host, configure the software to allow the application to receive the files and data that it needs for optimal performance in support your security goals:

- Add the application update server, updates.rapid7.com, to a whitelist, so that the application can receive updates.
- Add the application installation directory to a whitelist to prevent the anti-virus program from deleting vulnerability- and exploit-related files in this directory that it would otherwise regard as "malicious."

Consult your anti-virus vendor for more information on configuring the software to work with the application.

Planning a deployment

This chapter will help you deploy the application strategically to meet your organization's security goals. If you have not yet defined these goals, this guide will give you important questions to ask about your organization and network, so that you can determine what exactly you want to achieve.

The deployment and configuration options in the application address a wide variety of security issues, business models, and technical complexities. With a clearly defined deployment strategy, you can use the application in a focused way for maximum efficiency.

Understanding key concepts

Understanding the fundamentals of the application and how it works is key to determining how best to deploy it.

Understanding what the application does

Symantec CCS Vulnerability Manager is a unified vulnerability solution that scans networks to identify the devices running on them and to probe these devices for vulnerabilities. It analyzes the scan data and processes it for reports. You can use these reports to help you assess your network security at various levels of detail and remediate any vulnerabilities quickly.

The vulnerability checks identify security weaknesses in all layers of a network computing environment, including operating systems, databases, applications, and files. The application can detect malicious programs and worms, identify areas in your infrastructure that may be at risk for an attack, and verify patch updates and security compliance measures.

Understanding the components

The application consists of two main components:

- **Scan Engines** perform asset discovery and vulnerability detection operations. You can deploy Scan Engines outside your firewall, within your secure network perimeter, or inside your DMZ to scan any network *asset*.
- The **Security Console** communicates with Scan Engines to start scans and retrieve scan information. All exchanges between the Security Console and Scan Engines occur via encrypted SSL sessions over a dedicated TCP port that you can select. For better security and performance, Scan Engines do not communicate with each other; they only communicate with the Security Console.

  When the application scans an asset for the first time, the Security Console creates a repository of information about that asset in its database. With each ensuing scan that includes that asset, the Security Console updates the repository.

  The Security Console includes a Web-based interface for configuring and operating the application. An authorized user can log onto this interface securely, using HTTPS from any location, to perform any application-related task that his or her role permits. See *Understanding user roles and permissions* on page *18*. The authentication database is stored in an encrypted format on the Security Console server, and passwords are never stored or transmitted in plain text.

  Other Security Console functions include generating user-configured reports and regularly downloading patches and other critical updates from the Symantec central update system.

Symantec CCS Vulnerability Manager is "agentless"

The application performs all of its scanning operations over the network, using common Windows and UNIX protocols to gain access to target assets. This architecture makes it unnecessary for you to install and manage software agents on your target assets, which lowers the total cost of ownership (TCO) and eliminates security and stability issues associated with agents.

Understanding sites and asset groups

The Security Console interface enables you to plan scans effectively by organizing your network assets into *sites* and asset *groups*.

When you create a site, you identify the assets to be scanned, and then define scan parameters, such as scheduling and frequency. You also assign that site to a Scan Engine. You can only assign a given site to one Scan Engine. However, you can assign many sites to one Scan Engine.

You also define the type of scan you wish to run for that site. Each site is associated with a specific scan. The application supplies a variety of scan templates, which can expose different vulnerabilities at all network levels. Template examples include Penetration Test, Microsoft Hotfix, Denial of Service Test, and Full Audit. You also can create custom scan templates.

Another level of asset organization is an *asset group*. Like the site, this is a logical grouping of assets, but it is not defined for scanning. An asset group typically is assigned to a user who views scan reports about that group in order to perform any necessary remediation. An asset must be included within a site before you can add it to an asset group.

NOTE: If you are using RFC1918 addressing (192.168.x.x or 10.0.x.x addresses) different assets may have the same IP address. You can use site organization to enable separate Scan Engines located in different parts of the network to access assets with the same IP address.

Only designated global administrators are authorized to create sites and asset groups. For more details about access permissions, see *Understanding user roles and permissions* on page *18*.

Asset groups can include assets listed in multiple sites. They may include assets assigned to multiple Scan Engines, whereas sites can only include assets assigned to the same Scan Engine. Therefore, if you wish to generate reports about assets scanned with multiple Scan Engines, use the asset group arrangement. You also can configure reports for combination of sites, asset groups, and assets.

Understanding user roles and permissions

User access to Security Console functions is based on roles. You can assign default roles that include pre-defined sets of permissions, or you can create custom roles with permission sets that are more practical for your organization. See *Managing and creating user accounts* on page *39*. Once you give a role to a user, you restrict access in the Security Console to those functions that are necessary for the user to perform that role.

There are five default roles:

- *Global Administrator* on page *37*
- *Security Manager and Site Owner* on page *38*
- *Security Manager and Site Owner* on page *38*
- *Asset Owner* on page *38*
- *User* on page *39*

Define your goals

Knowing in advance what security-related goals you want to fulfill will help you design the most efficient and effective deployment for your organization.

Know your business case to know your goals

If you have not yet defined your goals for your deployment, or if you are having difficulty doing so, start by looking at your business model and your technical environment to identify your security needs.

Consider factors such as network topology, technical resources (hardware and bandwidth), human resources (security team members and other stake holders), time, and budget.

How big is your enterprise?

How many networks, subnetworks, and assets does your enterprise encompass?

The size of your enterprise is a major factor in determining how many Scan Engines you deploy.

What is the geography of your enterprise?

In how many physical locations is your network deployed? Where are these locations? Are they thousands or tens of thousands of miles away from each other, or across town from each other, or right next to each other? Where are firewalls and DMZs located?

These factors will affect how and where you deploy Scan Engines and how you configure your sites.

How is your network segmented?

What is the range of IP addresses and subnets within your enterprise?

Network segmentation is a factor in Scan Engine deployment and site planning.

What is your asset inventory?

What kinds of assets are you using? What are their functions? What operating systems, applications, and services are running on them? Which assets are physical hardware, and which are virtual? Where are these different assets located relative to firewalls and DMZs? What are your hidden network components that support other assets, such as VPN servers, LDAP servers, routers, switches, proxy servers, and firewalls? Does your asset inventory change infrequently? Or will today's spreadsheet listing all of your assets be out of date in a month?

Asset inventory influences site planning and scan template selection.

Does your asset inventory include laptops that employees take home? Laptops open up a whole new set of security issues that render firewalls useless. With laptops, your organization is essentially accepting external devices within your security perimeter. Network administrators sometimes unwittingly create back doors into the network by enabling users to connect laptops or home systems to a virtual private network (VPN).

Additionally, laptop users working remotely can innocently create vulnerabilities in many different ways, such as by surfing the Web without company-imposed controls or plugging in personal USB storage devices.

An asset inventory that includes laptops may require you to create a special site that you scan during business hours, when laptops are connected to your local network.

One possible environment: "Example, Inc."

As you answer the preceding questions, you may find it helpful to create a table. The following table lists network and asset information for a company called "Example, Inc."

Network segment	Address space	Number of assets	Location	Asset function
New York Sales	10.1.0.0/22	254	Building 1: Floors 1-3	Work stations
New York IT/Administration	10.1.10.0/23	50	Building 2: Floor 2	Work stations Servers
New York printers	10.1.20.0/24	56	Buildings 1 & 2	Printers
New York DMZ	172.16.0.0/22	30	Co-location facility	Web server Mail server
Madrid sales	10.2.0.0/22	65	Building 3: Floor 1	Work stations
Madrid development	10.2.10.0/23	130	Building 3: Floors 2 & 3	Work stations Servers
Madrid printers	10.2.20.0/24	35	Building 3: Floors 1-3	Printers
Madrid DMZ	172.16.10.0/24	15	Building 3: dark room	File server

What are the "hot spots" in your enterprise?

What assets contain sensitive data? What assets are on the perimeter of your network? Do you have Web, e-mail, or proxy servers running outside of firewalls?

Areas of specific concern may warrant Scan Engine placement. Also, you may use certain scan templates for certain types of high-risk assets. For example, a Web Audit scan template is most appropriate for Web servers.

What are your resources?

How much local-area network (LAN) and wide-area network (WAN) bandwidth do you have? What is your security budget? How much time do you have to run scans, and when can you run these scans without disrupting business activity?

These considerations will affect which scan templates you use, how you tune your scans, and when you schedule scans to run. See the chapter on setting up sites and scans in the user's guide.

What exactly are the security risks to your organization?

How easy is it for hackers to penetrate your network remotely? Are there multiple logon challenges in place to slow them down? How difficult is it for hackers to exploit vulnerabilities in your enterprise? What are the risks to data confidentiality? To data integrity? To data availability?

The triad of confidentiality, integrity, and availability (CIA) is a good metric by which to quantify and categorize risks in your organization.

Confidentiality is the prevention of data disclosure to unauthorized individuals or systems. What happens if an attacker steals customer credit card data? What if a trojan provides hacker access to your company's confidential product specifications, business plans, and other intellectual property?

Integrity is the assurance that data is authentic and complete. It is the prevention of unauthorized data modification. What happens when a virus wipes out records in your payroll database?

Availability refers to data or services being accessible when needed. How will a denial-of-service hack of your Web server affect your ability to market your products or services? What happens if a network attack takes down your phones? Will it cripple your sales team?

If your organization has not attempted to quantify or categorize risks, you can use reports to provide some guidelines. The algorithm that produces a risk score for each scanned asset calculates the score based on CIA factors.

Other risks have direct business or legal implications. What dangers does an attack pose to your organization's reputation? Will a breach drive away customers? Is there a possibility of getting sued or fined?

Knowing how your enterprise is at risk can help you set priorities for deploying Scan Engines, creating sites, and scheduling scans.

The application provides powerful tools for helping you to analyze and track risk so you prioritize remediation and monitor security trends in your environment over time. See the topics *Working with risk strategies to analyze threats* and *Working with risk trends in reports* in the user's guide.

What are your compliance requirements?

Many organizations have a specific reason for acquiring Symantec CCS Vulnerability Manager: they have to comply with a specific set of security requirements imposed by the government or by a private-sector entity that regulates their industry.

- Health care providers must protect the confidentiality of patient data as required by the Health Insurance Portability and Accountability Act (HIPAA).
- Many companies, especially those in the financial sector, are subject to security criteria specified in the Sarbanes-Oxley Act (SOX).
- U.S. government organizations and vendors who transact business with the government must comply with Federal Desktop Core Configuration (FDCC) policies for their Microsoft Windows systems.
- Merchants, who perform credit and debit card transactions, must ensure that their networks comply with Payment Card Industry (PCI) security standards.

The application provides a number of compliance tools, such as built-in scan templates that help you verify compliance with these standards. For a list of scan templates and their specifications, see *Where to find SCAP update information and OVAL files* on page *99*.

For official PCI scans the application provides additional tools, including PCI-sanctioned reports, Web interface features for PCI-specific site configuration and vulnerability exception management, and expanded application program interface (API) functionality for managing report distribution. For more information, see the *ASV Guide*, which you can request from Technical Support.

Verifying FDCC compliance

The application provides several tools to test and track FDCC policy compliance:

- a built-in FDCC scan template that includes Advanced Policy Engine checks for compliance with FDCC configuration policies (see the appendix on scan templates in the user's guide.)
- Web interface tools for tracking and overriding policy test results (see the chapter *Working with data from scans* in the *user's guide*.)
- XML and CSV reports for disseminating policy test result data (See the *user's guide*.)
- Web interface tools for viewing SCAP data and working with OVAL files (see *Where to find SCAP update information and OVAL files* on page *99*.)

These tools require a license that enables the Advanced Policy Engine and FDCC policy scanning.

What are your goals beyond compliance?

Compliance goals may help you to define your deployment strategy, but it's important to think beyond compliance alone to ensure security. For example, protecting a core set of network assets, such as credit card data servers in the case of PCI compliance, is important; but it may not be enough to keep your network secure—not even secure enough to pass a PCI audit.

Attackers will use any convenient point of entry to compromise networks. An attacker may exploit an Internet Explorer vulnerability that makes it possible to install a malicious program on an employee's computer when that employee browses the Web. The malware may be a remote execution program with which the hacker can access more sensitive network assets, including those defined as being critical for compliance.

Compliance, in and of itself, is not synonymous with security. On the other hand, a well implemented, comprehensive security plan will include among its benefits a greater likelihood of compliance.

Who is your security team?

Are you a one-person company or IT department? Are you the head of a team of 20 people, each with specific security-related tasks? Who in your organization needs to see asset/security data, and at what level of technical detail? Who's in charge of remediating vulnerabilities? What are the security considerations that affect who will see what information? For example, is it necessary to prevent a security analyst in your Chicago branch from seeing data that pertains to your Singapore branch?

These considerations will dictate how you set up asset groups, define roles and permissions, assign remediation tickets, and distribute reports. See *Managing users, roles, and permissions* on page *31*.

Ensuring complete coverage

The scope of your Symantec CCS Vulnerability Manager investment includes the type of license and the number of Scan Engines your purchase. Your license specifies a fixed, finite range of IP addresses. For example, you can purchase a license for 1,000 or 5,000 IP addresses.

Make sure your organization has a reliable, dynamic asset inventory system in place to ensure that your license provides adequate coverage. It may not be unusual for the total number of your organization's assets to fluctuate on a fairly regular basis. As staff numbers grow and recede, so does the number of workstations. Servers go on line and out of commission. Employees who are travelling or working from home plug into the network at various times using virtual private networks (VPNs).

This fluidity underscores the importance of having a dynamic asset inventory. Relying on a manually maintained spreadsheet is risky. There will always be assets on the network that are not on the list. And, if they're not on the list, they're not being managed. Result: added risk.

According to a paper by the technology research and advisory company, Gartner, Inc., an up-to-date asset inventory is as essential to vulnerability management as the scanning technology itself. In fact, the two must work in tandem:

"The network discovery process is continuous, while the vulnerability assessment scanning cycles through the environment during a period of weeks." (Source: *"A Vulnerability management Success Story" published by Gartner, Inc.*)

The paper further states that an asset database is a "foundation that enables other vulnerability technologies" and with which "remediation becomes a targeted exercise."

The best way to keep your asset database up to date is to perform discovery scans on a regular basis.

Planning your Scan Engine deployment

Your assessment of your security goals and your environment, including your asset inventory, will help you plan how and where to deploy Scan Engines. Keep in mind that if your asset inventory is subject to change on continual basis, you may need to modify your initial Scan Engine deployment over time.

Any deployment includes a Security Console and one or more Scan Engines to detect assets on your network, collect information about them, and test these assets for vulnerabilities. Scan Engines test vulnerabilities in several ways. One method is to check software version numbers, flagging out-of-date versions. Another method is a "safe exploit" by which target systems are probed for conditions that render them vulnerable to attack. The logic built into vulnerability tests mirrors the steps that sophisticated attackers would take in attempting to penetrate your network.

The application is designed to exploit vulnerabilities without causing service disruptions. It does not actually attack target systems.

One way to think of Scan Engines is that they provide strategic views of your network from a hacker's perspective. In deciding how and where to deploy Scan Engines, consider how you would like to "see" your network.

Distribute Scan Engines strategically

Distributed Scan Engines allow you to inspect your network from the inside. They are ideal for core servers and workstations. You can deploy distributed Scan Engines anywhere on your network to obtain multiple views. This flexibility is especially valuable when it comes to scanning a network with multiple subnetworks, firewalls, and other forms of segmentation.

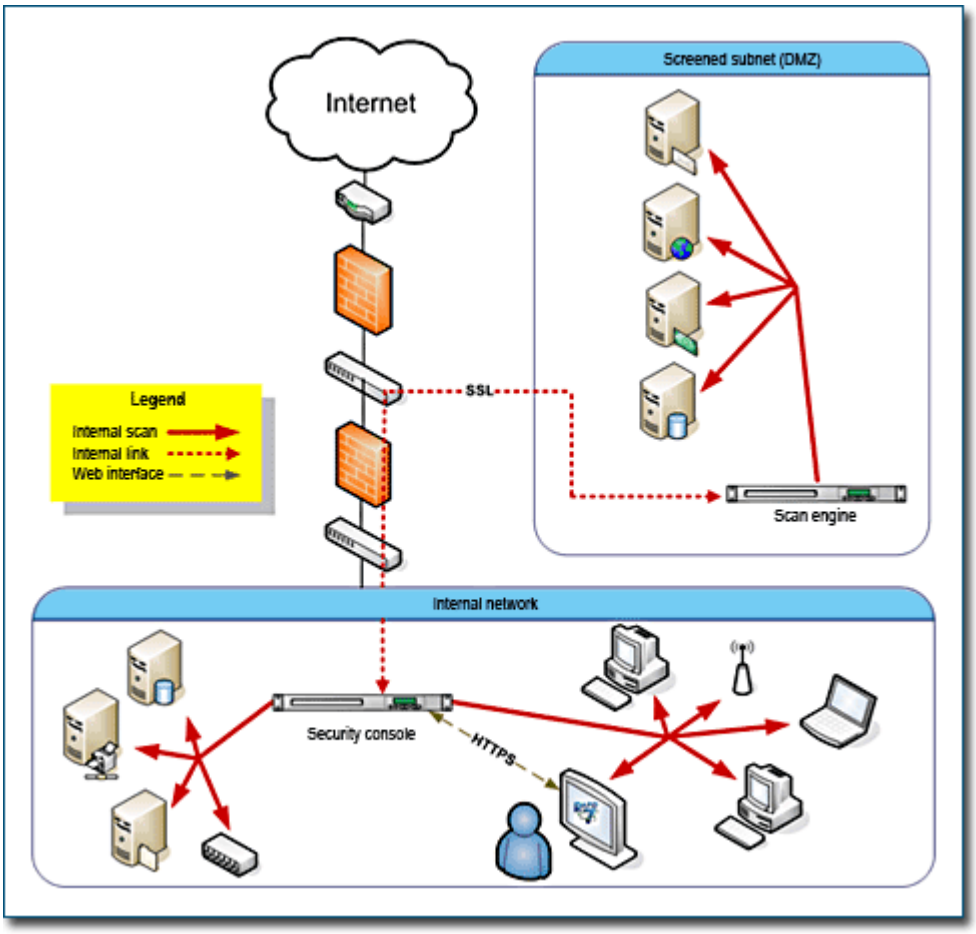

But, how many Scan Engines do you need? The question to ask first is, where you should you put them?

