

Symantec™ Critical System Protection Planning and Deployment Guide

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

Documentation version: 5.2.9

Legal Notice

Copyright © 2012 Symantec Corporation. All rights reserved.

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

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

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

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

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

Symantec Corporation
350 Ellis Street
Mountain View, CA 94043

<http://www.symantec.com>

Printed in the United States of America.

10 9 8 7 6 5 4 3 2 1

Technical Support

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

Symantec's support offerings include the following:

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

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

www.symantec.com/business/support/

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

Contacting Technical Support

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

www.symantec.com/business/support/

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

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

- Product release level

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

Licensing and registration

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

www.symantec.com/business/support/

Customer service

Customer service information is available at the following URL:

www.symantec.com/business/support/

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

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

Support agreement resources

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

Asia-Pacific and Japan	customercare_apac@symantec.com
------------------------	--

Europe, Middle-East, and Africa	semea@symantec.com
---------------------------------	--

North America and Latin America	supportsolutions@symantec.com
---------------------------------	--

Symantec™ Critical System Protection Planning and Deployment Guide

This document includes the following topics:

- [About Symantec Critical System Protection](#)
- [About the detection component and the prevention component](#)
- [About the Symantec Critical System Protection infrastructure](#)
- [Setting up a demonstration environment](#)
- [Planning an enterprise deployment](#)

About Symantec Critical System Protection

Symantec Critical System Protection provides a policy-based approach to endpoint security and compliance. Its intrusion prevention and detection features operate across a broad range of platforms and applications. It provides:

- A policy-based host security agent for monitoring and protection.
- Proactive attack prevention using the least privilege containment approach.
- A centralized management environment for enterprise systems that contain Windows, UNIX, and Linux computers.

Table 1-1 Symantec Critical System Protection capabilities

Security and protection	Compliance
<ul style="list-style-type: none">■ Real-time proactive enforcement■ Intrusion and malware prevention■ System hardening■ Application control■ Privileged user access control■ Vulnerability and patch mitigation■ Does not use signatures or require continual updates to content	<ul style="list-style-type: none">■ Real-time monitoring and auditing■ Host intrusion detection■ File integrity monitoring■ Configuration monitoring■ Tracking and monitoring of user access■ Logging and event reporting

The major features of Symantec Critical System Protection are as follows:

- Intrusion detection facility for compliance auditing
 - Real-time file integrity monitoring
 - Granular change detection of registry values, file contents, and attributes
 - Operating system and application log monitoring
 - Local event correlation and smart response actions
- Intrusion Prevention facility for malware prevention and system lockdown
 - Sandbox containment of operating system and application processes by an in-kernel reference monitor
 - Granular access control of network, file systems, registry, process-to-process memory access, system calls, and application and child process launches
 - Privileged user and program behavior
- Comprehensive out-of-the-box policies for complete system monitoring and protection of physical and virtual systems
- Centralized management environment for administering agents, policies, and events
- Integration with Security Information and Event Management (SIEM) and other security tools, as well as enterprise infrastructure components such as Active Directory, SMTP, and SNMP
- Broad platform support across Windows, Linux, UNIX and virtual environments for critical servers, workstations, laptops, and standalone systems

- Out-of-the-box policies that monitor and protect VMware vSphere components, including ESX/ESXi hypervisors, guest virtual machines and the vCenter Server.

The major benefits of Symantec Critical System Protection are as follows:

- Reduces emergency patching and minimizes patch-related downtime and IT expenses through proactive protection that does not require continuous updates.
- Reduces incidents and remediation costs with continuous security. Once the agent has a policy, it enforces the policy even when the computer is not connected to the corporate network. And even if a computer is unable to obtain the latest patches in a timely fashion, Symantec Critical System Protection continues to block attacks so that the computer is always protected.
- Provides visibility and control over the security posture of business-critical enterprise assets.
- Uses predefined compliance and hardening policies to provide efficient security management, reporting, alerting, and auditing of activities. Also provides compensating controls for compliance failures.

About the detection component and the prevention component

Symantec Critical System Protection agents have two enforcement components that you can independently activate on critical systems. The prevention component provides in-line prevention of potential security or compliance threats to the system before such access occurs. The detection component detects system configuration changes and monitors application and system logs for events of interest. Both components provide granular control over logging using policy settings and other agent configuration settings. Event logging provides visibility into actionable events as well as the efficient management of high volume events necessary for regulatory or forensic purposes.