In determining where to put Scan Engines, it's helpful to look at your network topology. What are the areas of separation? And where are the connecting points? If you can answer these questions, you have a pretty good idea of where to put Scan Engines.

It is possible to operate a Scan Engine on the same host computer as the Security Console. While this configuration may be convenient for product evaluation or small-scale production scenarios, it is not appropriate for larger production environments, especially if the Scan Engine is scanning many assets. Scanning is a RAM-intensive process, which can drain resources away from the Security Console.

Following are examples of situations that could call for the placement of a Scan Engine.

Firewalls, IDS, IPS, and NAT devices

You may have a firewall separating two subnetworks. If you have a Scan Engine deployed on one side of this firewall, you will not be able to scan the other subnetwork without opening the firewall. Doing so may violate corporate security policies.

An application-layer firewall may have to inspect every packet before consenting to route it. The firewall has to track state entry for every connection. A typical scan can generate thousands of connection attempts in a short period, which can overload the firewalls state table or state tracking mechanism.

Scanning through an Intrusion Detection System (IDS) or Intrusion Prevention System (IPS) can overload the device or generate an excessive number of alerts. Making an IDS or IPS aware that Symantec CCS Vulnerability Manager is running a vulnerability scan defeats the purpose of the scan because it looks like an attack. Also, an IPS can compromise scan data quality by dropping packets, blocking ports by making them "appear" open, and performing other actions to protect assets. It may be desirable to disable an IDS or IPS for network traffic generated by Scan Engines.

Having a Scan Engine send packets through a network address transition (NAT) device may cause the scan to slow down, since the device may only be able to handle a limited number of packets per second.

In each of these cases, a viable solution would be to place a Scan Engine on either side of the intervening device to maximize bandwidth and minimize latency.

VPNs

Scanning across virtual private networks (VPNs) can also slow things down, regardless of bandwidth. The problem is the workload associated with connection attempts, which turns VPNs into bottlenecks. As a Scan Engine transmits packets within a local VPN endpoint, this VPN has to intercept and decrypt each packet. Then, the remote VPN endpoint has to decrypt each packet. Placing a Scan Engine on either side of the VPN tunnel eliminates these types of bottlenecks, especially for VPNs with many assets.

Subnetworks

The division of a network into subnetworks is often a matter of security. Communication between subnetworks may be severely restricted, resulting in slower scans. Scanning across subnetworks can be frustrating because they are often separated by firewalls or have access control lists (ACLs) that limit which entities can contact internal assets. For both security and performance reasons, assigning a Scan Engine to each subnetwork is a best practice

Perimeter networks (DMZs)

Perimeter networks, which typically include Web servers, e-mail servers, and proxy servers, are "out in the open," which makes them especially attractive to hackers. Because there are so many possible points of attack, it is a good idea to dedicate as many as three Scan Engines to a perimeter network. A local Scan Engine can scan vulnerabilities related to outbound data traffic, since hacked DMZ assets could transmit viruses across the Internet. Another local Scan Engine can provide an interior view of the DMZ.

ACLs

Access Control Lists (ACLs) can create divisions within a network by restricting the availability of certain network assets. Within a certain address space, such as 192.168.1.1/254, Symantec CCS Vulnerability Manager may only be able to communicate with 10 assets because the other assets are restricted ay an ACL. If modifying the ACL is not an option, it may be a good idea to assign a Scan Engine to ACL-protected assets.

WANs and remote asset locations

Sometimes an asset inventory is distributed over a few hundred or thousand miles. Attempting to scan geographically distant assets across a Wide Area Network (WAN) can tax limited bandwidth. A Scan Engine deployed near remote assets can more easily collect scan data and transfer that data to more centrally located database. It is less taxing on network resources to perform scans locally. Physical location can be a good principle for creating a site. See the topic *Configuring scan credentials* in the user's guide. This is relevant because each site is assigned to one Scan Engine.

Other factors that might warrant Scan Engine placement include routers, portals, third-party-hosted assets, outsourced e-mail, and virtual local-area networks.

Understanding Scan Engine-omics

Another way to project how many Scan Engines you need is by tallying the total number of assets that you intend to scan and how quickly you need to scan them. The following calculation will give you a general idea. Keep in mind that average numbers for simultaneous scans and total scan hours will vary depending on the types of assets you are scanning and the scan template you are using.

In this example, you may wish to scan a range of 300 IP addresses within 24 hours. A Scan Engine can run approximately three simultaneous scans. Each scan of a live asset can take up to five hours or more, depending on conditions mentioned in the preceding paragraph. If you multiply 300 assets by five hours of scanning for each asset, the result is 1500 total Scan Engine hours.

The application can scan up to three sites simultaneously. Most scan templates, by default, are configured for 10 scan threads, or 10 assets scanned simultaneously. This means that it can scan 30 assets simultaneously. If you divide 1500 total Scan Engine hours by 30, which is the number of supported simultaneous scans, the result is 50 Scan Engine hours. If you divide this number by 24, which is the number of hours in which you wish to complete all scanning, the result rounds up to three Scan Engines.

You may wish to add a few more Scan Engines for redundancy and load balancing. A comfortable number of Scan Engines for this situation would be approximately 10.

Having more Scan Engines potentially means faster scanning. However, certain constraints may affect how quickly you can expect to complete a scan job, even with a healthy fleet of Scan Engines. Bandwidth often is the biggest issue. Faster scanning requires more bandwidth. If you have limited bandwidth, then you will have to allow for wider scan windows.

Fault tolerance is an argument for deploying multiple Scan Engines to a location, especially when scan data for critical assets is time sensitive.

Working with Dynamic Scan Pooling

If your license enables Dynamic Scan Pooling, you can use pools to enhance the consistency of your scan coverage. A scan pool is a group of Scan Engines that can be bound to a site so that the Scan Engines are shared and the load is distributed evenly across the Scan Engines in the pool. Dynamic Scan Pooling provides two main benefits:

- Scan load balancing prevents overload of individual Scan Engines that can cause gaps in scan coverage. When a pool is bound to a site, scan jobs are distributed throughout the pool with a round-robin scheme, reducing the load on any single pooled Scan Engine.
- Fault tolerance prevents scans from failing due to operational problems with individual Scan Engines. If the Security Console contacts one pooled Scan Engine to start a scan, but the Scan Engine is offline, the Security Console simply contacts the next pooled Scan Engine to start the scan.

To view scan history for an existing site that has been assigned to a Dynamic Scan Pool you must temporarily reassign the site to the local Scan Engine. See *Understanding sites and asset groups* on page *18* and the chapter *Setting up sites and scans* in the user's guide for more information.

NOTE: Dynamic Scan Pooling is only available in the extended API v1.2.

You must pair Scan Engines with the Security Console before you can pool them. Also, when pooling Scan Engines, make sure that they are similarly configured and all located within the same network to prevent inconsistent scan results. For example, if one pooled Scan Engine is located in Network A and another is located in Network B, they will report different results when scanning an asset in Network A. The Scan Engine that is located in the same network can perform a deep, credentialed scan for more comprehensive results. The Scan Engine in Network B, on the other hand, can perform an external scan with more limited results.

You can deploy Dynamic Scan Pools using the Symantec CCS Vulnerability Manager extended API v1.2. For more information, see the *API Guide*, which you can download from the *Support* page in Help.

Setting up the application and getting started

Once you've mapped out your Scan Engine deployment, you're more than halfway to planning your installation. The next step is to decide how you want to install the main components—the Security Console and Scan Engines.

Understanding deployment options

When you install Symantec CCS Vulnerability Manager software on a given host, your options include running the application as a just a Scan Engine or as a Security Console and Scan Engine.

Installation scenarios—which one are you?

The different ways to install Symantec CCS Vulnerability Manager address different business scenarios and production environments. You may find one of these to be similar to yours.

Small business, internal network

The owner of a single, small retail store has a network of 50 or 60 work stations and needs to ensure that they are PCI compliant. The assets include registers, computers for performing merchandise look-ups, and file and data servers. They are all located in the same building. A software-only Security Console/Scan Engine on a single server is sufficient for this scenario.

Mid-size company with some remote locations

A company has a central office and two remote locations. The headquarters and one of the other locations have only a handful of assets between them. The other remote location has 300 assets. Network bandwidth is mediocre, but adequate. It definitely makes sense to dedicate a Scan Engine to the 300-asset location. The rest of the environment can be supported by a Security Console and Scan Engine on the same host. Due to bandwidth limitations, it is advisable to scan this network during off-hours.

Global enterprise with multiple, large remote locations

A company headquartered in the United States has locations all over the world. Each location has a large number of assets. Each remote location has one or more dedicated Scan Engines. One bank of Scan Engines at the U.S. office covers local scanning and provides emergency backup for the remote Scan Engines. In this situation, it is advisable not to use the Scan Engine that shares the host with the Security Console, since the Security Console has to manage numerous Scan Engines and a great deal of data.

Where to put the Security Console

Unlike Scan Engines, the Security Console is not restricted in its performance by its location on the network. Consoles initiate outbound connections with Scan Engines to initiate scans. When a Security Console sends packets through an opening in a firewall, the packets originate from "inside" the firewall and travel to Scan Engines "outside." You can install the Security Console wherever it is convenient for you.

One Security Console is typically sufficient to support an entire enterprise, assuming that the Security Console is not sharing host resources with a Scan Engine. If you notice that the Security Console's performance is slower than usual, and if this change coincides with a dramatic increase in scan volume, you may want to consider adding a second Security Console.

Configuring a the environment involves pairing each installed Scan Engine with a Security Console. To say that multiple Security Consoles can share a Scan Engine means that a Scan Engine can be paired with multiple Security Consoles. For information on pairing Security Consoles and Scan Engines, see the topic *Specifying general static site information* in the user's guide.

A deployment plan for Example, Inc.

Let's return to the environment table for Example, Inc.

Network segment	Address space	Number of assets	Location	Asset function
New York Sales	10.1.0.0/22	254	Building 1: Floors 1-3	Work stations
New York IT/Administration	10.1.10.0/23	50	Building 2: Floor 2	Work stations Servers
New York printers	10.1.20.0/24	56	Buildings 1 & 2	Printers
New York DMZ	172.16.0.0/22	30	Co-location facility	Web server Mail server
Madrid sales	10.2.0.0/22	65	Building 3: Floor 1	Work stations
Madrid development	10.2.10.0/23	130	Building 3: Floors 2 & 3	Work stations Servers
Madrid printers	10.2.20.0/24	35	Building 3: Floors 1-3	Printers
Madrid DMZ	172.16.10.0/24	15	Building 3: dark room	File server

A best-practices deployment plan might look like this:

The eight groups collectively contain a total of 635 assets. Example, Inc., could purchase a fixed-number license for 635 licenses, but it would be wiser to purchase a discovery for the total address space. It is always a best practice to scan all assets in an environment according to standards such as PCI, ISO 27002, or ISO 27001. This practice reflects the hacker approach of viewing any asset as a possible attack point.

Example, Inc., should distribute Symantec CCS Vulnerability Manager components throughout its four physical locations:

- Building 1
- Building 2
- Building 3
- Co-Location facility

The IT or security team should evaluate each of the LAN/WAN connections between these locations for quality and bandwidth availability. The team also should audit these pipes for devices that may prevent successful scanning, such as firewalls, ACLs, IPS, or IDS.

Finally the team must address any logical separations, like firewalls and ACLs, which may prevent access.

The best place for the Security Console is in New York because the bulk of the assets are there, not to mention IT and administration groups.

Assuming acceptable service quality between the New York buildings, the only additional infrastructure would be a Scan Engine inside the Co-Location facility.

Example, Inc., should install at least one Scan Engine in the Madrid location, since latency and bandwidth utilization are concerns over a WAN link.

Finally, it's not a bad idea to add one more Scan Engine for the Madrid DMZ to bypass any firewall issues.

The following table reflects this plan.

Asset	Location
Security Console	New York: Building 2
Scan Engine #1	New York: Co-Location Facility
Scan Engine #2	Madrid: Building 3
Scan Engine #3	Madrid: dark room

Your deployment checklist

When you are ready to install, configure, and run Symantec CCS Vulnerability Manager, it's a good idea follow a general sequence. Certain tasks are dependent on others being completed.

You will find yourself repeating some of these steps:

- install components
- log onto the Security Console Web interface
- configure Scan Engines, and pair them with the Security Console
- perform vAsset discovery, if your license enables it
- create one or more sites
- assign each site to a Scan Engine
- select a scan template for each site
- schedule scans
- create user accounts, and assign site-related roles and permissions to these accounts
- run scans
- configure and run reports
- create asset groups to view reports and asset data
- create user accounts, and assign asset-group-related roles and permissions to these accounts
- assign remediation tickets to users
- re-run scans to verify remediation

Managing users, roles, and permissions

Effective use of scan information depends on how your organization analyzes and distributes it, who gets to see it, and for what reason. Managing access to information in the application involves creating asset groups and assigning roles and permissions to users. This chapter provides best practices and instructions for managing users, roles, and permissions.

Map roles to your organization

It is helpful to study how roles and permissions map to your organizational structure.

NOTE: A user authentication system is included. However, if your organization already uses an authentication service that incorporates Microsoft Active Directory or Kerberos, it is a best practice to integrate the application with this service. Using one service prevents having to manage two sets of user information.

In a smaller company, one person may handle all security tasks. He or she will be a Symantec CCS Vulnerability Manager Global Administrator, initiating scans, reviewing reports, and performing remediation. Or there may be a small team of people sharing access privileges for the entire system. In either of these cases, it is unnecessary to create multiple roles, because all network assets can be included in one site, requiring a single Scan Engine.

Example, Inc. is a larger company. It has a wider, more complex network, spanning multiple physical locations and IP address segments. Each segment has its own dedicated support team managing security for that segment alone.

One or two global administrators are in charge of creating user accounts, maintaining the system, and generating high-level, executive reports on all company assets. They create sites for different segments of the network. They assign security managers, site administrators, and system administrators to run scans and distribute reports for these sites.

The global administrators also create various asset groups. Some will be focused on small subsets of assets. Non-administrative users in these groups will be in charge of remediating vulnerabilities and then generating reports after follow-up scans are run to verify that remediation was successful. Other asset groups will be more global, but less granular, in scope. The non-administrative users in these groups will be senior managers who view the executive reports to track progress in the company's vulnerability management program.

Create custom roles and using preset roles

Whether you create a custom role or assign a preset role for an account depends on several questions: What tasks do you want that account holder to perform? What data should be visible to the user? What data should not be visible to the user.

For example, a manager of a security team that supports workstations may need to run scans on occasion and then distribute reports to team members to track critical vulnerabilities and prioritizing remediation tasks. This account may be a good candidate for an Asset Owner role with access to a site that only includes workstations and not other assets, such as database servers.

NOTE: Keep in mind that, except for the Global Administrator role, the assigning of a custom or preset role is interdependent with access to site and asset groups.

If you want to assign roles with very specific sets of permissions you can create custom roles. The following tables list and describe all permissions that are available. Some permissions require other permissions to be granted in order to be useful. For example, in order to be able to create reports, a user must also be able to view asset data in the reported-on site or asset group, to which the user must also be granted access.

The tables also indicate which roles include each permission. You may find that certain roles are granular or inclusive enough for a given account. A list of preset roles and the permissions they include follows the permissions tables. See *Give a user access to asset groups* on page *41*.

Permissions tables

TIP: Click any vertically aligned role in the permissions tables, such as Site Owner, to read information about it.

Global permissions

These permissions automatically apply to all sites and asset groups and do not require additional, specified access.

Permission	Description	Global Administrator	Security Manager and Site Owner	Security Manager and Site Owner	Asset Owner	User
Manage Sites	Create, delete, and configure all attributes of sites, except for user access. Implicitly have access to all sites. Manage shared scan credentials. **Other affected permissions:** When you select this permission, all site permissions automatically become selected. See *Site permissions* on page *33*.	x				
Manage Scan Templates	Create, delete, and configure all attributes of scan templates.	x				
Manage Report Templates	Create, delete, and configure all attributes of report templates.	x	x	x	x	x
Manage Scan Engines	Create, delete, and configure all attributes of Scan Engines; pair Scan Engines with the Security Console.	x				
Manage Policies	Copy existing policies; edit and delete custom policies.	x				
Appear on Ticket and Report Lists	Appear on user lists in order to be assigned remediation tickets and view reports. **Prerequisite:** A user with this permission must also have asset viewing permission in any relevant site or asset group: • *View Site Asset Data* on page *33* • *View Group Asset Data* on page *34*	x	x	x	x	x
Configure Global Settings	Configure settings that are applied throughout the entire environment, such as risk scoring and exclusion of assets from all scans.	x				

Site permissions

These permissions only apply to sites to which a user has been granted access.

Permission	Description	Global Administrator	Security Manager and Site Owner	Security Manager and Site Owner	Asset Owner	User
View Site Asset Data	View discovered information about all assets in accessible sites, including IP addresses, installed software, and vulnerabilities.	x	x	x	x	x
Specify Site Metadata	Enter site descriptions, importance ratings, and organization data.	x	x	x		
Specify Scan Targets	Add or remove IP addresses, address ranges, and host names for site scans.	x				
Assign Scan Engine	Assign a Scan Engine to sites.	x				
Assign Scan Template	Assign a scan template to sites.	x	x	x		
Manage Scan Alerts	Create, delete, and configure all attributes of alerts to notify users about scan-related events.	x	x	x		
Manage Site Credentials	Provide logon credentials for deeper scanning capability on password-protected assets.	x	x	x		
Schedule Automatic Scans	Create and edit site scan schedules.	x	x	x		
Start Unscheduled Scans	Manually start one-off scans of accessible sites (does not include ability to configure scan settings).	x	x	x	x	
Purge Site Asset Data	Manually remove asset data from accessible sites. **Prerequisites:** A user with this permission must also have one of the following permissions: • *View Site Asset Data* on page *33* • *View Group Asset Data* on page *34*	x				
Manage Site Access	Grant and remove user access to sites.	x				

Asset Group permissions

These permissions only apply to asset groups to which a user has been granted access.