Unlike the prevention component, which has proactive enforcement rules that can block activity before it occurs, the detection component monitors for system activity that has already occurred. However, the detection component can be configured to perform local response actions, such as writing a log, killing a process, or terminating a network session. It can also run programs or scripts when a specific condition has been triggered. Thus, in combination, the two components provide unique capabilities to both secure a system and to address regulatory compliance requirements.

For example, regulations such as PCI-DSS require that companies deploy file integrity monitoring for critical system and application files changes. The detection component of Symantec Critical System Protection can help to meet this requirement. If your objective is to deny such changes unless they occur by trusted mechanisms, you want to use the prevention component. For example, detecting that an important operating system binary like svchost.exe was recently modified is very different from preventing the modification in the first place. Symantec Critical System Protection lets you configure and use both detection capabilities and prevention capabilities as needed to address your auditing, compliance, and security requirements.

The detection component also contains operating system-specific policies that provide comprehensive operating system event monitoring and logging capabilities.

Symantec Critical System Protection policies provide thousands of pre-built rules that comprehensively monitor and harden the operating system of enterprise systems and require minimal tuning.

About the Symantec Critical System Protection infrastructure

The management components of Symantec Critical System Protection can be installed on one system or in a distributed model. Agents are generally deployed to every supported host to be monitored and protected, including the management server, management console, and SQL server database. Agentless monitoring of remote systems can extend file integrity monitoring and log monitoring functionality to systems where no native agent exists. For example, such systems include mainframe zLinux, AS 400, VAX, or VMS systems. The Symantec Critical System Protection management console is available in both a Web browser or as a standalone thick client.

[Figure 1-1](#) illustrates a high level architectural view of Symantec Critical System Protection.

Figure 1-1 Symantec Critical System Protection architecture

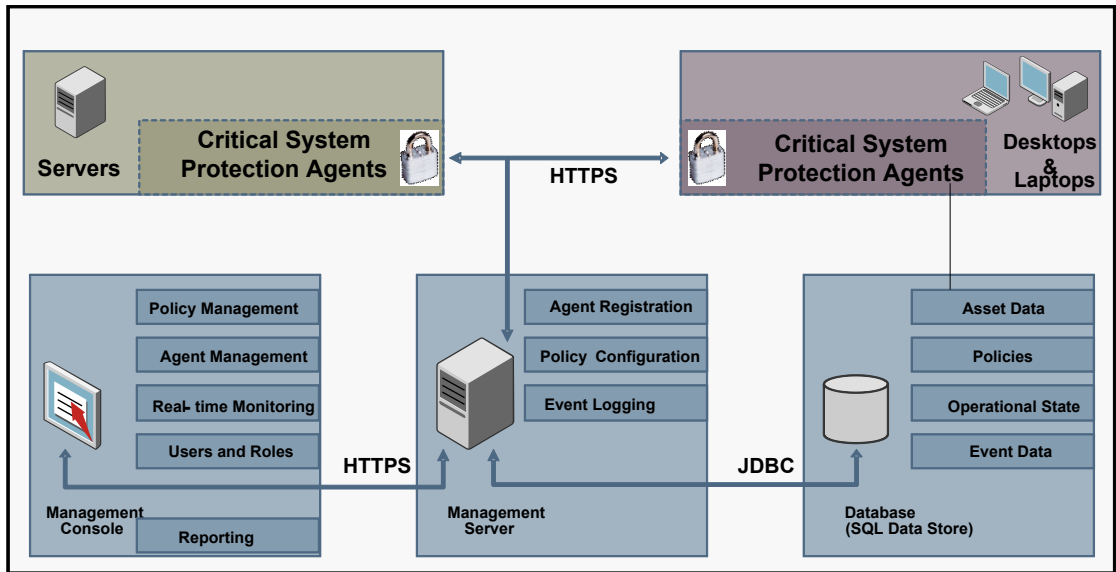


Table 1-2 Key components of Symantec Critical System Protection

Component	Description
Symantec Critical System Protection agent for behavior control	<p>The Symantec Critical System Protection agent for behavior control provides the following capabilities:</p> <ul style="list-style-type: none"> ■ Intercepts the system calls to enforce prevention policies ■ Contains multiple detection sensors for monitoring system change events and log files ■ Contains the tools for configuration and diagnostic support ■ Downloads the policies and settings from the management server and uploads events and status information to the management server ■ Natively supports Windows servers and workstations, Red Hat, SUSE, and CentOS Linux, and AIX, Solaris, HP-UX, and Tru64 UNIX ■ Supported on VMware guest systems with any of the operating systems that are natively supported. ■ Can be used to remotely monitor another host without a native agent, but note that only detection features are available in agentless mode <p>See the latest Symantec Critical System Protection Platform/Feature Matrix for up-to-date information about supported operating systems and virtual platforms.</p>

Table 1-2 Key components of Symantec Critical System Protection (*continued*)

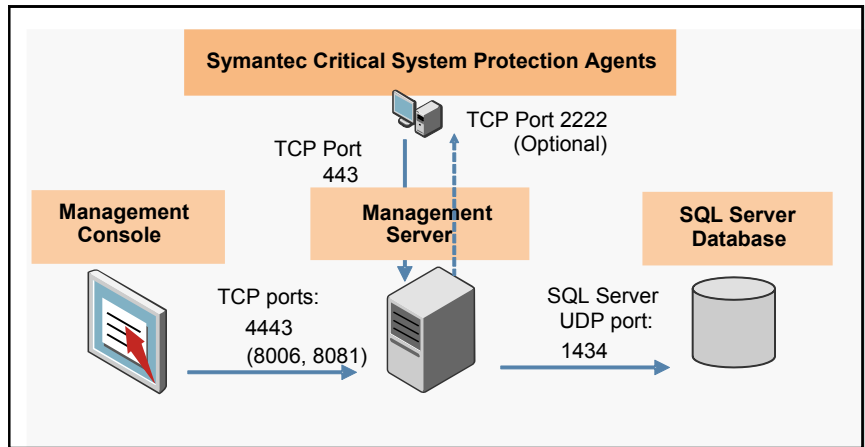
Component	Description
<p>Database</p> <p>An MS SQL Server 2005 or later</p>	<p>The database provides the following capabilities:</p> <ul style="list-style-type: none"> ■ Accessible thru JDBC/ODBC ■ Stores the policies, agent information, and real-time actionable events ■ Lets you configure encrypted communications between the database and the management server
<p>Management server</p> <p>A J2EE application server that supports high availability and scalability</p>	<p>The management server is based on Tomcat Application Server software.</p> <p>The management server provides the following capabilities:</p> <ul style="list-style-type: none"> ■ Secure communications with agent and console ■ Bulk event file storage management for efficient archival storage of all logged events ■ Alert processing (SMTP, SNMP, file), data purging, and other management functions
<p>Management console</p> <p>A rich client user interface that is written in Java Swing and is also available in a Web browser version</p>	<p>The management console provides the following capabilities:</p> <ul style="list-style-type: none"> ■ Policy, agent, and event management ■ Real-time event monitoring from the dashboard ■ Flexible hierarchy and agent grouping support ■ Event Wizard for quick policy adjustment ■ Querying, reporting, alerting ■ User and role management ■ Auditing console actions and server events
<p>Predefined Detection and Prevention policies</p>	<p>The predefined Detection and Prevention policies provide the following capabilities:</p> <ul style="list-style-type: none"> ■ Best practice policy content for operating system protection of Windows, Linux, UNIX, and vSphere ■ Common use case templates for creating customer-specific rules ■ Easy policy configuration interface ■ Flexible administration of the policies that are applied to agents

Key points to remember about the ports and the communication flow in Symantec Critical System Protection are as follows:

- Symantec Critical System Protection requires very few ports.
- All ports are configurable.
- Agents can communicate readily within a network address translation environment. A network address translation environment initiates connections

to the manager to transmit events and download policy updates or configuration updates.

Figure 1-2 Ports and communication flow



When you deploy Symantec Critical System Protection in your environment, you must ensure that the proper communications and connectivity are available for the following components:

- Server to database
- Agents to server
- Console to server

Agents continue to monitor and enforce security even if network outages occur between the agents and the server environment. In fact, you can also configure the agent to operate in a standalone or an unmanaged mode.

You can deploy Symantec Critical System Protection components on physical systems and in virtualized environments. A virtualized ecosystem such as the one supported by VMware has many parts. Its parts include management infrastructure, virtual guest machines, and hypervisors that span a variety of operating systems. To protect this heterogeneous environment, Symantec Critical System Protection relies on specific policies and enforcement agents that are appropriate to each component to be secured. The components include ESX, ESXi, and vCenter.