Permission	Description	Global Administrator	Security Manager and Site Owner	Security Manager and Site Owner	Asset Owner	User
Manage Dynamic Asset Groups	Create dynamic asset groups. Delete and configure all attributes of accessible dynamic asset groups except for user access. **Implicitly have access to all sites.** **Note:** A user with this permission has the ability to view all asset data in your organization.	x				
Manage Static Asset Groups	Create static asset groups. Delete and configure all attributes of accessible static asset groups except for user access. **Prerequisite:** A user with this permission must also have the following permissions and access to at least one site to effectively manage static asset groups: • *Manage Group Assets* on page *34* • *View Group Asset Data* on page *34*	x				
View Group Asset Data	View discovered information about all assets in accessible asset groups, including IP addresses, installed software, and vulnerabilities.	x	x		x	x
Manage Group Assets	Add and remove assets in static asset groups. **Note:** This permission does not include ability to delete underlying asset definitions or discovered asset data. **Prerequisite**: A user with this permission must also have of the following permission: • *View Group Asset Data* on page *34*	x				
Manage Asset Group Access	Grant and remove user access to asset groups.	x				

Report permissions

The Create Reports permission only applies to assets to which a user has been granted access. Other report permissions are not subject to any kind of access.

Permission	Description	Global Administrator	Security Manager and Site Owner	Security Manager and Site Owner	Asset Owner	User
Create Reports	Create and own reports for accessible assets; configure all attributes of owned reports, except for user access. **Prerequisites:** A user with this permission must also have one of the following permissions: • *View Site Asset Data* on page *33* • *View Group Asset Data* on page *34*	x	x	x	x	x
Use Restricted Report Sections	Create report templates with restricted sections; configure reports to use templates with restricted sections. **Prerequisites:** A user with this permission must also have one of the following permissions: • *Manage Report Templates* on page *32*	x				
Manage Report Access	Grant and remove user access to owned reports.	x				

Ticket permissions

These permissions only apply to assets to which a user has been granted access.

Permission	Description	Global Administrator	Security Manager and Site Owner	Security Manager and Site Owner	Asset Owner	User
Create Tickets	Create tickets for vulnerability remediation tasks. **Prerequisites:** A user with this permission must also have one of the following permissions: • *View Site Asset Data* on page *33* • *View Group Asset Data* on page *34*	x	x	x	x	x
Close Tickets	Close or delete tickets for vulnerability remediation tasks. **Prerequisites:** A user with this permission must also have one of the following permissions: • *View Site Asset Data* on page *33* • *View Group Asset Data* on page *34*	x	x	x	x	x

Vulnerability exception permissions

These permissions only apply to sites or asset groups to which a user has been granted access.

Permission	Description	Global Administrator	Security Manager and Site Owner	Security Manager and Site Owner	Asset Owner	User
Submit Vulnerability Exceptions	Submit requests to exclude vulnerabilities from reports. **Prerequisites:** A user with this permission must also have one of the following permissions: • *View Site Asset Data* on page *33* • *View Group Asset Data* on page *34*	x	x	x	x	x
Review Vulnerability Exceptions	Approve or reject requests to exclude vulnerabilities from reports. **Prerequisites:** A user with this permission must also have one of the following permissions: • *View Site Asset Data* on page *33* • *View Group Asset Data* on page *34*	x				
Delete Vulnerability Exceptions	Delete vulnerability exceptions and exception requests. **Prerequisites:** A user with this permission must also have one of the following permissions: • *View Site Asset Data* on page *33* • *View Group Asset Data* on page *34*	x				

List of roles

Global Administrator

The Global Administrator role differs from all other preset roles in several ways. It is not subject to site or asset group access. It includes all permissions available to any other preset or custom role. It also includes permissions that are not available to custom roles:

- Manage all functions related to user accounts, roles, and permissions.
- Manage vConnections and vAsset discovery.
- Manage configuration, maintenance, and diagnostic routines for the Security Console.
- Manage shared scan credentials.

Security Manager and Site Owner

The Security Manager and Site Owner roles include the following permissions:

- *Manage Report Templates* page *32*
- *Appear on Ticket and Report Lists* page *32*
- *View Site Asset Data* page *33*
- *Specify Site Metadata* page *33*
- *Assign Scan Template* page *33*
- *Manage Scan Alerts* page *33*
- *Manage Site Credentials* page *33*
- *Schedule Automatic Scans* page *33*
- *Start Unscheduled Scans* page *33*
- *View Group Asset Data* page *34* (Security Manager only)
- *Create Reports* page *35*
- *Create Tickets* page *36*

The only distinction between these two roles is the Security Manager's ability to work in accessible sites *and* assets groups. The Site Owner role, on the other hand, is confined to sites.

Asset Owner

The Asset Owner role includes the following permissions in accessible sites and asset groups:

- *Manage Report Templates* page *32*
- *Appear on Ticket and Report Lists* page *32*
- *View Site Asset Data* page *33*
- *Start Unscheduled Scans* page *33*
- *View Group Asset Data* page *34*
- *Create Reports* page *35*

User

Although "user" can refer generically to any owner of a Symantec CCS Vulnerability Manager account, the name *User*, with an upper-case *U,* refers to one of the preset roles. It is the only role that does not include scanning permissions. It includes the following permissions in accessible sites and asset groups:

- *Manage Report Templates* page *32*
- *Appear on Ticket and Report Lists* page *32*
- *View Site Asset Data* page *33*
- *View Group Asset Data* page *34* (Security Manager only)
- *Create Reports* page *35*
- *Create Tickets* page *36*

Managing and creating user accounts

The **Users** links on the *Administration* page provide access to pages for creating and managing user accounts. Click **manage** next to *Users* to view the *Users* page. On this page, you can view a list of all accounts within your organization.

To edit a user account:

1. Click **Edit** for any listed account, and change its attributes.

   The application displays the *User Configuration* panel. The process for editing an account is the same as the process for creating a new user account. See *Configure general user account attributes* on page *40*.

To delete an account:

1. Click **Delete** for that account.

   If that account has been used to create a report, or if that account has been assigned a ticket, the application displays a dialogue box prompting you to reassign or delete the report or ticket in question. Doing the latter might make sense for a ticket that concerns a closed issue or an old report that contains out-of-date information.

To assign listed items to an alternate account:

1. Select an account from the drop-down list, and click **Reassign items**.

   OR

2. Click **Delete items** to remove these items from the database.

Configure general user account attributes

You can specify attributes for general user accounts on the *User Configuration* panel.

To configure user account attributes:

1. Click **New User** on the *Users* page.
2. (Optional) Click **Create** next to *Users* on the *Administration* page. The Security Console displays the *General* page of the *User Configuration* panel.
3. Enter all requested user information in the text fields.
4. (Optional) Select the appropriate source from the drop-down list to authenticate the user with external sources.

   Before you can create externally authenticated user accounts you must define external authentication sources. See *Using external sources for user authentication* on page *46*.
5. Check the **Account enabled** check box.

   You can later disable the account without deleting it by clicking the check box again to remove the checkmark.
6. Click **Save** to save the new user information.

Assign a role and permissions to a user

Assigning a role and permissions to a new user allows you to control that user's access to Security Console functions.

To assign a role and permissions to a new user:

1. Go to the *Roles* page.
2. Choose a role from the drop-down list.

   When you select a role, the Security Console displays a brief description of that role.

   If you choose one of the five default roles, the Security Console automatically selects the appropriate check boxes for that role.

   If you choose **Custom Role**, select the check box for each permission that you wish to grant the user.
3. Click **Save** to save the new user information.

Give a user access to specific sites

A global administrator automatically has access to all sites. A security manager, site administrator, system administrator, or nonadministrative user has access only to those sites granted by a global administrator.

To grant a user access to specific sites:

1. Go to the *Site Access* page.
2. (Optional) Click the appropriate radio button to give the user access to all sites.
3. (Optional) Click the radio button for creating a custom list of accessible sites to give the user access to specific sites,
4. Click **Add Sites**.

   The Security Console displays a box listing all sites within your organization.
5. Click the check box for each site that you want the user to access.
6. Click **Save**.

   The new site appears on the *Site Access* page.
7. Click **Save** o save the new user information.

Give a user access to asset groups

A global administrator automatically has access to all asset groups. A site administrator user has no access to asset groups. A security manager, system administrator, or nonadministrative user has access only to those access groups granted by a global administrator.

To grant a user access to asset group:

1. Go to the *Asset Group Access* page.
2. (Optional) Click the appropriate radio button to give the user access to all asset groups.
3. (Optional) Click the radio button for creating a custom list of accessible asset groups to give the user access to specific asset groups.
4. Click **Add Groups**.

   The Security Console displays a box listing all asset groups within your organization.
5. Click the check box for each asset group that you want this user to access.
6. Click **Save**.

   The new asset group appears on the *Asset Group Access* page.
7. Click **Save** to save the new user information.

Managing and maintaining the application

This chapter provides instructions for managing and maintaining Symantec CCS Vulnerability Manager through controls in the Web interface and a set of command prompts.

Configure data warehousing settings

NOTE: Currently, this warehousing feature only supports PostgreSQL databases.

You can configure warehousing settings to store scan data or to export it to a PostgreSQL database. You can use this feature to obtain a richer set of scan data for integration with your own internal reporting systems.

NOTE: Due to the amount of data that can be exported, the warehousing process may take a long time to complete.

This is a technology preview of a feature that is undergoing expansion.

To configure data warehouse settings:

1. Click **manage** next to *Data Warehousing* on the *Administration* page.
2. Enter database server settings on the **Database** page.
3. Go to the *Schedule* page, and select the check box to enable data export.
   You can also disable this feature at any time.
4. Select a date and time to start automatic exports.
5. Select an interval to repeat exports.
6. Click **Save.**

Manage Security Console settings

TIP: Click **Manage** next to Security Console on the *Administration* page to launch the Security Console Configuration panel.

Although the default Security Console settings should work for a broad range of network environments, you can change settings to meet specific scanning requirements.

Viewing general configuration settings

On the *General* page, you can view the version and serial numbers for the instance of the Security Console that you are using.

You also can enable the process auto-stop feature, which pauses scans and stops report generation when the memory on the host server is dangerously low. The benefit of this feature is that it reduces the possibility of the server failing. However, it may cause the application to stop scans or reporting activities before they are complete.

Changing the Security Console Web server default settings

The Security Console runs its own Web server, which delivers the user interface.

To change the Security Console web server default settings:

1. Go to the *Web Server* page.
2. Type a new number for the access port if desired.
3. Type a new session time-out if desired.

   This value is the allowed number of seconds of user inactivity after which the Security Console times out, requiring a new logon.
4. Type new numbers for initial request and maximum request handler threads, if necessary.

   It is recommended that you consult Technical Support first. In this context, threads refer to the number of simultaneous connections that the Security Console will allow. Typically a single browser session accounts for one thread. If simultaneous user demand is high, you can raise the thread counts to improve performance. The Security Console will increase the thread count dynamically if required, so manual increases may be unnecessary.
5. Type a new number for failed logon threshold if desired. This is the number of failed logon attempts that the Security Console permits before locking out the would-be user.
6. Click **Save**.

Managing the HTTPS certificate

The application, by default, uses a self-signed X.509 certificate which is created during installation. You can replace this certificate with a custom, self-signed certificate or a certificate signed by a trusted.

NOTE: The signed certificate must be based on an application -generated CSR. The application does not allow you to import an arbitrary key pair/certificate that you generated

You can view the current HTTPs certificate on the Web server page. You also can change the certificate in three different ways.

You can create a new, self-signed SSL certificate to overwrite the current one for the application Web server. You can use the new certificate 'as-is' or you can have it can be signed by a CA by generating a certificate signing request (CSR).

To manage certificates:

1. Click **Manage HTTPS Certificate** on the Web Server page.
   The Security Console displays a box titled *Manage HTTPs Certificate*.
2. Click **Create New Certificate**.
   The Security Console displays a box for new certificate information.
3. Type the information, and click **Create**. Then, click **Done**.
   The new certificate name appears on the *Web Server* page.

To generate a CSR once you have created a new certificate.

1. Click **Manage HTTPS Certificate** on the Web Server page.
2. Click **Generate CSR** in the *Manage HTTPs Certificate* box.
3. Copy the generated CSR and send to your CA.
4. Import the certificate after it is signed by your CA.
5. Copy the certificate.
6. Click **Manage HTTPS Certificate** on the Web Server page.
7. Paste it in the text box and click **Import**.
8. Click **Done**.
9. Click **Save** to save the new Security Console information.

NOTE: The certificate can contain one or more subject-alternative X.509 name extensions.

Manage updates

By default, Symantec CCS Vulnerability Manager automatically downloads and applies two types of updates.

Content updates

Content updates include new checks for vulnerabilities, patch verification, and security policy compliance. Content updates always occur automatically when they are available.

Product updates

Product updates include performance improvements, bug fixes, and new product features. Unlike content updates, it is possible to disable automatic product updates and update the product manually.

Disabling automatic product updates

You can disable automatic product updates and initiate one-time product updates on an as-needed basis. This gives your organization the time and flexibility to train staff or otherwise prepare for updates that might cause changes in workflow. For example, a new feature may streamline a particular workflow by eliminating certain steps.

To disable automatic product updates:

1. Go to the **Administration** tab.
2. Click **manage** next to *Security Console*.

   The *Security Console Configuration* panel appears.
3. Select **Updates** from the menu on the left-hand side.
4. Clear the checkbox labeled **Enable automatic product updates**.

   A warning dialog box appears about the risks of disabling automatic product updates.
5. Click **Disable automatic product updates** to confirm that you want to turn off this feature.
6. (Optional) Click **Cancel** to leave automatic product updates enabled.
7. Click **Save**.

   Whenever you change this setting and click **Save**, the application downloads any available product updates. If you have disabled the setting, it does not apply any downloaded product updates.

Enabling automatic product updates

NOTE: Your PostgreSQL database must be version 9. Otherwise, the application will not apply product updates. If you are using an earlier version of PostgreSQL, *see Migrating the database* on page *57*.

Enabling automatic product updates ensures that you are always running the most current version of the application.

To enable automatic product updates after they have been previously disabled:

1. Go to the **Administration** tab.
2. Click **manage** next to *Security Console*.

   The *Security Console Configuration* panel appears.
3. Select **Updates** from the menu on the left-hand side.
4. Select the **Enable automatic product updates** checkbox.
5. Click **Save**.

   Whenever you change this setting and click **Save**, the application downloads any available product updates. If you have enabled the setting, it also applies any downloaded product updates and restarts.

Manual product updates

When automatic product updates have been disabled, you can manually download product updates.

NOTE: By using this one-time update feature, you are not enabling future automatic product updates if they are not currently enabled.

To manually download a new product update:

1. Go to the **Administration** page.
2. Click **manage** next to *Security Console*.

   The *Security Console Configuration* screen appears.
3. Select **Updates** from the menu on the left-hand side.

   Current available updates appear on the *Updates* page.
4. Click **Manual Update** to install them.

   A warning dialog box appears, indicating that the time to update will vary depending on the number and complexity of updates, and that future automatic product updates will remain disabled.
5. Click **Complete this one-time update** to perform the update.
6. (Optional) Click **Cancel** if you do not want to perform the update.

Configuring proxy settings

The proxy settings function allows the Security Console to redirect its update requests to an alternative server. Technical Support will advise if you need to change this setting.

If the Security Console has no direct connection to the Internet, a pre-configured proxy server can allow indirect access.

NOTE: For information on configuring updates for an Appliance, see the *Appliance Guide* which you can download from the *Support* page of *Help*.

To configure proxy settings:

1. Go to the *Proxy Settings* page.
2. Type the information for the proxy server in the appropriate fields.
3. Initiate an auto-update using the **`update now`** command.
   See *Using the command console* on page *72*.
4. Click **Save** to save the new Security Console information.

Using external sources for user authentication

You can integrate Symantec CCS Vulnerability Manager with external authentication sources. If you use one of these sources, leveraging your existing infrastructure will make it easier for you to manage user accounts.

The application provides single-sign-on external authentication with two sources:

- **LDAP (including Microsoft Active Directory):** Active Directory (AD) is an LDAP-supportive Microsoft technology that automates centralized, secure management of an entire network's users, services, and resources.
- **Kerberos:** Kerberos is a secure authentication method that validates user credentials with encrypted keys and provides access to network services through a "ticket" system.

The application also continues to support its two internal user account stores:

- XML file lists default "built-in" accounts. A Global Administrator can use a built-in account to log on to the application in maintenance mode to troubleshoot and restart the system when database failure or other issues prevent access for other users.
- Datastore lists standard user accounts, which are created by a global administrator.

Before you can create externally authenticated user accounts you must define external authentication sources.

To define external authentication sources:

1. Go to the *Authentication* page in the *Security Console Configuration* panel.
2. Click **Add...** in the area labelled *LDAP/AD authentication sources* to add an LDAP/Active Directory authentication source

   The Security Console displays a box labeled *LDAP/AD Configuration*.
3. Click the checkbox labeled **Enable authentication source**.
4. Type the name, address or fully qualified domain name, and port of the LDAP server that you wish to use for authentication.

   Default LDAP port numbers are 389 or 636, the latter being for SSL. Default port numbers for Microsoft AD with Global Catalog are 3268 or 3269, the latter being for SSL.

NOTE: It is recommended that you enter a fully qualified domain name in all capital letters for the LDAP server configuration. Example: SERVER.DOMAIN.EXAMPLE.COM

5. (Optional) Select the appropriate check box to require secure connections over SSL.
6. (Optional) Specify permitted authentication methods, type them in the appropriate text field. Separate multiple methods with commas (,), semicolons (;), or spaces.

   Simple Authentication and Security Layer (SASL) authentication methods for permitting LDAP user authentication are defined by the Internet Engineering Task Force in document RFC 2222 *(http://www.ietf.org/rfc/rfc2222.txt* (http://www.ietf.org/rfc/rfc2222.txt)). The application supports the use of GSSAPI, CRAM-MD5, DIGEST-MD5, SIMPLE, and PLAIN methods.

NOTE: It is not recommended that you use PLAIN for non-SSL LDAP connections.