For information about VMware, see VMware product documentation, such as the "*vSphere Support Guide*," located on the [VMware Web site](#).

Setting up a demonstration environment

You can quickly set up a demonstration environment that contains all of the Symantec Critical System Protection components on one computer. These components include the management server, the database, the management console, and the agent. Such environments are useful early on to gain hands-on experience and a deeper understanding of the product features and capabilities. Although the built-in database is limited to 4GB and typically should not be used for more than 500 agents, you can still exercise most major capabilities of the solution. You can deploy all components in one virtual machine on a low end server with a two-processor Xeon, 4GB memory, and 10GB free disk space. You can also use workstation virtual machines to deploy a small number of agents.

Note: High availability and scalability tests require dedicated systems for Symantec Critical System Protection, and should not be conducted in the demonstration environment.

Table 1-3 Quick installation overview for an evaluation environment

Step	Task
1	Deploy the Symantec Critical System Protection server, database, and console. Installation of the management environment should take approximately 30 minutes or less.
2	Deploy the Symantec Critical System Protection agents to the computers that need protection. Installation of each agent should take approximately five minutes or less.

For installation details, see the *Symantec Critical System Protection Installation Guide*.

Symantec Critical System Protection provides host-based agents with broad platform coverage across Windows, Linux, Solaris, AIX, and HP-UX. It protects critical enterprise systems with policy-based endpoint security and compliance. Although you can quickly set up a demonstration environment, it is important to remember that up-front planning and continued operational rigor regarding test verification are required before a production rollout.

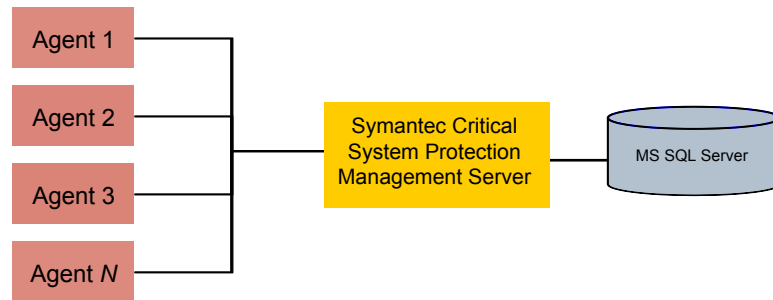
Planning an enterprise deployment

While installation of a small demonstration environment can be completed within an hour, enterprise deployments require careful planning. You first need to gain

a deeper understanding of the product and how it best integrates with your organization's overall security and compliance strategy.

Most development or test environments deploy at least a two-node management environment setup. One node hosts the Symantec Critical System Protection management server. The other node hosts the MS SQL Server instance that is needed for the Symantec Critical System Protection database. [Figure 1-3](#) shows a diagram of a simple setup.

Figure 1-3 Simple system setup

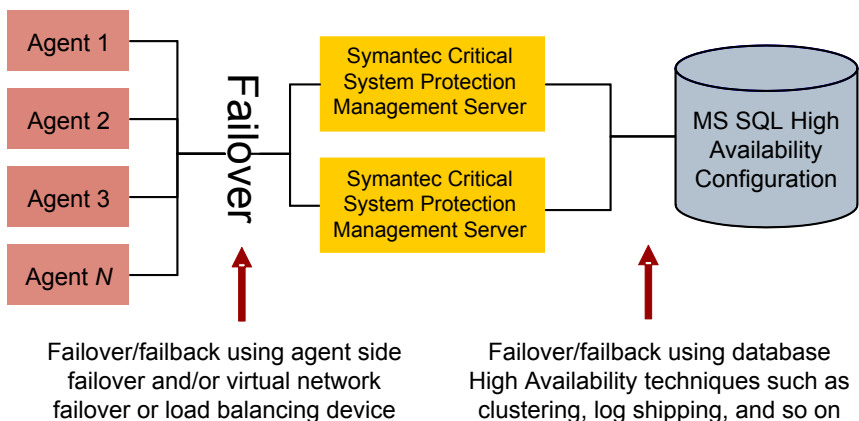


With appropriate hardware and configuration settings, this scenario can support 10,000 deployed agents. The Symantec Critical System Protection manager node, which is itself stateless, does not require excessive system resources and uses little disk I/O. It can also operate efficiently as a Guest Virtual Machine (GVM). Thus, configuring multiple manager systems on virtual machines provides for simple failover and easy setup.

A best practice is to install a Symantec Critical System Protection agent on each computer that hosts a management component. This setup lets you monitor and secure the systems from tampering. For example, in this two-node scenario, you should deploy a Symantec Critical System Protection Windows agent to protect the manager. You should deploy a second agent on the computer where the Symantec Critical System Protection SQL Server database resides. You should deploy appropriate protection and detection policies to these agents to activate enforcement.

Symantec Critical System Protection's J2EE application server and Microsoft SQL Server management components support high availability scenarios using common enterprise techniques. Many users add an additional Symantec Critical System Protection manager node to provide high availability. With the proper configuration, when one Symantec Critical System Protection manager fails, the remaining node can still service agent communications and console activities. [Figure 1-4](#) shows this configuration.

Figure 1-4 Symantec Critical System Protection high availability



For details, see the section that describes specifying alternate management servers in the *Symantec Critical System Protection Installation Guide*.

To further enhance the high availability scenario, some users choose to add a fourth system to avoid database system outages. Several SQL Server high availability techniques are available. For example, you can configure an active/passive SQL Server database cluster or use a SQL Server log shipping approach for near-real time database redundancy.

See the various documents on [Microsoft's Web site](#) for information about MS SQL high availability scenarios. For example:

[High Availability Solutions Overview](#)

[High Availability Solutions \(SQL Server\)](#)

About scaling

For a configuration that supports many more agents, such as 20,000 to 100,000 endpoints, you should add more front-end Symantec Critical System Protection management servers. You need the additional managers to handle agent communications. As a general rule, you can have 10,000 agents per system.

For these larger environments, you must configure the agent communication settings appropriately. The most notable setting is the agent polling check-in interval, which by default is five minutes. A setting of five minutes is appropriate for evaluation purposes, but not for large-scale deployments. You need to set it to a larger time interval, such as 30 minutes to 60 minutes for most agents. In the production environments that have few policy or configuration changes, agents do not need to check so frequently for changes. Most agents do not need to listen

for notification changes from the Symantec Critical System Protection manager. If the agents do not need to listen, then you can disable the default listening port, which is 2222 on most agents. Some examples of polling and listening configurations are as follows:

- Systems whose policy and configuration are relatively stable and that are typical with respect to your Symantec Critical System Protection deployment. Such systems should have a relatively long polling interval, such as 30-60 minutes. They should not listen for policy and configuration change notifications. This scenario should be the majority of agents in your deployment.
- Systems whose policy and configuration settings change significantly more often than your typical agents. This scenario can represent a small set of the systems that you use as the initial phase of rolling out policy or configuration changes. Therefore these systems have more frequent changes due to policy tuning. Or, this scenario can represent the high risk systems, such as in the DMZ, for which you want a fast reaction time to policy changes. These agents should have a shorter polling interval, 5-15 minutes, for example. They should not listen for policy and configuration change notifications. This scenario should represent a small percentage of your deployment to avoid putting a constant, heavy load on the management servers.
- Systems whose policy and configuration are relatively stable, but need a fast reaction time to infrequent policy or configuration changes. These systems can have a long polling interval, 30-60 minutes, for example. These systems should listen for policy and configuration change notifications and therefore you should maintain the default listening port, 2222. This scenario represents an alternative for high risk systems, such as in the DMZ, and for internal systems with unusually sensitive applications or data. This scenario should represent a small percentage of your deployment to avoid putting a heavy load on the management servers when you do change a policy or configuration.

Key database server installation considerations

A key question during the initial server installation is whether or not you should configure the Symantec Critical System Protection database to “Enable Unicode storage.” This option is enabled by default and is designed to work with all agents and policies that require double-byte character support. Examples of these languages are Japanese, Chinese, Korean, and other multi-byte languages.

The advantage of enabling unicode storage is that it allows Symantec Critical System Protection to support data from international operating systems. The disadvantage is that it affects system sizing and performance. Unicode storage consumes two bytes of data space for every character in the strings that are stored

in the database. Two bytes are used even when the source data is itself in a single-byte format, such as ASCII data. Symantec Critical System Protection stores strings such as Syslog and Windows event log messages, path names, and policy names. As a result, non-Unicode customers can experience increased database disk space usage and subsequent increased processing and backup. Many customers have only single-byte language requirements and do not need to use Unicode storage. You should choose to enable or disable this setting based on your understanding of your organization's deployment requirements.