7. Click the checkbox labeled **Follow LDAP referrals** if desired.

   As the application attempts to authenticate a user, it queries the target LDAP server. The LDAP and AD directories on this server may contain information about other directory servers capable of handling requests for contexts that are not defined in the target directory. If so, the target server will return a referral message to the application, which can then contact these additional LDAP servers. For information on LDAP referrals, see the document *LDAPv3 RFC 2251* (http://www.ietf.org/rfc/rfc2251.txt).
8. Type the base context for performing an LDAP search if desired. You can initiate LDAP searches at many different levels within the directory.

   To force the application to search within a specific part of the tree, specify a search base, such as CN=sales,DC=acme,DC=com.
9. Click one of the three buttons for LDAP attributes mappings, which control how LDAP attribute names equate, or map, to attribute names.

   Your attribute mapping selection will affect which default values appear in the three fields below. For example, the LDAP attribute `Login ID` maps to the user's login ID. If you select AD mappings, the default value is `sAMAccountName`. If you select AD Global Catalog mappings, the default value is `userPrincipalName`. If you select Common LDAP mappings, the default value is `uid`.
10. Click **Save**.

    The Security Console displays the *Authentication* page with the LDAP/AD authentication source listed.

To add a Kerberos authentication source:

1. Click **Add...** in the area of the Authentication page labeled *Kerberos Authentication sources*.

   The Security Console displays a box labeled *Kerberos Realm Configuration*.
2. Click the checkbox labeled **Enable authentication source**.
3. Click the appropriate checkbox to set the new realm that you are defining as the default Kerberos realm.

   The Security Console displays a warning that the default realm cannot be disabled.
4. Type the name of the realm in the appropriate text field.
5. Type the name of the key distribution center in the appropriate field.
6. Select the check box for every encryption type that your authentication source supports. During authentication, the source runs through each type, attempting to decrypt the client's credentials, until it uses a type that is identical to the type used by the client.
7. Click **Save**.

   The Security Console displays the *Authentication* page with the new Kerberos distribution center listed.

   Once you have defined external authentication sources, you can create accounts for users who are authenticated through these sources.
8. Click the **Administration** tab on the *Home* page.
9. Click **Create** next to *Users* on the *Administration* page,

   The Security Console displays the *User Configuration* panel.

   On the *General* page, the **Authentication** method drop-down list contains the authentication sources that you defined in the Security Console configuration file.
10. Select an authentication source.

    The built-in user store authentication is represented by the *Symantec CCS Vulnerability Manager user* option.

    The "Active Directory" option indicates the LDAP authentication source that you specified in the Security Console configuration file.

    If you select an external authentication source, the application disables the password fields. It does not support the ability to change the passwords of users authenticated by external sources.
11. Fill in all other fields on the *General* page.
12. Click **Save**.

NOTE: If you log on to the interface as a user with external authentication, and then click your user name link at the top right corner of any page, the Security Console displays your account information, including your password; however, if you change the password on this page, the application will not implement the change.

Manually setting Kerberos encryption types

If you are authenticating users with Kerberos, you can increase security for connections to the Kerberos source, by specifying the types of ticket encryptions that can be used in these connections. To do so, take the following steps:

1. Using a text editor, create a new text file named *kerberos.properties*.
2. Add a line that specifies one or more acceptable encryption types. For multiple types, separate each types with a character space:

   ```
   default_tkt_enctypes=<encryption_type encryption_type>
   ```

   You can specify any of the following ticket encryption types:

   - des-cbc-md5
   - des-cbc-crc
   - des3-cbc-sha1
   - rc4-hmac
   - arcfour-hmac
   - arcfour-hmac-md5
   - aes128-cts-hmac-sha1-96
   - aes256-cts-hmac-sha1-96

   Example:

   ```
   default_tkt_enctypes= aes128-cts-hmac-sha1-96 aes256-cts-hmac-sha1-96
   ```
3. Save the file in the installation_directory/nsc/conf directory.

   The changes are applied at the next startup.

Viewing Security Console database information

You can view the name and type of the Security Console database on the *Database* page of the Security Console configuration panel. You also can change displayed database settings.

To save the changes, click **Save**.

Changing default Scan Engine settings

The Security Console communicates with distributed Scan Engines over a network to initiate scans and retrieve scan results. If you want to obtain scan status information more quickly or reduce bandwidth or resource consumption required for Security Console-to-Scan-Engine communication, you can tune various settings on the Scan Engines page of the *Security Console Configuration* panel. See the following sections:

- *Configuring Security Console connections with distributed Scan Engines* on page *50*
- *Allocating threads for monitoring scans* on page *50*
- *Retrieving incremental scan results from distributed Scan Engines* on page *51*

Configuring Security Console connections with distributed Scan Engines

The Security Console establishes connections with distributed Scan Engines to launch scans and retrieve scan results. This communication can be disrupted by low network bandwidth, high latency, or situations in which Scan Engines are performing high numbers of simultaneous scans. If any of these conditions exist in your environment, you may want to consider increasing connection settings on the *Scan Engines* configuration page:

- The **Connection timeout** setting controls how long the Security Console waits for the creation of a connection with a distributed Scan Engine.
- The **Response timeout** setting controls how long the Security Console waits for a response from an Scan Engine that it has contacted.

NOTE: It is recommended that you consult with Technical Support before tuning these settings.

To configure these settings, take the following steps.

1. Go to the *Scan Engines* page in the *Security Console Configuration* panel.
2. Click the *Administration* tab.
3. On the *Administration* page, click **manage** for the Security Console.
4. Click **Scan Engines** in the *Security Console Configuration* panel.
5. Adjust the *Connections* settings.
6. Edit the value in the Connection timeout field to change the number of milliseconds that elapse before a connection timeout occurs.
7. Edit the value in the Response timeout field to change the number of milliseconds that elapse before the Security Console no longer waits for a response from an Scan Engine.
8. Click **Save** in the top bar of the panel to save the changes.
9. Restart the Security Console so that the configuration changes can take effect.

TIP: Because millisecond values can be difficult to read, a time value that is easier to read appears to the right of each value field. As you change either timeout value, note how the equivalent value changes.

Allocating threads for monitoring scans

The Security Console allocates a thread pool for retrieving scan status information. You can adjust the number of threads, which corresponds to the number of scan status messages that the Security Console can retrieve simultaneously. For example, if you increase the number of distributed Scan Engines and the number of scans running simultaneously, you can increase the threads in the pool so that the Security Console can retrieve more status messages at the same time.

NOTE: It is recommended that you consult with Technical Support before tuning these settings.

Keep in mind that retrieval time is subject to network conditions such as bandwidth and latency. Whenever the number of active threads in use exceeds the overall number of threads in the pool, the Security Console removes unused scan status threads after specific time interval. If you notice an overall decrease in the frequency of scan status messages, you may want to consider increasing the timeout value.

To adjust pool settings for scan status threads, take the following steps:

1. Go to the *Scan Engines* page in the *Security Console Configuration* panel.
2. Click the **Administration** tab.
3. Click **manage** for the Security Console on the *Administration* page.
4. Click **Scan Engines** in the *Security Console Configuration* panel.
5. Adjust the *Scan Status* settings.
6. Edit the value in the **Thread idle timeout** field to change the number of milliseconds that elapse before the Security Console removes unused scan threads.
7. Edit the value in the **Thread pool size** field to change the number of threads in the pool for monitoring scan status.
8. Click **Save** in the top bar of the panel to save the changes.
9. Restart the Security Console so that the configuration changes can take effect.

TIP: Because millisecond values can be difficult to read, a time value that is easier to read appears to the right of each value field. As you change either timeout value, note how the equivalent value changes.

Retrieving incremental scan results from distributed Scan Engines

The Security Console communicates with Scan Engines over a network to retrieve scan results. By default, the Security Console retrieves the full set of results after each scan completes. This can cause a significant, temporary increase in bandwidth usage, especially with scans that collect large sets of data. You can modulate bandwidth usage by enabling the Security Console to retrieve incremental results from distribute Scan Engines as scans are in progress.

If you enable this feature, you will see scan information from distributed Scan Engines on scan progress pages in the Security Console Web interface. By default, the Security Console displays incremental scan results only for the local Scan Engine that is installed on the same host as the Security Console.

NOTE: Enabling incremental scan results may increase bandwidth usage during a scan, but it reduces bandwidth usage at the end of a scan.

To enable incremental scan status retrieval for distributed Scan Engines, take the following steps:

1. Go to the *Scan Engines* page in the *Security Console Configuration* panel.
2. Click the **Administration** tab.
3. Click **manage** for the Security Console on the *Administration* page.
4. Click **Scan Engines** in the *Security Console Configuration* panel.
5. Enable the incremental scan results setting.
6. Click the check box labeled **Retrieve incremental scan results** from distributed Scan Engines.
7. Click **Save** in the top bar of the panel to save the changes.
8. Restart the Security Console so that the configuration change can take effect.

Viewing, activating, renewing, or changing your license

On the *Licensing* page, you can see license-related information about your Security Console. You also can activate a new license or start the process to modify or renew your license. Your Security Console must be connected to the Internet to activate your license.

NOTE: If your Security Console is not connected to the Internet *see Performing offline activation and updates* on page *67*.

The *License Activation* area displays general information about your license.

- If your license is *active*, you will see a link for contacting Symantec to modify your license, which is optional.
- If your license is *expired*, you will see a link for contacting Symantec to renew your license. You will need an active license in order to run scans and create reports.

To activate your license, you can take the following steps:

1. Click **Activate a New License** if you received a product key.

   The Security Console displays a text box.
2. Enter the key in the text box.

   You can copy the key from the e-mail.
3. Click **Activate**.

   You do not have to click **Save**. The application does not have to restart to complete the activation.

In the *License Details* area, you can see more information about your license:

- The value for *License Status* is one of four different modes, depending on the status of your license.
- The value for *Expiration* is the date that your current license expires.
- The value for *Max. Scan Engines* is the total number of internal Scan Engines that you can use. These Scan Engines can be installed on any host computers on your network.
- The value for *Max. Assets to Scan* is the total number of assets that you can scan with your internal Scan Engines.
- The value for *Max. Assets to Scan w/ Hosted Engine* is the total number of assets that you can scan with a Hosted Scan Engine.
- If the value for *SCADA Scanning* is *Enabled*, you can scan assets with the SCADA scan template. For a description of this template, see *Where to find SCAP update information and OVAL files* on page *99*.
- If the value for *Discovery Scanning* is *Enabled*, you can run discovery scans to determine what assets are available on your network without performing vulnerability checks on those assets.
- If the value for *PCI Reporting* is *Enabled*, you can create reports using the PCI Executive Overview and PCI Audit report templates. If this feature is disabled, it will still appear to be available in your scan template, but it will not be active during scans.
- If the value for *Web Application Scanning* is *Enabled*, you can scan Web applications with the spider. If this feature is disabled, it will still appear to be available in your scan template, but it will not be active during scans.
- If the value for *Policy Scanning* is *Enabled*, you can scan assets to verify compliance with configuration policies. If this feature is disabled, it will still appear to be available in your scan template, but it will not be active during scans.

Performing a backup and restore

Running regularly scheduled backup and restore routines ensures full recovery of the Security Console in the event of hardware failure. The application performs actual backup and restore procedures in maintenance mode. It cannot run these procedures while scans are in progress. See *Running in maintenance mode* on page *57*. However, you set up backup and restore operations while the application is in normal mode.

Important notes on backup and restore

There are four possible options on the backup/restore page:

- Backup data onto the application's file system.
- Restore an installation from a prior backup already on the application's file system.
- Copy an existing backup to external media using **Browse**.
- Restore an installation from a prior backup on external storage.

NOTE: You should copy backup data to external storage media to prevent loss in the event of a hardware failure.

What is saved and restored

A backup will save the following items:

- database
- configuration files (*nsc.xml*, *nse.xml*, *userdb.xml*, and *consoles.xml*)
- licenses
- keystores
- report images
- custom report templates
- custom scan templates
- generated reports
- scan logs

Backing up data

To back up data:

1. Go to the *Administration* page.
2. Click **Maintenance**.
3. Go to the *Maintenance—Backup/Restore* page
4. Type a short description of the new backup for your own reference.
5. Click **Start Backup**.

   A message appears indicating that if you proceed with the backup, it will then restart in Maintenance Mode.

   The Security Console will restart in maintenance mode and run the backup. If you're a global administrator, you can log on to monitor the backup process. You will see a page that lists each backup activity.
6. Click **Restart Server** when you see a notification that the backup is complete.

   A message appears indicating that any pending maintenance activities will be canceled if you proceed with the restart.
7. Click **Restart** at the bottom of that message.

Restoring data

The restore procedure will re-establish the Symantec CCS Vulnerability Manager system to its exact state immediately preceding the backup. This means that the application will retain the serial number, license limits, sites, templates, and all other relevant files and data.

If a hardware failure has rendered a system unusable, reinstall the application.

To restore Symantec CCS Vulnerability Manager:

1. Go to the *Administration* page.
2. Click **Maintenance**.
3. Go to the *Maintenance—Backup/Restore* page
4. Click **Browse...** to locate the file on the backup media.
5. Click **Start Upload and Restore**.

   As with backup, the application goes into maintenance mode to restore the file.
6. Restart the application in normal operating mode.

Performing database maintenance

You can initiate several maintenance operations to maximize database performance and drive space.

Database maintenance operations can take from a few minutes to a few hours, depending on the size of the database. Once you start these operations, the application shuts down and restarts in Maintenance Mode. Any in-progress scans or reports will stop before completion and any related data will be lost. You will have to rerun any reports or scans after the application completes maintenance operations and restarts in Normal Mode. For more information, see *Running in maintenance mode* on page 57.

To perform database maintenance:

1. Go to the *Administration* page.
2. Click **Maintenance**.
3. Go to the *Database Maintenance* page and select any of the following options:
   - **Clean up Database** removes leftover data that is associated with deleted objects such as sites, assets, or users.
   - **Compress Database** tables frees up unused table space.
   - **Reindex Database** rebuilds indexes that may have become fragmented or corrupted over time.
   - **Regenerate Statistics** refreshes scan statistics that appear in the Security Console Web interface.
4. Click **Start Database Maintenance**.

Running diagnostics

You can run several diagnostic functions to catch issues that may be affecting system performance.

Selecting diagnostic routines

To run diagnostics for internal application issues:

1. Click the **Administration** tab.
2. The Security Console displays the *Administration* page.
3. Click **Diagnose** next to *Troubleshooting*.

   The Security Console displays the *Troubleshooting* page.
4. Click the check box for each diagnostics routine you want to perform.

After performing the requested diagnostics, the Security Console displays a table of results. Each item includes a red or green icon, indicating whether or not an issue exists with the respective system component.

Sending scan logs

You can transmit logs generated by Scan Engines to Technical Support by clicking **Send Logs** on the Troubleshooting page.

NOTE: An optional SMTP (e-mail) transport mechanism is also supported when a direct link is unavailable. Contact Technical Support for more information.

To send scan logs:

1. Click **Send Logs** on the *Troubleshooting* page.

   The Security Console displays a box for uploading the logs.
2. Select an upload method from the drop-down list.
   - **HTTPS upload**. The application encrypts the logs using PGP before sending them directly over an SSL connection to the Symantec DMZ, and subsequently to the support database. This method bypasses third-party servers.
   - **SMTP**. You can e-mail the reports. Contact Technical Support to inquire about this option before attempting to use it.
3. Type a message to send with the logs.

   The message may refer to scan errors, a support case, or a report of aberrant system behavior.
4. Click **Send Logs**.

Running in maintenance mode

NOTE: Only global administrators are permitted to run the application in maintenance mode.

Maintenance mode is a startup mode in which the application performs general maintenance tasks and recovers from critical failures of one or more of its subsystems. During maintenance mode, you cannot run scans or reports. Available functions include logging, the database, and the Security Console Web interface.

NOTE: The application automatically runs in maintenance mode when a critical internal error occurs.

When the application is running in maintenance mode, you see the page */admin/maintenance/index.html* upon logging on. This page shows all available maintenance tasks and indicates the current status of the task that is being performed. You cannot select a new task until the current task is completed. Afterward, you can switch tasks or click **Restart** to return to normal operating mode.

To work in Maintenance mode:

1. Click the *Administration* tab.
2. On the *Administration* page, click **Maintenance**.
   The Security Console displays the *Maintenance Mode* page.

Migrating the database

The the application's database is a core component for all major operations, such as scanning, reporting, asset and vulnerability management, and user administration. The efficiency with which it executes these tasks depends heavily on database performance. The current PostgreSQL database version, 9, features a number of performance and stability enhancements. The application takes full advantage of these improvements to scale flexibly with the needs of your environment. Future releases will include powerful features that will require the latest PostgreSQL version.

NOTE: Only administrators can migrate the database.

If your installation is running an earlier version of PostgreSQL, you can easily migrate it to the latest version, using a tool in the Security Console Web interface.

Migration involves five required tasks:

1. *Preparing for migration* on page *58*
2. *Starting and monitoring the migration* on page *60*
3. *Verifying the success of the migration* on page *61*
4. *Ensuring database consistency* on page *62*
5. *Backing up the post-migration database* on page *63*

Restoring a backup of a PostgreSQL 8.2 database after migrating to PostgreSQL 9 is not supported. After you perform and verify the migration to PostgreSQL 9 and ensure database consistency, it is very important that you back up the database immediately to prevent the need to restore an earlier version of the database. See *Backing up the post-migration database* on page *63*.

This document also provides instructions for optional post-migration tasks:

- restoring backups
- restoring tuned PostgreSQL settings

Preparing for migration

Some preparation will ensure that migration takes the least possible amount of time and is successful:

- Make sure port 5431 is open, as this is the port reserved for the migration. If you attempt to run the migration and you see a message on the migration page that port 5431 is in use, contact your system administrator to determine what service is running on port 5431. Shut that service down until migration is complete.
- Make sure you have sufficient disk space. During the migration, the old database is backed up and a new one is created, so you need to accommodate both databases. If you are unable to run the migration because of insufficient disk space, you can take several steps to free up space. See *Freeing up disk space for migration* on page *59*.
- Make sure that you have an authenticated global administrator account, so that you can monitor the migration in the Security Console Web interface. If your account is authenticated by an external source, such as LDAP or Kerberos, you will be able to start the migration. However, when the application restarts in Maintenance Mode, you will not be able to log on to the Security Console Web interface to monitor the migration. The database stores information about the external authentication sources, and it won't be operational during the migration. If you do not monitor the migration, you will not know if any issues have occurred requiring you to restart it or take some other action. You have several options for monitoring the migration:
- When the application restarts in Maintenance Mode, log on with an authenticated global administrator account instead of an external source. This will allow you to monitor status messages in the Security Console Web interface. If you do not have an authenticated account, you can create one or modify an existing account accordingly. See *Managing users, roles, and permissions* on page *31*. You also can log on with the default administrator account that was created during installation.
- If you do not have an authenticated administrator account you can monitor the migration by reading the nsc.log file, which is located in [installation_directory]/nsc. If you need to restart the application manually, you can do so from the command prompt using the **`restart`** command.