Note: It is more difficult to later convert a Unicode database to a non-Unicode database than the other way round. You can experience a data loss of 2-byte character data when it is constrained to a 1-byte format.

Agent installation considerations for key features

Agent installation configures the appropriate networking for the environment. The agent installation configuration includes which Symantec Critical System Protection managers to communicate with, what ports to use, and how often to poll for changes. The initial Symantec Critical System Protection installation also determines whether key product features are enabled or not. The key agent feature installation considerations are discussed in greater detail in each section, as follows:

- Enabling the intrusion prevention feature
See [“About enabling intrusion prevention”](#) on page 18.
- Enabling the real-time file integrity monitoring feature in intrusion detection
See [“About enabling the real-time file integrity monitoring feature in intrusion detection”](#) on page 20.
- Creating agent registration groups
See [“About creating agent registration groups”](#) on page 21.

About enabling intrusion prevention

The intrusion prevention capability is enabled by default during agent installation on all Symantec Critical System Protection agents that support it.

See the latest [Symantec Critical System Protection Platform/Feature Matrix](#) on this Web page for up-to-date information about supported operating systems and virtual platforms.

Symantec recommends that you have intrusion prevention enabled. The prevention enforcement engine consists of the kernel mode drivers. When intrusion prevention is enabled, it is loaded on the agent computer on the next restart.

Immediately after installation, the intrusion prevention engine is configured to run with a Null policy. A Null policy contains no prevention rules and hence does not block any actions on the system. So, while the kernel driver intercepts all resource access operations by the operating system, it essentially operates in pass-through mode.

The key advantage of enabling the intrusion prevention feature with a Null policy is that the appropriate kernel drivers are loaded. Thus, the agent is ready to accept other prevention policies when you decide to deploy them. This configuration avoids any disruption to the normal operation of the critical server, since it does not require a restart to deploy the policy. If you plan to deploy prevention policies on an agent, you must leave the intrusion prevention feature enabled so that the agent can accept a prevention policy.

The slight disadvantage of leaving this setting enabled is its effect on performance. The drivers continue to intercept every resource request even when there is no benefit. This setting results in unnecessary overhead if the only requirement for the system is for file integrity monitoring or log monitoring. Symantec recommends that you deploy the agents in a lab or a quality assurance pre-production environment before you deploy the agent on production systems. Deploying the agents on a pre-production environment that closely represents the configuration and load of your production systems is the best way to understand the performance or throughput impact on a specific server that may have high number of processes, disk I/O activity, or network activity.

The intrusion prevention deployment by default also requires a system restart for the driver to load initially and observe all process launches from startup. Some production systems have very long periods of time before they can be stopped and restarted. You may prefer to deploy only detection features to avoid having to restart.

If you choose to continue with intrusion prevention enabled, it is best practice to restart the system immediately after installation activities have completed. Restarting immediately provides the best chance to discover any adverse system effects so that you can take action quickly. For example, you can disable the intrusion prevention capability until the underlying issue is corrected. Deferring a system restart after installation activities is a significant risk, since a system restart may not occur for a long time following the installation. Operational staff may erroneously attribute the issues that are seen during a later restart to factors unrelated to the earlier Symantec Critical System Protection installation. This mistake can in turn increase the time to resolution and overall costs. Thus, a best practice is to do at least one if not both of the following tasks:

- Restart immediately when you install or upgrade the agent with intrusion prevention enabled.

- Perform thorough testing up front to be sure there are no issues on production systems.

You can use the Symantec Critical System Protection console to visually determine the status of agents needing to be restarted. A blue triangle icon in the Asset display indicates that the prevention driver is due to be enabled at the next system restart. In addition, you can modify and run the Agent Detail query to see the status of agents pending restarts.

For more information about queries, see the [Symantec Critical System Protection Administration Guide](#).

Note: At the time of installation, you can selectively decide which agents should have intrusion prevention enabled.

For example you may initially roll out the prevention feature to a subset of the systems, such as a DMZ or particularly valuable system. You can then increase the prevention usage over time. If you disable intrusion prevention and want to enable it in the future, you must run the `sisipsconfig.exe` tool with the `-i` option and restart the computer. The `-i` option toggles the intrusion prevention service feature on and off. The `sisipsconfig.exe` tool is located in the `\Agent\IPS\bin` directory.

For information about the `sisipsconfig.exe` tool, you can see the following sources:

[Symantec Critical System Protection Administration Guide](#)

[Disable Intrusion Prevention \[IPS\] for Solaris|AIX|Linux|Windows](#)

Symantec recommends that all policy tuning be conducted on representative systems in a test or pre-production environment. This will help identify and correct any Symantec or customer-defined policy rules that inadvertently block resources required for the normal functioning of applications or operating system services.

About enabling the real-time file integrity monitoring feature in intrusion detection

During agent installation, the real-time file integrity monitoring capability is enabled by default on all Symantec Critical System Protection agents that support the feature. Symantec recommends that you enable the real-time file integrity monitoring capability. When this feature is enabled, the detection engine loads kernel mode drivers on the agent to aid in file monitoring activities. When real-time file integrity monitoring capability is enabled as part of intrusion detection, the driver functions are observational only. However, Symantec still recommends a

phased rollout on representative systems, to ensure that the operating system or business critical applications operate as expected.

See the latest [Symantec Critical System Protection Platform/Feature Matrix](#) on this Web page for up-to-date information about which operating systems support real-time file integrity monitoring.

The use of the drivers in the real-time file integrity monitoring feature has the following key advantages:

- The richness of the file integrity monitoring event data for local file systems. Data such as real time logging and the user and process that are involved in the file integrity events are logged.
- The flexibility in determining which changes are valid and which are invalid by adding exceptions for trusted installers or users performing the changes.

You can disable the real-time file integrity monitoring feature at the initial installation, if desired or necessary. The detection engine then does not load the real-time file integrity monitoring kernel mode drivers. It automatically uses polling-based file integrity monitoring instead. The file monitoring polling interval is defined in the policy or in the global default settings.

If you disable real-time file integrity monitoring and want to enable it in the future, you must run the `sisipsconfig.exe` tool with the **-rtfim** option. The `sisipsconfig.exe` tool is located in the `\Agent\IPS\bin` directory. The **-rtfim** option toggles the real-time file integrity monitoring feature on and off. On some operating systems, you may also need to restart the agent to load the real-time file integrity monitoring kernel mode drivers.

For information about the `sisipsconfig.exe` tool, see the following sources:

[How to disable File Integrity Monitoring \(FIM\) driver in Symantec Critical System Protection \(SCSP\) for Windows & AIX](#)

[Symantec Critical System Protection Administration Guide.](#)

About creating agent registration groups

You can distribute different policies and configuration settings to different computers by creating agent groups in the management console. First create the groups using the management console, set the different policies for the groups, and then associate the agents with the groups during installation. When the agent first registers with the Symantec Critical System Protection manager, it is placed into the groups that were specified during the agent installation. It immediately takes on the policies and settings that were defined for those groups.

The specification of group settings is optional during installation. If you leave the group settings empty, the agents register to the topmost node in the group

hierarchy. Later, you must manually move agents into appropriate groups in the console. However, this method can be tedious for a large number of agents. If your deployment numbers grow into the thousands, it is very beneficial to plan in advance. You should know how you want to manage policy deployments. Typically, you should determine how you want to group and organize your agents and specify your group settings during agent installation.

About automating agent installations

For larger organizations, you can automate the agent installation process effectively with provisioning tools, scripts, and other mechanisms. The agent installation kits all support silent, unattended installation. However, it is very important that you fully understand the types of computers that you install the agent on. You should understand their hardware specifications and software specifications, as well as other possible software that might conflict with the installation process.

The agent's operating system requirements are detailed in the *Symantec Critical System Protection Installation Guide* and in the Platform Support Matrix. Use Symantec Critical System Protection only with supported hardware and operating systems. Using the product with unsupported equipment and operating systems may result in unexpected behavior.

See the latest [Symantec Critical System Protection Platform/Feature Matrix](#) on this Web page for up-to-date information about supported operating systems and virtual platforms.

Once you decide on your installation choices, encode these choices in your automated deployment scripts. You can begin enforcing the Symantec Critical System Protection policies on agents immediately after agent installation and registration with the management server.

For information about all installation options, see the *Symantec Critical System Protection Installation Guide*.