Freeing up disk space for migration

In most cases the **Start Migration** button is disabled if you do not have enough disk space to perform the migration. However, in some environments, the **Start Migration** button may be enabled but disk space issues may still occur during migration. For example, see the section about Linux file systems. To free up disk space, try the solutions listed in the following sequence. Check if the **Start Migration** button is enabled after each step.

NOTE: It is recommended that your free disk space be equal to **1.6 GB + (1.3 x database_size)**.

Run the following database maintenance tasks, which remove unnecessary data and free up unused table space:

- re-indexing the database
- cleaning up the database
- compressing tables

If you have not run these tasks recently, doing so may free up considerable space. It is recommended that you run each task individually in the following sequence. After completing each task, try running the migration again. Re-indexing may provide all the space you need, making it unnecessary to clean up the database or compress tables, which can take significant time depending on the size of your database. See *Performing database maintenance* on page *55*.

Move the following directories from the host system, and restore them after the migration is complete:

NOTE: These directories and files take up increasing amounts of disk space over time as the application accumulates data.

- backup files—[installation_directory]/nsc/*.bak and—[installation_directory]/nsc/*.zip
- reports directory—[installation_directory]/nsc/htroot/reports
- access log directory, including the Tomcat log subdirectory—[installation_directory]/nsc/logs
- scan event data files directory—[installation_directory]/nse/scans
- Security Console logs—[installation_directory]/nsc/*.log and /nsc/*.log*
- PostgreSQL log—[installation_directory]/nsc/nxpsql/nxpgsql.log
- scan logs directory—[installation_directory]/nse/nse.log and /nse/nse.log.

To create free space for migration:

1. Move from the host system any files or directories extraneous to the application that are not required by other applications to run. You can restore them after the migration is complete.
2. Delete the contents of the java.io.tmpdir directory on the host system. The location depends on the operating system.

NOTE: If the disk space problem occurs after a previous migration attempt failed, see *Addressing a failed migration* on page *64*.

After taking the preceding steps, try starting the migration again. If you still don't have enough disk space, contact Technical Support.

Creating free space in Linux

By default, Linux file systems reserve 5 percent of disk space for privileged, or root, processes. Although this reserved space is not available for database migration, the application includes it in the pre-migration calculation of total available space. As a result, a migration may be able to start, but then fail, because the actual amount of free disk space is lower than what was detected in the calculation. You can lower the amount of reserved disk space to make the amount of actual free space more consistent with the results of the pre-migration calculation. To do so, use the tune2fs utility. The command includes parameters for the percentage of reserved disk space and the partition on which the application is installed.

Example**: `tune2fs -m 1 /dev/sdf1`**

Starting and monitoring the migration

To monitor the migrations:

1. Go to the *Administration* page.
2. Click **Maintenance**.
3. Go to the *Database Migration* page.
4. Review your database migration status on the migration page.
5. Click **Start Migration** if it indicates that your installed PostgreSQL version is earlier than 9.0.3 and that your system is ready for migration.

After you click **Start Migration**, the application goes into Maintenance Mode. Normal operations, such as scanning and running reports, are unavailable. See *Running in maintenance mode* on page 57. If you're an administrator, you can log on to monitor migration status messages.

During migration, the application copies data from the PostgreSQL 8.2 database to the new PostgreSQL 9 database. The migration requires enough disk space for both of these databases.

It also backs up the PostgreSQL 8.2 database and stores it in the directory [installation_directory]/nsc/nxpgsql-backup-[timestamp] after the migration completes.

The estimated migration time is based on the size of the database.

After all migration processes finish, the application restarts, and you can resume normal operations.

The PostgreSQL database applies its default settings after migration. If you modified *postgresql.conf* or *pg_hba.conf* prior to the migration you will need to reapply any custom settings to those configuration files. See *Restoring custom PostgreSQL settings* on page 64. You can refer to the modified configuration files in the old, archived version of PostgreSQL for custom settings.

If you click the **Cancel** button to stop the migration before it completes, the application will discontinue the migration process. You can then restart the application in normal, operational mode.

If the migration fails, your current version of the PostgreSQL database will remain intact, and you can continue using the application without interruption. See *Addressing a failed migration* on page 64.

In very rare instances the application may display the migration FAQs while in Maintenance Mode after the migration process has been executed, instead of a status message detailing the results of the migration. If this occurs, contact Technical Support for assistance before restarting the server. You should also contact Technical Support if this situation occurred and you inadvertently restarted your server, or at any time after the migration if you note on the Database Migration page that the version of PostgreSQL running in your environment is earlier than 9.0.3.

Addressing migration that takes a long time with no new status messages

Depending on the amount of data, the PostgreSQL 8.2 to PostgreSQL 9 migration can take several hours or even days. Therefore, a long migration time is not unusual, and extended periods without new status messages do not necessarily indicate that the migration is "hanging."

You can perform a couple of quick checks to confirm that the migration is still proceeding when no status messages are visible:

1. Run top in Linux or Task manager in Windows, and check if a PostgreSQL process is running and using CPU resources.
2. Check migration log files, located in [installation_directory]/nsc/nxpgsql/pgsql/pg_transfer_8.2_dump.log, for messages about database tables being copied.

The Security Console will display a notification when the processes have completed.

Verifying the success of the migration

To verify migration success, take the following steps:

1. Go to the *Administration* page.
2. Click **Maintenance**.
3. Go to the *Database Migration* page.
4. Read the installed version of PostgreSQL, which is displayed on the page.

   If the migration was successful, the installed version will be 9.0.3.

   OR
5. Open the nsc.log file, located in [installation_directory]/nsc, to verify that PostgreSQL 9 is running.
6. Search for the string *PostgreSQL*. You will find the active PostgreSQL version number with an instance of that string.

   It will appear on a line that looks like the following example:

   ```
   NSC 2011-03-11T18:45:01 PostgreSQL 9.0.3, compiled by Visual C++ build 1500, 64-bit
   ```

Upon confirming that the migration was successful, take the following steps:

1. Move back any files or directories that you moved off the host system to free up disk space for migration, See *Freeing up disk space for migration* on page *59*.
2. Move the [installation_directory]/nsc/nxpgsql-backup-[timestamp] directory to an external location for storage.

   It contains the pre-migration database, including the postgresql.conf file.

NOTE: Before you resume normal operations, make sure to verify database consistency as described in the following section.

If you modified postgresql.conf or pg_hba.conf prior to the migration you will need to reapply any custom settings to those configuration files. See *Restoring custom PostgreSQL settings* on page *64*. You can refer to the modified configuration files in the old, archived version of PostgreSQL for custom settings.

Ensuring database consistency

This procedure involves two, or possibly three, steps. The first two, checking database consistency and cleaning up the database, take little time.

You would only perform the third step, regenerating statistics, if the consistency check detects data duplication. Regenerating statistics may take a long time depending on your scan data. If you have consistency issues and do not regenerate statistics, your reports and the Web interface may display inaccurate information. In this event, regenerating statistics will ensure that the most recent scan statistics are correct. If you prefer, you can regenerate statistics during a separate maintenance window. Not regenerating statistics at this time will have no impact on the accuracy of statistics gathered in the future.

To verify database consistency and respond appropriately:

1. Go to the *Administration* page.
2. Click **Diagnose**.

   the Security Console displays the *Troubleshooting* page.
3. Select only the **Database Diagnostics** check box.
4. Click **Perform Diagnostics**.

   A table appears on the page, listing the results of all diagnostic tests. Read the results for duplication-related tests: *Node Synopsis Duplication*, *Scan Synopsis Duplication*, *Device Synopsis Duplication*, *Site Synopsis Duplication*, *Group Synopsis Duplication*, *Device Nodes Duplication*, *Device Scans Duplication*, and *Site Scans Duplication.* Red circles containing the letter *X* indicate consistency issues.
5. Go to the *Administration* page.
6. Click **Maintenance**.

   The Security Console displays the *Database Maintenance* page.
7. (Optional) Select only the **Clean up Database** task to remove any unnecessary data if the diagnostics in the preceding step indicated no duplication.
8. (Optional) Select both **Clean up Database** *and* **Regenerate Statistics** if the diagnostics in the preceding step indicated *any* duplication.
9. Click **Start Database Maintenance**.

NOTE: All diagnostics options are selected by default, but only database diagnostics are necessary for verifying database consistency after migration. To see only the information you need for this task, clear any other selected check boxes.

Once you start these operations, the application shuts down and restarts in Maintenance Mode. Any in-progress scans or reports will stop before completion and any related data will be lost. You will have to rerun any reports or scans after it completes maintenance operations and restarts in Normal Mode. For more information, see *Running in maintenance mode* on page *57*.

If after regenerating statistics you notice any discrepancies in reports or the Security Console Web interface, contact Technical Support.

Backing up the post-migration database

It is very important that you back up the database immediately after you verify the success of the migration and ensure database consistency. This preserves a baseline instance of the post-migration database and prevents the need to restore a backup of a PostgreSQL 8.2 database, which is not supported.

Perform this step only after you have completed the preceding steps:

- migrating the database
- verifying the success of the migration
- ensuring database consistency

For instructions on backing up the database, see *Performing a backup and restore* on page *53*.

Restoring a database that was backed up as part of a migration

After migration, the application backs up the PostgreSQL 8.2 database and stores it in the directory [installation_directory]/nsc/nxpgsql-backup-[timestamp].

If you want to restore this particular database, take the following steps:

1. Shut down the application.
2. Rename the pgsql directory of the post-migration database.

   It is located in [installation_directory]/nsc/nxpgsql.
3. Copy the backup directory, named nxpgsql-backup-[timestamp], into the [installation_directory]/nsc directory, and rename it nxpgsql.
4. Start the application and resume operations.

TIP: Move the backup directory with all original permissions attributes preserved. Doing so prevents the requirement on Linux that nxpgsql be the owner of the directory, as well as the necessity on Windows to give the system user access to the directory.

If you are planning to restore the database that was backed up during the migration, keep several things in mind:

If you run scans or reports after the migration and then restore the backup database, the Security Console Web interface will not list scan or report instances from the period between the migration and the restoration because the restored database does not contain those records.

When you start to run scans or reports after the restoration, the associated scan or report instances that are being populated in the restored database will overwrite the instances that were generated in the file system prior to the restoration.

Graphic charts will initially be out of synch with the restored database because they always reflect the latest site, scan, or asset group information. Each chart will refresh and synchronize with the restored database after an event associated with it. For example, running a scan will refresh and synchronize the charts for any associated sites or asset groups.

Restoring custom PostgreSQL settings

WARNING: Do not simply copy the old configuration files into the new database location. This may prevent the database from starting due to compatibility issues. For each file, compare each setting one by one, and edit only the properties that you modified in the PostgreSQL 8.2 installation.

The PostgreSQL database applies its default settings after migration. If you previously modified the postgresql.conf file to tune database performance or the pg_hba.conf to enable remote connections to the database, you will need to reapply those modified settings.

After the migration is complete, you can refer to the configuration files in the old, archived version of PostgreSQL, which is stored in the directory [instsallation_directory]/nsc/nxpgsql-backup-[time-stamp].

Addressing a failed migration

If the migration fails, your current version of the PostgreSQL database will remain intact. Simply restart the application and resume normal operations.

Before you run the migration again, find out if the failure occurred due to disk space errors. In certain cases, migration may exceed available disk space before finishing, even if the automatic pre-migration check determined that sufficient disk space was available.

To troubleshoot a failed migration:

- Check your available space for the disk that the application is installed on.
- (Optional) In Windows, right-click the icon for the disk and then click **Properties** from the pop-up menu. Read the amount of available disk space.
- (Optional) In Linux, run the command to show disk space: **`df -h`** in the [installation_directory]/nsc directory. Read the amount of available disk space.
- If the available disk space is less than the database size, free up disk space, and run the migration again. See *Freeing up disk space for migration* on page *59*.
- If your free disk space is equal to at least 1.6 GB + (1.3 x database_size), this suggests that the migration did not fail due to disk space issues but for a different reason. Contact Technical Support for assistance in completing the migration.

NOTE: If you do not wish to retry migration after failure, you should still delete the /nxpgsql-temp directory, because it uses up considerable disk space.

If the migration fails due to a system failure or power outage and you attempt to run the migration again, you may encounter a disk space limitation issue. This is because during the failed migration attempt, the application created an nxpgsql-temp directory. Simply delete this directory and start the migration again. The temp directory is located in [installation_directory]/nsc.

Enabling FIPS mode

If you are operating Symantec CCS Vulnerability Manager in an environment where the use of FIPS-enabled products is mandatory, or if you want the security of using a FIPS-certified encryption module, you should enable FIPS mode. The application supports the use of Federal Information Processing Standard (FIPS) 140-2 encryption, which is required by government agencies and companies that have adopted FIPS guidelines.

What is FIPS?

The FIPS publications are a set of standards for best practices in computer security products. FIPS certification is applicable to any part of a product that employs cryptography. A FIPS-certified product has been reviewed by a lab and shown to comply with FIPS 140-2 (Standard for Security Requirements for Cryptographic Modules), and to support at least one FIPS-certified algorithm.

Government agencies in several countries and some private companies are required to use FIPS-certified products.

What is FIPS mode?

FIPS mode is a configuration that uses FIPS-approved algorithms only. When the application is configured to operate in FIPS mode, it implements a FIPS-certified cryptographic library to encrypt communication between the Security Console and the user for both the browser and API interfaces.

FIPS mode considerations

It is important to note that due to encryption key generation considerations, the decision to run in FIPS mode or non-FIPS mode is irrevocable. The application must be configured to run in FIPS mode immediately after installation and before it is started for the first time, or else left to run in the default non-FIPS mode. Once the application has started with the chosen configuration, you will need to reinstall it to change between modes.

Activating FIPS mode

When Symantec CCS Vulnerability Manager is installed, it is configured to run in non-FIPS mode by default. The application must be configured to run in FIPS mode before being started for the first time. See *Activating FIPS mode in Linux* on page *65*.

When FIPS mode is enabled, communication between the application and non-FIPS enabled applications such as Web browsers or API clients cannot be guaranteed to function correctly.

Activating FIPS mode in Linux

You must follow these steps after installation, and BEFORE starting the application for the first time.

1. Install rng-utils.

   The encryption algorithm requires that the system have a large entropy pool in order to generate random numbers. To ensure that the entropy pool remains full, the rngd daemon must be running while the application is running. The rngd daemon is part of the rng-utils Linux package.
2. Download and install the rng-utils package using the system's package manager.
3. Run the command **`rngd -b -r /dev/urandom`**.
4. Create a properties file for activating FIPS mode.
5. Create a new file using a text editor.
6. Enter the following line in this file:

   `fipsMode=1`
7. Save the file in the [install_directory]/nsc directory with the following name:

   CustomEnvironment.properties
8. Start the Security Console.

TIP: Add the rngd command to the system startup files so that it runs each time the server is restarted.

Activating FIPS mode in Windows

You must follow these steps after installation, and before starting the application for the first time.

1. Create a properties file for activating FIPS mode.
2. Create a new file using a text editor.
3. Enter the following line in this file:

   `fipsMode=1`
4. Save the file in the [install_directory]\nsc directory with the following name: CustomEnvironment.properties
5. Start the Security Console.

NOTE: You can disable database consistency checks on startup using the CustomEnvironment.properties file. Do this only if instructed by Technical Support.

Verifying that FIPS mode is enabled

To ensure that FIPS mode has been successfully enabled, check the Security Console log files for the following messages:

```
FIPS 140-2 mode is enabled. Initializing crypto provider
Executing FIPS self tests...
```

Performing offline activation and updates

By default, the application is configured to activate its license and download updates through an Internet connection. Some business environments have security policies that do not permit sensitive assets to be exposed to the Internet. For such environments, you can activate and update the application without connecting it directly to the Internet. This procedure is known as "offline activation and updates."

How activation and updates occur

A look at the activation and update process will help you understand how the offline procedure works.

How online activations and updates are initiated

When the Security Console is connected to the Internet, updates can be initiated three different ways:

- The Security Console automatically runs an update at regular time intervals.
- The Security Console restarts.
- A Global Administrator runs the **`update now`** command. See *Using the command console* on page *72*.

In all three cases, the transfer of applicable update files to the application and the installation of these files happen automatically.

Activation and updates occur in the same process

Online activation and updates occur during the same process:

1. The Security Console initiates an HTTP exchange with the Update Server.
2. The Update Server transfers to the Security Console all required files for that Security Console and for any Scan Engines with which it is paired.

   In environments where multiple Security Consoles share a paired Scan Engine, *only* the first Security Console to contact the Update Server will receive files for the mutually paired Scan Engine.
3. When the Security Console receives the complete set of files, it sends a message of acknowledgement to the Update Server. This acknowledgement ensures that the Update Server never transfers *these* files to *this* Security Console again. If this Security Console queries the Update Server again before new updates become available, the Update Server will respond with a message that no updates are available.

Understand and perform the offline procedure

The offline procedure works in a similar way to the online process, except that an Internet-connected proxy computer communicates with the update server instead of the Security Console.

NOTE: This procedure results in license activation for new installations and updates for installations that were previously activated.

The offline procedure requires the following items:

- the Security Console
- an Internet-connected computer that acts as the Security Console's proxy

With online updates, required information is automatically gathered for an update request and sends it to the Update Server without storing it. With offline updates, only the gathering of the information is automatic.

NOTE: The transfer of files must be performed manually.

When the Security Console initiates an update in offline mode, it creates directories and files within the [installation_path]/updates/offline directory that provide necessary update information for the proxy computer. To enable this offline mode, you must make a configuration change to the Security Console.

The directory contains the following items:

- the *update.properties* file, which provides update instructions
- two keystores, which generate keys that authenticate the Security Console to the update server

When initiating offline mode, the Security Console also creates *updatetable,* which lists the required update files.

WARNING: Do not make any changes to the directory structure. The Security Console and the computer that will be used as its proxy both parse this directory structure.

If the Security Console is paired with any Scan Engines, the directory will include a similarly structured sub-directory for each Scan Engine.

Coordinate updates with Scan Engine pairing

For brand-new installations that have not yet been activated, if you want to pair the Security Console with Scan Engines, you will first need to perform the offline procedure to activate the license on the Security Console and apply updates. Afterward, you can pair Scan Engines with the Security Console. See the topic *Configuring Scan Engines* in the user's guide.

NOTE: If you do not click **Update** for a newly paired Scan Engine before performing an offline update, the update attempt will fail.

After you pair a Scan Engine, and before you perform an offline update, make sure to click **Update** for the paired Scan Engine in the Security Console Web interface. See the topic on Scan Engine configuration in the user's guide. Doing so activates the license on the newly paired Scan Engine and provides the Security Console with information it will require for future updates of the Scan Engine.

You will need to perform the offline procedure once again to update the Scan Engines.

For deployments that already have been activated, make sure to pair the Security Console with any Scan Engines before performing the offline procedure to obtain updates.

Prepare the proxy computer

The proxy computer requires the same Java Runtime Environment (JRE) as Symantec CCS Vulnerability Manager. Using a different version is not supported and may cause activation and updates to fail.

These steps need to be performed only once for a proxy computer:

1. Select an installation with the same architecture and operating system as the proxy computer.
2. Copy the JRE files from the installation selected in step 1 to the proxy computer.
3. Copy the [installation_path]/_jvm directory from the Security Console host to the proxy computer. For future reference, the destination directory path will be represented by the variable [proxy_env_path].
4. Copy the following files from [installation_path]/shared/lib/ to [proxy_env_path]:
   - r7shared.jar
   - sslj.jar
   - cryptojFIPS.jar
   - slf4j-api-1.6.4.jar

WARNING: Do not make any changes to the files or directories.

NOTE: You must have root privileges to enable and perform offline activations and updates. You can log on as root, or begin each command with `sudo`.

By default, the application is configured to perform activation and updates automatically with an Internet connection. You must configure it for the offline procedure by taking the following steps.

The application must have run at least once before you can configure it. If you have not yet started the Security Console, start it before performing the following steps.

To enable offline activation:

1. Change an attribute in the nsc.xml file to enable offline updates.

   This file controls the Security Console configuration settings, including if, when, and how the Security Console initiates updates.
2. Stop the application.
3. Open the configuration file [installation_path]/nsc/conf/nsc.xml
4. Find the tag for the `AutoUpdate` element. Set the value for the method attribute from `http` to `offline`.
5. Ensure that the `enabled` attribute is set to **`1`** to support offline updates.

   ```
   AutoUpdate enabled="1" manual="0" method="offline"
   ```
6. Save and close the file.
7. Restart the application.

   The configuration change and restart should generate Security Console output that includes the following content:

   ```
   Offline update configured to write to: [installation_path]/updates/offline.
   ```

   If the license is not activated, the application will attempt to initiate activation as it starts. If the Security Console is not connected to the Internet, the activation attempt will generate Security Console output with the following content:

   ```
   Product activation failed (No route to host: no further information)
   ```

   These messages do not indicate an offline activation failure.
8. Repeat this procedure for each Security Console if you are using multiple Security Consoles.

Transfer files required for update requests to a proxy computer

After restarting, the application generates the offline directory. You need to transfer these files to a computer with an Internet connection. With these files, the computer contacts the Update Server as a proxy for the Security Console.

To transfer files:

1. Copy the [installation_path]/updates/offline directory and its contents from the Security Console host into a directory on the proxy computer.

   For future reference, the destination directory path will be represented by the variable [proxy_update_path].
2. Delete the proxy_update_path directory and its contents after the offline update procedure is complete for each Security Console.

   If you are using the same proxy computer for multiple Security Consoles, you will need to delete proxy_update_path directory and its contents after performing each update for each Security Console.

Initiate the update request from the proxy computer

If you are using multiple Security Consoles, repeat these steps for each one.

To initiate the update request:

1. Run the following command on the proxy computer from [proxy_env_path] to initiate the update request.

   In Windows:

   ```
   _jvm\bin\java.exe -cp r7shared.jar;sslj.jar;cryptoj-FIPS.jar;slf4j-api-1.6.4.jar com.rapid7.updater.OfflineUpdateProxy [proxy_update_path]
   ```

   In Linux:

   ```
   _jvm/bin/java -cp r7shared.jar:sslj.jar:cryptoj-FIPS.jar:slf4j-api-1.6.4.jar com.rapid7.updater.OfflineUpdateProxy [proxy_update_path]
   ```

   The proxy computer contacts the Update Server, which uses the keystores to verify the proxy computer's identity (as the "console"). Then, the Update Server transfers appropriate files to the proxy computer.

NOTE: If there is no output from the command, then no offline updates were attempted. Ensure that `[proxy_update_path]` was specified correctly.

2. Inspect the output from the command to verify that the offline update proxy command ran correctly.

NOTE: It is strongly recommended that you back up the update files after the transfer from the Update Server.

Apply the license activation or updates to the Security Console(s)

To apply the license activation or updates to the Security Console(s) you must make sure the Security Console is still running. If you are using multiple Security Consoles, repeat the following procedure for each one.

NOTE: License activation and updates cannot occur simultaneously in offline environments.

To apply the license activation:

1. Copy the appropriate update directory from [proxy_update_path] over the [installation_path]/updates/offline directory of the Security Console host.
2. Run the `update now` command once to activate the license on the Security Console, if necessary, and once to apply any updates.
3. Apply the update to any paired Scan Engines.
4. Start a browser, and go to the URL of the Security Console Web interface: https://[console_host_url]:3780.
5. Log onto the application.
6. Click the **Administration** tab.

   The Security Console displays the *Administration* page.
7. Click **Manage** for Scan Engines.
8. Click **Update** for each Scan Engine.

   Alternatively, you can use the `update engines` Security Console command to update all paired Scan Engines in one step.

Verify the success of an activation or update

To verify the success of an activation:

1. Start a browser, and go to the URL of the Security Console Web interface: https://[console_host_url]:3780.
2. Log on to the application.
3. Click the **Administration** tab.

   The Security Console displays the *Administration* page.
4. Click **Manage** for the Security Console.

   The Security Console displays the *Security Console Configuration* panel.
5. Click **Licensing**.
6. Verify that the license status is *activated*.
7. (Optional) Use the **`show licenses`** Security Console command to see the license status.

To verify the success of an update:

1. Start a browser, and go to the URL of the Security Console Web interface: https://[console_host_url]:3780.
2. Log on to the application.
3. Click the **Administration** tab.

   The Security Console displays the *Administration* page.
4. Click **Manage** for the Security Console.

   The Security Console displays the *Security Console Configuration* panel.
5. Click **General**.
6. Verify that the date and time for *Last Update* are current.
7. (Optional) Use the `version` Security Console command to see the last update ID.

After the offline update procedure is complete for each Security Console, make sure to delete the proxy_update_path directory and its contents. If you are using the same proxy computer for multiple Security Consoles, you will need to delete proxy_update_path directory and its contents after performing each update for each Security Console.

Using the command console

If you are a Global Administrator, you can perform certain Security Console operations using the command console. You can see real-time diagnostics and a behind-the-scenes view of the application when you use this tool.

You can type `help` to see a list of all available commands and their descriptions. For more detailed information, see *Available commands* on page *73*.

Accessing the command console

Global Administrators have access to the Security Console to perform administrative functions. For a list of commands, see *Available commands* on page *73*.

Using the command console

1. Click the **Administration** tab in the Security Console Web interface.

   The Security Console displays the *Administration* page.
2. Click the link to **Run** console commands, which is displayed with the *Troubleshooting* item.

   The command console page appears with a box for entering commands.
3. Enter a command.
4. Click **Execute**.

In Linux

To use the Security Console Web interface in Linux:

1. Start a console screen session if one is not already in progress.
   If the host is remote, use SSH to log on first.
2. Type commands and click **ENTER**.

Available commands

A list of available commands follows. Text in square brackets contain optional parameters, as explained in the action descriptions. Text in arrow brackets contain variables.

Command	Action
`activate`	Activate the application with a product key.
`database diagnostics`	Check the database for inconsistencies like multiple entries for an asset.
`[show] diag[nostics]`	Display diagnostic information about the Security Console.
`exit`	Stop the Security Console service.
`garbagecollect`	Start the garbage collector, a Java application that frees up drive space no longer used to store data objects.
`get property [<name>]`	View the value assigned to a parameter associated with the Scan Engine. Example: `get property os.version`. The Security Console would return: `os.version=5.1.` If you type `get property` without a parameter name, the Security Console will list all properties and associated values. You can view and set certain properties, such as the IP socket number, which the application uses for communication between the Security Console and the Scan Engine. Other properties are for system use only; you may view them but not set them.
`heap dump`	"Dump" or list all the data and memory addresses "piled up" by the Java garbage collector. The dump file is saved as heap.hprof in the nsc directory.
`help`	Display all available commands.
`license request from-email-address [mail-relay-server]`	E-mail a request for a new license. The email-address parameter is your address as the requestor. The optional `mail-relay-server` parameter designates an internally accessible mail server to which the license server should connect to send the e-mail. After you execute this command, the application displays a message that the e-mail has been sent. When you receive the license file, store it in the `nsc/licenses` directory without modifying its contents. Licenses have a `.lic` suffix.
`log rotate`	Compress and save the nsc.log file and then create a new log.
`ping host-address [tcp-port]`	Ping the specified host using an ICNMP ECHO request, ICP ACK packet, and TCP SYN packet. The default TCP port is 80.
`quit`	Stop the Security Console service.
`restart`	Stop the Security Console service and then start it again.
`[show] scan configs`	Show all defined scan configurations.
`[show] schedule`	Display the currently scheduled jobs for scans, auto-update retriever, temporal risk score updater, and log rotation.

Command	Action
`send support [from-email-address] [mail-relay-server] [message-body]`	Send logs generated by the Security Console and Scan Engine(s) for troubleshooting support. By default, the application sends the request to a log server via HTTPS. Alternatively, you can e-mail the request by specifying a sender's e-mail address or outbound mail relay server. You also can type a brief message with the e-mail request. When you execute the command, the Security Console displays a scrolling list of log data, including scheduled scans, auto-updates, and diagnostics.
`server diagnostics`	Display diagnostic information that may be useful for debugging or simply monitoring the application.
`show licenses`	Display information about all licenses currently in use. Multiple licenses may operate at once.
`show locked accounts`	List all user accounts locked out by the Security Console. The application can lock out a user who attempts too many logons with an incorrect password.
`show mem`	List statistics about memory use.
`[show] threads`	Display the list of active threads in use.
`traceroute host-address`	Determine the IP address route between your local host and the host name or IP address that you specify in the command. When you execute this command, the Security Console displays a list of IP addresses for all "stops" or devices on the given route.
`unlock account <name>`	Unlock the user account named in the command.
`update now`	Check for and apply updates manually and immediately, instead of waiting for the Security Console to automatically retrieve the next update.
`update engines`	Send pending updates to all defined Scan Engines.
`[ver] version`	Display the current software version, serial number, most recent update, and other information about the Security Console and local Scan Engine. Add "console" to the command to display information about the Security Console only. Add "engines" to the command to display information about the local Scan Engine and all remote Scan Engines paired with the Security Console.

Troubleshooting

This chapter provides descriptions of problems commonly encountered when using the application and guidance for dealing with them. If you do need to contact Technical Support, this guide will help you gather the information that Support needs to assist you.

Working with log files

If you are encountering problems with the Security Console or Scan Engine, you may find it helpful to consult log files for troubleshooting. Log files can also be useful for routine maintenance and debugging purposes.

The section does not cover the scan log, which is related to scan events. See *Viewing the scan log* on page *80*.

Locating each log file and understanding its purpose

Log files are located in [installation_directory]/nsc/logs directory on the Security Console and [installation_directory]/nse/logs on Scan Engines. The following log files are available:

NOTE: In earlier product versions, API information was stored in nsc.log.

- *access.log* (on the Security Console only): This file captures information about resources that are being accessed, such as pages in the Web interface. At the INFO level, access.log captures useful information about API events, such as APIs that are being called, the API version, and the IP address of the API client. This is useful for monitoring API use and troubleshooting API issues. The file was called *access_log* in earlier product versions.
- *auth.log* (on the Security Console only): This file captures each logon or logoff as well as authentication events, such as authentication failures and lockouts. It is useful for tracking user sessions. This file was called *um_log* in earlier product versions.
- *nsc.log* (on the Security Console only): This file captures system- and application-level events in the Security Console. It is useful for tracking and troubleshooting various issues associated with updates, scheduling of operations, or communication issues with distributed Scan Engines. Also, if the Security Console goes into Maintenance Mode, you can log on as a global administrator and use the file to monitor Maintenance Mode activity.
- *nse.log* (on the Security Console and distributed Scan Engines): This file is useful for troubleshooting certain issues related to vulnerability checks. For example, if a check produces an unexpected result, you can look at the nse.log file to determine how the scan target was fingerprinted. On distributed Scan Engines only, this file also captures system- and application-level events not recorded in any of the other log files.
- *mem.log* (on the Security Console and distributed Scan Engines): This file captures events related to memory use. It is useful for troubleshooting problems with memory-intensive operations, such as scanning and reporting.

Structure and contents of log files

Log files have the following format:

```
[yyyy-mm-ddThh:mm:ss GMT] [LEVEL] [Thread: NAME] [MESSAGE]
```

Example:

```
2011-12-20T16:54:48 [INFO] [Thread: Security Console] Security Console
started in 12 minutes 54 seconds
```

The date and time correspond to the occurrence of the event that generates the message.

Every log message has a severity level:

Level	Meaning	Example
ERROR	an abnormal event that prevents successful execution of system processes and can prevent user operations, such as scanning	the Security Console's failure to connect to the database
WARN	an abnormal event that prevents successful execution of system processes but does not completely prevent a user operation, such as scanning	disruption in communication between the Security Console and a remote Scan Engine
INFO	a normal, expected event that is noteworthy for providing useful information about system activity	the Security Console's attempts to establish a connection with a remote Scan Engine
DEBUG	a normal, expected event that need not be viewed except for debugging purposes	the execution of operations within the Security Console/Scan Engine protocol

When reading through a log file to troubleshoot major issues, you may find it useful look for ERROR- and WARN-level messages initially.

Thread identifies the process that generated the message.

Configuring which log severity levels are displayed

By default, all log files display messages with severity levels of INFO and higher. This means that they display INFO, WARN, ERROR messages and do not display DEBUG messages. You can change which severity levels are displayed in the log files. For example, you might want to filter out all messages except for those with WARN and ERROR severity levels. Or, you may want to include DEBUG messages for maintenance and debugging purposes.

Configuration steps are identical for the Security Console and distributed Scan Engines. To configure which log severity levels are displayed, take the following steps:

NOTE: In the user-log-settings.xml file, *default* refers to the nsc.log file or nse.log file, depending on whether the installed component is the Security Console or a distributed Scan Engine.

1. In a text editor, open the user-log-settings.xml file, which is located in the [installation_directory]/nsc/conf directory.
2. Un-comment the following line by removing the opening and closing comment tags: `<!--` and `-->`:
   `<!-- <property name="default-level" value="INFO"/> -->`
3. If you want to change the logging level for the nsc.log (for Security Console installations) or nse.log file (for Scan Engine installations), leave the value *default* unchanged. Otherwise, change the value to one of the following to specify a different log file:
   - *auth*
   - *access*
   - *mem*
4. Change the value in the line to your preferred severity level: DEBUG, INFO, WARN, or ERROR.
   Example: `<property name="default-level" value="DEBUG"/>`
5. To change log levels for additional log files, simply copy and paste the un-commented line, changing the values accordingly.
   Examples:
   ```
   <property name="default-level" value="DEBUG"/>
   <property name="auth-level" value="DEBUG"/>
   <property name="access-level" value="DEBUG"/>
   <property name="mem-level" value="DEBUG"/>
   ```
6. Save and close the file.

The change is applied after approximately 30 seconds.

Addressing a failure during startup

If a subsystem critical error occurs during startup, then the application will attempt to queue an appropriate maintenance task to respond to that failure. Afterward, it restarts in maintenance mode.

If you are an administrator, you can log on and examine the cause of failure. If required, you can take certain steps to troubleshoot the issue.

Two types of recovery tasks are available:

- *DBConfig* task is triggered when the application is unable to connect to the configured database. It allows you to test the database configuration settings and save it upon success.
- *Recovery* task is a general recovery task that is triggered when an unknown failure occurs during startup. This is very rare and happens only when one or more of the configuration files is not found or is invalid. This task allows you to view the cause of the failure and upload support logs to a secure log server, where they can be used for troubleshooting.

The application may fail to restart in maintenance mode in case of extremely critical failures if the maintenance Web server does not have the default port 3780 available. This may happen if there is already an instance of it running, or if one or more of the key configuration files is invalid or missing. These files have extensions such as *.nsc*, *.xml*, and *.userdb*.

Addressing failure to refresh a session

When the Web interface session times out in an idle session, the Security Console displays the logon window so that the user can refresh the session. If a communication issue between the Web browser and the Security Console Web server prevents the session from refreshing, a user will see an error message. If the user has unsaved work, he or she should not leave the page or close the browser because the work may not be lost after the communication issue is resolved.

A communication failure may occur for one of the following reasons. If any of these is the cause, take the appropriate action:

- The Security Console is offline. Restart the Security Console.
- The Security Console has been disconnected from the Internet. Reconnect the Security Console to the Internet.
- The user's browser has been disconnected from the Internet. Reconnect the browser to the Internet.
- The Security Console address has changed. Clear the address resolution protocol (ARP) table on the computer hosting the browser.