Best practice for agent installation

As a best practice, Symantec recommends that you use Symantec Critical System Protection's full protection and monitoring features unless you have no immediate need to use the underlying features. Users who focus only on intrusion detection for the foreseeable future should disable the intrusion prevention features. Similarly, users who focus only on intrusion prevention should disable the real-time file integrity monitoring feature. Regulatory needs or a need for increased visibility into file integrity monitoring events should drive use of real-time file integrity monitoring by intrusion detection users.

Some users may be under pressure to broadly and quickly deploy only the file integrity monitoring and log monitoring features. They may not have time to perform adequate internal quality assurance testing. In this scenario, it would be safest to disable intrusion prevention and real-time file integrity monitoring and look into re-enabling those features when time permits.

About development, test, and production instances

You need to consider how your organization intends to enforce code promotion and change management practices in regard to your development, test, and production deployments. You should treat Symantec Critical System Protection like any other enterprise application in the following aspects:

- Source code version control
- Change processes related to policy or configuration changes
- Separation of duties

Some organizations have distinctly separate development, test, pre-production, and production environments for Symantec Critical System Protection. Policy changes are made first in the development environment and then promoted to the test environment. Eventually, they are promoted to the production environment. In classic source code provisioning, policies are exported from the Symantec Critical System Protection console and checked into a source control system. Then, in each phase of the code promotion cycle, the policy is retrieved from source control and imported into the Symantec Critical System Protection console. It is deployed from the console to its scope of systems. Eventually, the policy is verified and then imported into the production Symantec Critical System Protection console and deployed to operational production systems.

Some customers have combined their development environment and test environment into one Symantec Critical System Protection environment, while the production environment remains separate. Under some circumstances, users manage all classes of system from one central console. Symantec generally does not recommend a single control console across types of system. This can and has led to abuses and inadvertent situations where production systems are mistakenly used for test bed activities. Such mistakes can lead to errors in the policy rules or the configuration settings that affect production systems.

In all environments, you should grant only the access rights that any one individual requires. You can use the user roles in the Symantec Critical System Protection console to control who has access to specific agents or you can use access rights to agent groups. In addition, explicit user roles can limit access to basic functionality in the console. These predefined roles include Administrators, Authors, Managers, and Guests. By default, the built-in Administrator's role has

complete, unrestricted access to all available Symantec Critical System Protection features and tasks and agents. The Guest's role has only read access to the console. Because you need at least one role with unrestricted access, Symantec recommends that you do not modify the built-in Administrator's role.

For more information about default roles and defining custom roles, read the chapter titled "Using the Admin page," in the *Symantec Critical System Protection Administration Guide*.

About the incremental rollout of agent installations and policy deployment

In any software deployment, you should first deploy the software to a representative sample of systems and validate it before you deploy it more broadly. In many organizations, installation or operations teams deploy Symantec Critical System Protection agents, while security or compliance functions maintain the detection policies and prevention policies. Plan the agent installation and deployment activities in phases to account for the differences in the workload and activities on production systems. By planning the agent installation and deployment activities in phases, you can proactively address any issues that you encounter.

Symantec recommends that you first deploy agents configured with the appropriate policies on a small number of representative systems, to ensure that the systems and any business applications operate properly. Then, from the Symantec Critical System Protection console, deploy policies to a representative pool of agents. Verify that the systems and business functions operate properly with the appropriate security policies applied.

See the *Symantec Critical System Protection Administration Guide* for information about applying policies to agents.

Incremental rollout is especially valuable for blocks of like systems such as a dozen domain controllers or many similar database, Web servers, or cluster members. Symantec considers it a best practice to deploy to a small subset of like systems first, say one to three agents, to validate that critical business applications function as desired. Incremental rollouts provide the opportunity to manage each checkpoint and not have to troubleshoot both installation issues and policy issues at the same time.

Incremental testing and incremental deployment are especially valuable for the small set of extreme systems that are found in some environments. These extreme systems exhibit high end scaling that is out of proportion in comparison to the rest of the enterprise. Such systems are candidates for additional scrutiny and testing before production deployment. Examples of extreme systems include the following:

- Systems that have more than 1,000 concurrent processes where most other systems may have less than 100 processes.
- Systems that monitor changes on 500,000 files where most other systems generally have less than 10,000 files to monitor.
- Domain controllers, database servers, or other application servers that handle an order of magnitude more requests than typical systems.
- Systems that are near the limit of CPU or memory usage before Symantec Critical System Protection is installed.
- A clustered environment

See the following Web page for current information about the known issues that might affect your rollout:

[Symantec Critical System Protection known issues](#)