An extreme delay in the Security Console's response to the user's request to refresh the session also may cause the failure message to appear.

Resetting account lockout

When a user attempts to log on too many times with an incorrect password, the application locks out the user until the lockout is reset for that user.

The default lockout threshold is 4 attempts. A global administrator can change this parameter on the *Security Console Configuration—Web Server* page. See *Changing the Security Console Web server default settings* on page *43*.

You can reset the lockout using one of the following three methods:

- If you're a global administrator, go to the *Users* page, and click the padlock icon that appears next to the locked out user's name.
- Run the console command **`unlock account`**. See *Using the command console* on page *72*.
- Restart the Security Console. This is the only method that will work if the locked out user is the only global administrator in your organization.

Long or hanging scans

Occasionally, a scan will take an unusually long time, or appear to have completely stopped.

It is not possible to predict exactly how long a scan should take. Scan times vary depending on factors such as the number of target assets and the thoroughness or complexity of the scan template. However, you can observe whether a scan is taking an exceptionally long time to complete by comparing the scan time to that of previous scans.

In general, if a scan runs longer than eight hours on a single host, or 48 hours on a given site, it is advisable to check for certain problems.

Tip for addressing delayed scan operations

If you attempt to start, pause, resume, or stop a scan, and a message appears for a long time indicating that the operation is in progress, this may be due to a network-related delay in the Security Console's communication with the Scan Engine. In networks with low bandwidth or high latency, delayed scan operations may result in frequent time-outs in Security Console/Scan Engine communication, which may cause lags in the Security Console receiving scan status information. To reduce time-outs, you can increase the Scan Engine response timeout setting. See *Configuring Security Console connections with distributed Scan Engines* on page *50*.

Scan memory issues

Scans can be slow, or can fail, due to memory issues. See *Out-of-memory issues* on page *81*.

Scan complexity

For every target host that it discovers, the application scans its ports before running any vulnerability checks. The range of target ports is a configurable scan template setting. Scan times increase in proportion to the number of ports scanned.

In particular, scans of UDP ports can be slow, since the application, by default, sends no more than two UDP packets per second in order to avoid triggering the ICMP rate-limiting mechanisms that are built into TCP/IP stacks for most network devices.

To increase scan speed, consider configuring the scan to only examine well-known ports, or specific ports that are known to host relevant services. See the chapter on scan templates in the user's guide.

Scan Engine offline

If the Scan Engine goes off line during the scan, the scan will appear to hang. When a Scan Engine goes off line during the scan, the database will need to remove data from the incomplete scan. This process leaves messages similar to the following the scan log:

```
DBConsistenc3/10/09 12:05 PM: Inconsistency discovered for dispatched scan
ID 410, removing partially imported scan results...
```

If a Scan Engine goes offline, restart it. Then, go the Scan Engine Configuration panel to confirm that the Scan Engine is active. See the topic *Configuring Scan Engines* in the user's guide.

Viewing the scan log

You can view an activity log for a scan that is in progress or complete.

To view the scan log:

1. Click **View scan log**.
   The console displays the scan log.
2. Click your browser's **Back** button to return to the *Scan Progress* page.

Scan stopped by a user

If another user stops a scan, the scan will appear to have hung. To determine if this is the case, examine the log for a message similar to the following:

```
Symantec3/16/09 7:22 PM: Scan [] stopped: "maylor" <>
```

See *Viewing the scan log* on page *80*.

Long or hanging reports

Occasionally, report generation will take an unusually long time, or appear to have completely stopped. You can find reporting errors in the Security Console logs.

Reporting memory issues

Report generation can be slow, or can fail, due to memory issues. See *Out-of-memory issues* on page *81*.

Stale scan data

Database speed affects reporting speed. Over time, data from old scans will accumulate in the database. This causes the database to slow down.

If you find that reporting has become slow, look in the Security Console logs for reporting tasks whose durations are inconsistent with other reporting tasks, as in the following example:

```
nsc.log.0:Reportmanage1/5/09 3:00 AM: Report task serviceVulnStatistics
finished in 2 hours 1 minute 23 seconds
```

You can often increase report generation speed by cleaning up the database. Regular database maintenance removes leftover scan data and host information. See *Viewing the scan log* on page *80* and *Performing a backup and restore* on page *53*.

Out-of-memory issues

Scanning and reporting are memory-intensive tasks, so errors related to these activities may often be memory issues. You can control memory use by changing settings. Some memory issues are related how system resources are controlled.

Memory Protection

By default, the process auto-stop feature pauses in-progress scans and reports, or prevents them from starting, to protect the server from crashing if it is running low on free memory. If scans or reports are hanging, check to see if this feature is enabled. You can disable it; however, be aware that low memory can cause the server to fail. It is preferable to determine and correct the source of the low-memory issue.

To access the auto-stop feature:

1. Go to the *Administration* page.
2. Click **Manage** next to *Security Console*.
3. Click **General**.

   The Security Console displays the *Security Console Configuration* panel.
4. Select or clear the check box labeled **Enable process auto-stop**.

java.lang.OutofMemoryError

If the process auto-stop feature is not enabled, excessive Scan Engine or Security Console activity can cause the application to run out of memory and crash. If the application has crashed, you can verify that the crash was due to lack of memory by checking the log files for the following message:

```
java.lang.OutOfMemoryError: Java heap space
```

If you see this message, contact Technical Support. Do not restart the application unless directed to do so.

Fixing memory problems

Since scanning is memory-intensive and occurs frequently, it is important to control how much memory scans use so that memory issues do not, in turn, affect scan performance. There are a number of strategies for ensuring that memory limits do not affect scans.

Reduce scan complexity

As the number of target hosts increases, so does the amount of memory needed to store scan information. If the hosts being scanned have an excessive number of vulnerabilities, scans could hang due to memory shortages.

To reduce the complexity of a given scan, try a couple of approaches:

- Reduce the number of target hosts by excluding IP addresses in your site configuration.
- Reduce the number of target vulnerabilities by excluding lower-priority checks from your scan template.

After patching any vulnerabilities uncovered by one scan, add the excluded IP addresses or vulnerabilities to the site configuration, and run the scan again.

For more information see the topic *Configuring scan credentials for sites* and the chapter *Working with scan templates and tuning scan performance* in the user's guide.

Reduce Scan Count

Running several simultaneous scans can cause the Security Console to run out of memory. Reduce the number of simultaneous scans to conserve memory.

Upgrade Hosts

If scans are consistently running out of memory, consider adding more memory to the servers.

More information on managing scan-related resources

See the following chapters for more detailed information on making scans more memory-friendly:

- *Planning a deployment* on page *16*
- The chapter *Working with scan templates and tuning scan performance* in the user's guide

Update failures

Occasionally, system updates will be unsuccessful. You can find out why by examining the system logs.

Corrupt update table

The application keeps track of previously-applied updates in an update table. If the update table becomes corrupt, the application will not know which updates need to be downloaded and applied.

If it cannot install updates due to a corrupt update table, the Scan Console log will contain messages similar to the following:

```
AutoUpdateJo3/12/09 5:17 AM: NSC update failed: com.rapid7.updater.Update-
Exception: java.io.EOFException
at com.rapid7.updater.UpdatePackageProcessor.getUpdateTable(Unknown Source)
at com.rapid7.updater.UpdatePackageProcessor.getUpdates(Unknown Source)
at com.rapid7.updater.UpdatePackageProcessor.getUpdates(Unknown Source)
at com.rapid7.nexpose.nsc.U.execute(Unknown Source)
at com.rapid7.scheduler.Scheduler$_A.run(Unknown Source)
```

If this occurs, contact Technical Support. See *Viewing the scan log* on page *80*.

Interrupted update

By default, the application automatically downloads and installs updates. The application may download an update, but its installation attempt may be unsuccessful.

You can find out if this happened by looking at the scan log.

Check for update time stamps that demonstrate long periods of inactivity.

```
AU-BE37EE72A11/3/08 5:56 PM: updating file: nsc/htroot/help/html/757.htm
```

```
NSC 11/3/08 9:57 PM: Logging initialized (system time zone is SystemV/
PST8PDT)
```

You can use the **`update now`** command prompt to re-attempt the update manually:

1. Click the **Administration** tab to go to the *Administration* page.
2. Go to the *diag_console* page, where you can run command prompts.
3. Replace *index.html* with *diag_console.html*: https://[console_url]:3780/admin/global/diag_console.html in the navigation bar,
4. Enter the command **`update now`** in the text box and click **Execute**.

   The Security Console displays a message to indicate whether the update attempt was successful. See *Viewing the scan log* on page *80*.

Corrupt File

If the application cannot perform an update due to a corrupt file, the Scan Console log will contain messages similar to the following:

```
AU-892F7C6793/7/09 1:19 AM: Applying update id 919518342
```

```
AU-892F7C6793/7/09 1:19 AM: error in opening zip file
```

```
AutoUpdateJo3/7/09 1:19 AM: NSC update failed: com.rapid7.updater.UpdateEx-
ception:
```

```
java.util.zip.ZipException: error in opening zip file
```

```
at com.rapid7.updater.UpdatePackageProcessor.B(Unknown Source)
```

```
at com.rapid7.updater.UpdatePackageProcessor.getUpdates(Unknown Source)
```

```
at com.rapid7.updater.UpdatePackageProcessor.getUpdates(Unknown Source)
```

```
at com.rapid7.nexpose.nsc.U.execute(Unknown Source)
```

```
at com.rapid7.scheduler.Scheduler$_A.run(Unknown Source)
```

If the update fails due to a corrupt file, it means that the update file was successfully downloaded, but was invalid. If this occurs, contact Technical Support.See *Viewing the scan log* on page *80*.

Interrupted connection to the update server

If a connection between the Security Console and the update server cannot be made, it will appear in the logs with a message similar to the following.

```
AU-A7F0FF3623/10/09 4:53 PM: downloading update: 919518342
AutoUpdateJo3/10/09 4:54 PM: NSC update failed: java.net.SocketTimeoutEx-
ception
```

The java.net.SocketTimeoutException is a sign that a connection cannot be made to the update server. If the connection has been interrupted, other updates prior to the failure will have been successful.

You can use the **`update now`** command prompt to re-attempt the update manually. See *Interrupted update* on page *83* and see *Viewing the scan log* on page *80*.

Glossary

For more detailed information on any term in this glossary, search for the term in *Help*.

Advanced Policy Engine

Advanced Policy Engine is a license-enabled scanning feature that performs checks for compliance with Federal Desktop Core Configuration (FDCC), United States Government Configuration Baseline (USGCB), and other configuration policies. For information about other tools related to compliance with FDCC configuration policies, see *What are your compliance goals?* in the *administrator's guide*.

API (application programming interface)

An API is a function that a developer can integrate with another software application by using program calls. The term *API* also refers to one of two sets of XML APIs, each with its own included operations: API v1.1 and Extended API v1.2. To learn about each API, see the API documentation, which you can download from the *Support* page of Help.

Asset

An asset is a single device on a network that the application discovers during a scan. In the Web interface and API, an asset may also be referred to as a *device*. See *Managed asset* on page *89* and *Unmanaged asset* on page *94*. An asset's data has been integrated into the scan database, so it can be listed in sites and asset groups. In this regard, it differs from a *node*. See *Node* on page *90*.

Asset group

An asset group is a logical collection of managed assets to which specific members have access for creating or viewing reports or tracking remediation tickets. An asset group may contain assets that belong to multiple sites or other asset groups. An asset group is either static or dynamic. An asset group is not a site. See *Site* on page *93*. See *Dynamic asset group* on page *88* and *Static asset group* on page *93*.

Asset Owner

Asset Owner is one of the preset roles. A user with this role can view data about discovered assets, run manual scans, and create and run reports in accessible sites and asset groups.

Asset search filter

An asset search filter is a set of criteria with which a user can refine a search for assets to include in a dynamic asset group. An asset search filter is different from a *vAsset discovery filter* on page *95*.

Authentication

Authentication is the process of a security application verifying the logon credentials of a client or user that is attempting to gain access. By default the application authenticates users with an internal process, but you can configure it to authenticate users with an external LDAP or Kerberos source.

Average risk

Average risk is a setting in risk trend report configuration. It is based on a calculation of your risk scores on assets over a report date range. For example, average risk gives you an overview of how vulnerable your assets might be to exploits whether it's high or low or unchanged. Some assets have higher risk scores than others. Calculating the average score provides a high-level view of how vulnerable your assets might be to exploits.

Benchmark

In the context of scanning for FDCC policy compliance, a benchmark is a combination of policies that share the same source data. Each policy in the Advanced Policy Engine contains some or all of the rules that are contained within its respective benchmark. See *Federal Desktop Core Configuration (FDCC)* on page *89* and *United States Government Configuration Baseline (USGCB)* on page *94*.

Breadth

Breadth refers to the total number of assets within the scope of a scan.

Category

In the context of scanning for FDCC policy compliance, a category is a grouping of policies in the Advanced Policy Engine configuration for a scan template. A policy's category is based on its source, purpose, and other criteria. See *Advanced Policy Engine* on page *85*, *Federal Desktop Core Configuration (FDCC)* on page *89*, and *United States Government Configuration Baseline (USGCB)* on page *94*.

Check type

A check type is a specific kind of check to be run during a scan. Examples: The Unsafe check type includes aggressive vulnerability testing methods that could result in Denial of Service on target assets; the Policy check type is used for verifying compliance with policies. The check type setting is used in scan template configurations to refine the scope of a scan.

Command console

The command console is a page in the Security Console Web interface for entering commands to run certain operations. When you use this tool, you can see real-time diagnostics and a behind-the-scenes view of Security Console activity. To access the command console page, click the **Run console commands** link next to the *Troubleshooting* item on the *Administration* page.

Common Configuration Enumeration (CCE)

Common Configuration Enumeration (CCE) is a standard for assigning unique identifiers known as CCEs to configuration controls to allow consistent identification of these controls in different environments. CCE is implemented as part of its compliance with SCAP criteria for an Unauthenticated Scanner product.

Common Platform Enumeration (CPE)

Common Platform Enumeration (CPE) is a method for identifying operating systems and software applications. Its naming scheme is based on the generic syntax for Uniform Resource Identifiers (URI). CCE is implemented as part of its compliance with SCAP criteria for an Unauthenticated Scanner product.

Common Vulnerabilities and Exposures (CVE)

The Common Vulnerabilities and Exposures (CVE) standard prescribes how the application should identify vulnerabilities, making it easier for security products to exchange vulnerability data. CVE is implemented as part of its compliance with SCAP criteria for an Unauthenticated Scanner product.

Common Vulnerability Scoring System (CVSS)

Common Vulnerability Scoring System (CVSS) is an open framework for calculating vulnerability risk scores. CVSS is implemented as part of its compliance with SCAP criteria for an Unauthenticated Scanner product.

Compliance

Compliance is the condition of meeting standards specified by a government or respected industry entity. The application tests assets for compliance with a number of different security standards, such as those mandated by the Payment Card Industry (PCI) and those defined by the National Institute of Standards and Technology (NIST) for Federal Desktop Core Configuration (FDCC).

Continuous scan

A continuous scan starts over from the beginning if it completes its coverage of site assets within its scheduled window. This is a site configuration setting.

Coverage

Coverage indicates the scope of vulnerability checks. A coverage improvement listed on the *News* page for a release indicates that vulnerability checks have been added or existing checks have been improved for accuracy or other criteria.

Depth

Depth indicates how thorough or comprehensive a scan will be. Depth refers to level to which the application will probe an individual asset for system information and vulnerabilities.

Discovery (scan phase)

Discovery is the first phase of a scan, in which the application finds potential scan targets on a network. Discovery as a scan phase is different from *vAsset discovery* on page *95*.

Dynamic asset group

A dynamic asset group contains scanned assets that meet a specific set of search criteria. You define these criteria with asset search filters, such as IP address range or operating systems. The list of assets in a dynamic group is subject to change with every scan or when vulnerability exceptions are created. In this regard, a dynamic asset group differs from a static asset group. See *Asset group* on page *85* and *Static asset group* on page *93*.

Dynamic Scan Pool

The Dynamic Scan Pool feature allows you to use Scan Engine pools to enhance the consistency of your scan coverage. A Scan Engine pool is a group of shared Scan Engines that can be bound to a site so that the load is distributed evenly across the shared Scan Engines. You can configure scan pools using the Extended API v1.2.

Dynamic site

A dynamic site is a collection of assets that are targeted for scanning and that have been discovered through vAsset discovery. Asset membership in a dynamic site is subject to change if the discovery connection changes or if filter criteria for asset discovery change. See *Static site* on page *94*, *Site* on page *93*, and *vAsset discovery* on page *95*.

Exploit

An exploit is an attempt to penetrate a network or gain access to a computer through a security flaw, or vulnerability. Malicious exploits can result in system disruptions or theft of data. Penetration testers use benign exploits only to verify that vulnerabilities exist. The Metasploit product is a tool for performing benign exploits. See *Metasploit* on page *90*. See *Published exploit* on page *91*.

Exposure

An exposure is a vulnerability, especially one that makes an asset susceptible to attack via malware or a known exploit.

Extensible Configuration Checklist Description Format (XCCDF)

As defined by the National Institute of Standards and Technology (NIST), Extensible Configuration Checklist Description Format (XCCDF) "is a specification language for writing security checklists, benchmarks, and related documents. An XCCDF document represents a structured collection of security configuration rules for some set of target systems. The specification is designed to support information interchange, document generation, organizational and situational tailoring, automated compliance testing, and compliance scoring." Advanced Policy Engine checks for FDCC policy compliance are written in this format.

False positive

A false positive is an instance in which the application flags a vulnerability that doesn't exist. A false negative is an instance in which the application fails to flag a vulnerability that does exist.

Federal Desktop Core Configuration (FDCC)

The Federal Desktop Core Configuration (FDCC) is a grouping of configuration security settings recommended by the National Institute of Standards and Technology (NIST) for computers that are connected directly to the network of a United States government agency. The Advanced Policy Engine provides checks for compliance with these policies in scan templates. Performing these checks requires a license that enables the Advanced Policy Engine feature and FDCC scanning.

Fingerprinting

Fingerprinting is a method of identifying the operating system of a scan target or detecting a specific version of an application.

Global Administrator

Global Administrator is one of the preset roles. A user with this role can perform all operations that are available in the application and they have access to all sites and asset groups.

Host

A host is a physical or virtual server that provides computing resources to a guest virtual machine. In a high-availability virtual environment, a host may also be referred to as a node. The term *node* has a different context in the application. See *Node* on page *90*.

Latency

Latency is the delay interval between the time when a computer sends data over a network and another computer receives it. Low latency means short delays.

Malware

Malware is software designed to disrupt or deny a target systems's operation, steal or compromise data, gain unauthorized access to resources, or perform other similar types of abuse. The application can determine if a vulnerability renders an asset susceptible to malware attacks.

Malware kit

Also known as an exploit kit, a malware kit is a software bundle that makes it easy for malicious parties to write and deploy code for attacking target systems through vulnerabilities.

Managed asset

A managed asset is a network device that has been discovered during a scan and added to a site's target list, either automatically or manually. Only managed assets can be checked for vulnerabilities and tracked over time. Once an asset becomes a managed asset, it counts against the maximum number of assets that can be scanned, according to your license.

Manual scan

A manual scan is one that you start at any time, even if it is scheduled to run automatically at other times. Synonyms include *ad-hoc scan* and *unscheduled scan*.

Metasploit

Metasploit is a product that performs benign exploits to verify vulnerabilities. See *Exploit* on page *88*.

MITRE

The MITRE Corporation is a body that defines standards for enumerating security-related concepts and languages for security development initiatives. Examples of MITRE-defined enumerations include Common Configuration Enumeration (CCE) and Common Vulnerability Enumeration (CVE). Examples of MITRE-defined languages include Open Vulnerability and Assessment Language (OVAL). A number of MITRE standards are implemented, especially in verification of FDCC compliance.

National Institute of Standards and Technology (NIST)

National Institute of Standards and Technology (NIST) is a non-regulatory federal agency within the U.S. Department of Commerce. The agency mandates and manages a number of security initiatives, including Security Content Automation Protocol (SCAP). See *Security Content Automation Protocol (SCAP)* on page *93*.

Node

A node is a device on a network that the application discovers during a scan. After the application integrates its data into the scan database, the device is regarded as an *asset* that can be listed in sites and asset groups. See *Asset* on page *85*.

Open Vulnerability and Assessment Language (OVAL)

Open Vulnerability and Assessment Language (OVAL) is a development standard for gathering and sharing security-related data, such as FDCC policy checks. In compliance with an FDCC requirement, each OVAL file that the application imports during configuration policy checks is available for download from the *SCAP* page in the Security Console Web interface.

Override

An override is a change made by a user to the result of a check for compliance with a configuration policy rule. For example, a user may override a Fail result with a Pass result.

Payment Card Industry (PCI)

The Payment Card Industry (PCI) is a council that manages and enforces the PCI Data Security Standard for all merchants who perform credit card transactions. The application includes a scan template and report templates that are used by Approved Scanning Vendors (ASVs) in official merchant audits for PCI compliance.

Permission

A permission is the ability to perform one or more specific operations. Some permissions only apply to sites or asset groups to which an assigned user has access. Others are not subject to this kind of access.

Policy

A policy is a set of primarily security-related configuration guidelines for a computer, operating system, software application, or database. Compliance is verified with a number of different policies, including those encompassed in the United States Government Configuration Baseline (USGCB) and the Federal Desktop Core Configuration (FDCC). See *Advanced Policy Engine* on page *85*, *Federal Desktop Core Configuration (FDCC)* on page *89*, *United States Government Configuration Baseline (USGCB)* on page *94*, and *Scan* on page *92*.

Policy Result

In the context of FDCC policy scanning, a result is a state of compliance or non-compliance with a rule or policy. Possible results include *Pass*, *Fail*, or *Not Applicable*.

Policy Rule

A rule is one of a set of specific guidelines that make up an FDCC configuration policy. See *Federal Desktop Core Configuration (FDCC)* on page *89*, *United States Government Configuration Baseline (USGCB)* on page *94*, and *Policy* on page *91*.

Published exploit

In the context of the application, a published exploit is one that has been developed in Metasploit or listed in the Exploit Database. See *Exploit* on page *88*.

Real Risk strategy

Real Risk is one of the built-in strategies for assessing and analyzing risk. It is also the recommended strategy because it applies unique exploit and malware exposure metrics for each vulnerability to Common Vulnerability Scoring System (CVSS) base metrics for likelihood (access vector, access complexity, and authentication requirements) and impact to affected assets (confidentiality, integrity, and availability). See *Risk strategy* on page *92*.

Risk

In the context of vulnerability assessment, risk reflects the likelihood that a network or computer environment will be compromised, and it characterizes the anticipated consequences of the compromise, including theft or corruption of data and disruption to service. Implicitly, risk also reflects the potential damage to a compromised entity's financial well-being and reputation.

Risk score

A risk score is a rating that the application calculates for every asset and vulnerability. The score indicates the potential danger posed to network and business security in the event of a malicious exploit. You can configure the application to rate risk according to one of several built-in risk strategies, or you can create custom risk strategies.

Risk strategy

A risk strategy is a method for calculating vulnerability risk scores. Each strategy emphasizes certain risk factors and perspectives. Four built-in strategies are available: *Real Risk strategy* on page *91*, *TemporalPlus risk strategy* on page *94*, *Temporal risk strategy* on page *94*, and *Weighted risk strategy* on page *96*. You can also create custom risk strategies.

Risk trend

A risk trend graph illustrates a long-term view of your assets' probability and potential impact of compromise that may change over time. Risk trends can be based on average or total risk scores. The highest-risk graphs in your report demonstrate the biggest contributors to your risk on the site, group, or asset level. Tracking risk trends helps you assess threats to your organization's standings in these areas and determine if your vulnerability management efforts are satisfactorily maintaining risk at acceptable levels or reducing risk over time. See *Average risk* on page *86* and *Total risk* on page *94*.

Role

A role is a set of permissions. Five preset roles are available. You also can create custom roles by manually selecting permissions. See *Asset Owner* on page *85*, *Security Manager* on page *93*, *Global Administrator* on page *89*, *Site Owner* on page *93*, and *User* on page *95*.

Scan

A scan is a process by which the application discovers network assets and checks them for vulnerabilities. See *Exploit* on page *88* and See *Vulnerability check* on page *95*.

Scan credentials

Scan credentials are the user name and password that the application submits to target assets for authentication to gain access and perform deep checks. Many different authentication mechanisms are supported for a wide variety of platforms. See *Shared scan credentials* on page *93* and *Site-specific scan credentials* on page *93*.

Scan Engine

The Scan Engine is one of two major application components. It performs asset discovery and vulnerability detection operations. Scan engines can be *distributed* within or outside a firewall for varied coverage. Each installation of the Security Console also includes a local engine, which can be used for scans within the console's network perimeter.

Scan template

A scan template is a set of parameters for defining how assets are scanned. Various preset scan templates are available for different scanning scenarios. You also can create custom scan templates. Parameters of scan templates include the following:

- methods for discovering assets and services
- types of vulnerability checks, including safe and unsafe
- Web application scanning properties
- verification of compliance with policies and standards for various platforms

Scheduled scan

A scheduled scan starts automatically at predetermined points in time. The scheduling of a scan is an optional setting in site configuration. It is also possible to start any scan manually at any time.

Security Console

The Security Console is one of two major application components. It controls Scan Engines and retrieves scan data from them. It also controls all operations and provides a Web-based user interface.

Security Content Automation Protocol (SCAP)

Security Content Automation Protocol (SCAP) is a collection of standards for expressing and manipulating security data. It is mandated by the U.S. government and maintained by the National Institute of Standards and Technology (NIST). The application complies with SCAP criteria for an Unauthenticated Scanner product.

Security Manager

Security Manager is one of the preset roles. A user with this role can configure and run scans, create reports, and view asset data in accessible sites and asset groups.

Shared scan credentials

One of two types of credentials that can be used for authenticating scans, shared credentials are created by Global Administrators and can be applied to multiple assets in any number of sites. See *Site-specific scan credentials* on page *93*.

Site

A site is a collection of assets that are targeted for a scan. Each site is associated with a list of target assets, a scan template, one or more Scan Engines, and other scan-related settings. See *Dynamic site* on page *88* and *Static site* on page *94*. A site is not an asset group. See *Asset group* on page *85*.

Site-specific scan credentials

One of two types of credentials that can be used for authenticating scans, a set of single-instance credentials is created for an individual site configuration and can only be used in that site. See *Scan credentials* on page *92* and *Shared scan credentials* on page *93*.

Site Owner

Site Owner is one of the preset roles. A user with this role can configure and run scans, create reports, and view asset data in accessible sites.

Static asset group

A static asset group contains assets that meet a set of criteria that you define according to your organization's needs. Unlike with a dynamic asset group, the list of assets in a static group does not change unless you alter it manually. See *Dynamic asset group* on page *88*.

Static site

A static site is a collection of assets that are targeted for scanning and that have been manually selected. Asset membership in a static site does not change unless a user changes the asset list in the site configuration. For more information, see *Dynamic site* on page *88* and *Site* on page *93*.

Temporal risk strategy

One of the built-in risk strategies, Temporal indicates how time continuously increases likelihood of compromise. The calculation applies the age of each vulnerability, based on its date of public disclosure, as a multiplier of CVSS base metrics for likelihood (access vector, access complexity, and authentication requirements) and asset impact (confidentiality, integrity, and availability). Temporal risk scores will be lower than TemporalPlus scores because Temporal limits the risk contribution of partial impact vectors. See *Risk strategy* on page *92*.

TemporalPlus risk strategy

One of the built-in risk strategies, TemporalPlus provides a more granular analysis of vulnerability impact, while indicating how time continuously increases likelihood of compromise. It applies a vulnerability's age as a multiplier of CVSS base metrics for likelihood (access vector, access complexity, and authentication requirements) and asset impact (confidentiality, integrity, and availability). TemporalPlus risk scores will be higher than Temporal scores because TemporalPlus expands the risk contribution of partial impact vectors. See *Risk strategy* on page *92*.

Total risk

Total risk is a setting in risk trend report configuration. It is an aggregated score of vulnerabilities on assets over a specified period.

United States Government Configuration Baseline (USGCB)

The United States Government Configuration Baseline (USGCB) is an initiative to create security configuration baselines for information technology products deployed across U.S. government agencies. USGCB evolved from FDCC, which it replaces as the configuration security mandate in the U.S. government. The Advanced Policy Engine provides checks for Microsoft Windows 7, Windows 7 Firewall, and Internet Explorer for compliance with USGCB baselines. Performing these checks requires a license that enables the Advanced Policy Engine feature and USGCB scanning. See *Advanced Policy Engine* on page *85* and *Federal Desktop Core Configuration (FDCC)* on page *89*.

Unmanaged asset

An unmanaged asset is a device that has been discovered during a scan but not correlated against a managed asset or added to a site's target list. The application is designed to provide sufficient information about unmanaged assets so that you can decide whether to manage them. An unmanaged assets does not count against the maximum number of assets that can be scanned according to your license.

Unsafe check

An unsafe check is a test for a vulnerability that can cause a denial of service on a target system. Be aware that the check itself can cause a denial of service, as well. It is recommended that you only perform unsafe checks on test systems that are not in production.

Update

An update is a released set of changes to the application. By default, two types of updates are automatically downloaded and applied:

- *Content* updates include new checks for vulnerabilities, patch verification, and security policy compliance. Content updates always occur automatically when they are available.
- *Product* updates include performance improvements, bug fixes, and new product features. Unlike content updates, it is possible to disable automatic product updates and update the product manually.

User

User is one of the preset roles. An individual with this role can view asset data and run reports in accessible sites and asset groups.

vAsset discovery

vAsset discovery is a process by which the application automatically discovers virtual assets through a connection with a vSphere server or virtual machine host. You can refine or limit asset discovery with criteria filters. See *vAsset discovery filter* on page *95* and *vConnection* on page *95*. vAsset discovery is different from *Discovery (scan phase)* on page *87*.

vAsset discovery filter

A vAsset discovery filter is a set of criteria refining or limiting vAsset discovery results. This type of filter is different from an *Asset search filter* on page *85*.

vConnection

A vConnection is a connection that is initiated with a server that manages virtual machines in order to discover those assets. A Global Administrator can configure a vConnection. See *vAsset discovery filter* on page *95*.

Vulnerability

A vulnerability is a security flaw in a network or computer.

Vulnerability category

A vulnerability category is a set of vulnerability checks with shared criteria. For example, the Adobe category includes checks for vulnerabilities that affect Adobe applications. There are also categories for specific Adobe products, such as Air, Flash, and Acrobat/Reader. Vulnerability check categories are used to refine scope in scan templates. Vulnerability check results can also be filtered according category for refining the scope of reports.

Vulnerability check

A vulnerability check is a series of operations that are performed to determine whether a security flaw exists on a target asset.

Vulnerability exception

A vulnerability exception is the removal of a vulnerability from a report and from any asset listing table. Excluded vulnerabilities also are not considered in the computation of risk scores.

Weighted risk strategy

One of the built-in risk strategies, Weighted is based primarily on asset data and vulnerability types, and it takes into account the level of importance, or weight, that you assign to a site when you configure it. See *Risk strategy* on page *92*.

Appendix A: SCAP compliance

Symantec CCS Vulnerability Manager complies with Security Content Automation Protocol (SCAP) criteria for an Unauthenticated Scanner product. SCAP is a collection of standards for expressing and manipulating security data in standardized ways. It is mandated by the U.S. government and maintained by the National Institute of Standards and Technology (NIST).

This appendix provides information about how the SCAP standards are implemented for an Unauthenticated Scanner:

- The **Common Platform Enumeration (CPE)** naming scheme, based on the generic syntax for Uniform Resource Identifiers (URI), is a method for identifying operating systems and software applications.
- The **Common Vulnerabilities and Exposures (CVE)** standard prescribes how the product should identify vulnerabilities, making it easier for security products to exchange vulnerability data.
- The **Common Vulnerability Scoring System (CVSS)** is an open frame work for calculating vulnerability risk scores.
- **Common Configuration Enumeration (CCE)** is a standard for assigning unique identifiers known as CCEs to configuration controls to allow consistent identification of these controls in different environments.

How CPE is implemented

During scans, Symantec CCS Vulnerability Manager utilizes its fingerprinting technology to recognize target platforms and applications. After completing scans and populating its scan database with newly acquired data, it applies CPE names to fingerprinted platforms and applications whenever corresponding CPE names are available.

Within the database, CPE names are continually kept up to date with changes to the National Institute of Standards (NIST) CPE dictionary. With every revision to the dictionary, the application maps newly available CPE names to application descriptions that previously did not have CPE names.

The Security Console Web interface displays CPE names in scan data tables. You can view these names in listings of assets, software, and operating systems, as well as on pages for specific assets. CPE names also appear in reports in the XML Export format.

How CVE is implemented

When Symantec CCS Vulnerability Manager populates its scan database with discovered vulnerabilities, it applies Common Vulnerabilities and Exposures (CVE) identifiers to these vulnerabilities whenever these identifiers are available.

You can view CVE identifiers on vulnerability detail pages in the Security Console Web interface. Each listed identifier is a hypertext link to the CVE online database at nvd.nist.gov, where you can find additional relevant information and links.

You can search for vulnerabilities in the application interface by using CVE identifiers as search criteria.

CVE identifiers also appear in the Discovered Vulnerabilities sections of reports.

The application uses the most up-to-date CVE listing from the CVE mailing list and changelog. Since the application always uses the most up-to-date CVE listing, it does not have to list CVE version numbers. The application updates its vulnerability definitions every six hours through a subscription service that maintains existing definitions and links and adds new ones continuously.

How CVSS is implemented

For every vulnerability that it discovers, Symantec CCS Vulnerability Manager computes a Common Vulnerability Scoring System (CVSS) Version 2 score. In the Security Console Web interface, each vulnerability is listed with its CVSS score. You can use this score, severity rankings, and risk scores based on either temporal or weighted scoring models—depending on your configuration preference—to prioritize vulnerability remediation tasks.

The application incorporates the CVSS score in the PCI Executive Summary and PCI Vulnerability Details reports, which provide detailed Payment Card Industry (PCI) compliance results. Each discovered vulnerability is ranked according to its CVSS score. Symantec is an Approved Scanning Vendor (ASV); and Symantec CCS Vulnerability Manager is a Payment Card Industry (PCI)-sanctioned tool for conducting compliance audits. CVSS scores correspond to severity rankings, which ASVs use to determine which determine whether a given asset is compliant with PCI standards.

The application also includes the CVSS score in report sections that appear in various report templates. The *Highest Risk Vulnerability Details* section lists highest risk vulnerabilities and includes their categories, risk scores, and their CVSS scores. The *Index of Vulnerabilities* section includes the severity level and CVSS rating for each vulnerability.

The *PCI Vulnerability Details* section contains in-depth information about each vulnerability included in a PCI Audit (legacy) report. It quantifies the vulnerability according to its severity level and its CVSS rating.

How CCE is implemented

Symantec CCS Vulnerability Manager tests assets for compliance with configuration policies. It displays the results of compliance tests on the scan results page of every tested asset. The *Advanced Policy Listing* table on this page displays every policy against which the asset was tested.

Every listed policy is a hyperlink to a page about that policy, which includes a table of its constituent rules. Each listed rule is a hyperlink to a page about that rule. The rule page includes detailed technical information about the rule and lists its CCE identifier.

CCE entries can be found via the search feature. See *Using the Search feature* in the *user's guide*.

Where to find SCAP update information and OVAL files

Symantec CCS Vulnerability Manager automatically includes any new SCAP content with each content update. For a definition of content updates, see *Manage updates* on page *44*. You can view SCAP update information on the *SCAP* page, which you can access from the *Administration* page in Security Console Web interface.

Four tables appear on the SCAP page:

- CPE Data
- CVE Data
- CVSS Data
- CCE Data

Each table lists the most recent content update that included new SCAP data and the most recent date that NIST generated new data.

On the SCAP page you also can view a list of Open Vulnerability and Assessment Language (OVAL) files that it has imported during configuration policy checks. In compliance with an FDCC requirement, each listed file name is a hyperlink that you can click to download the XML-structured check content.

Index