



Using Management Center for Cisco Security Agents 4.0

Corporate Headquarters

Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
<http://www.cisco.com>
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 526-4100

Customer Order Number: OL-3957-01



THE SPECIFICATIONS AND INFORMATION REGARDING THE PRODUCTS IN THIS MANUAL ARE SUBJECT TO CHANGE WITHOUT NOTICE. ALL STATEMENTS, INFORMATION, AND RECOMMENDATIONS IN THIS MANUAL ARE BELIEVED TO BE ACCURATE BUT ARE PRESENTED WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. USERS MUST TAKE FULL RESPONSIBILITY FOR THEIR APPLICATION OF ANY PRODUCTS.

THE SOFTWARE LICENSE AND LIMITED WARRANTY FOR THE ACCOMPANYING PRODUCT ARE SET FORTH IN THE INFORMATION PACKET THAT SHIPPED WITH THE PRODUCT AND ARE INCORPORATED HEREIN BY THIS REFERENCE. IF YOU ARE UNABLE TO LOCATE THE SOFTWARE LICENSE OR LIMITED WARRANTY, CONTACT YOUR CISCO REPRESENTATIVE FOR A COPY.

The Cisco implementation of TCP header compression is an adaptation of a program developed by the University of California, Berkeley (UCB) as part of UCB's public domain version of the UNIX operating system. All rights reserved. Copyright © 1981, Regents of the University of California.

NOTWITHSTANDING ANY OTHER WARRANTY HEREIN, ALL DOCUMENT FILES AND SOFTWARE OF THESE SUPPLIERS ARE PROVIDED "AS IS" WITH ALL FAULTS. CISCO AND THE ABOVE-NAMED SUPPLIERS DISCLAIM ALL WARRANTIES, EXPRESSED OR IMPLIED, INCLUDING, WITHOUT LIMITATION, THOSE OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE.

IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THIS MANUAL, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

CCIP, CCSP, the Cisco Arrow logo, the Cisco *Powered* Network mark, Cisco Unity, Follow Me Browsing, FormShare, and StackWise are trademarks of Cisco Systems, Inc.; Changing the Way We Work, Live, Play, and Learn, and iQuick Study are service marks of Cisco Systems, Inc.; and Aironet, ASIST, BPX, Catalyst, CCDA, CCDP, CCIE, CCNA, CCNP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, the Cisco IOS logo, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Empowering the Internet Generation, Enterprise/Solver, EtherChannel, EtherSwitch, Fast Step, GigaStack, Internet Quotient, IOS, IP/TV, iQ Expertise, the iQ logo, iQ Net Readiness Scorecard, LightStream, MGX, MICA, the Networkers logo, Networking Academy, Network Registrar, *Packet*, PIX, Post-Routing, Pre-Routing, RateMUX, Registrar, ScriptShare, SlideCast, SMARTnet, StrataView Plus, Stratm, SwitchProbe, TeleRouter, The Fastest Way to Increase Your Internet Quotient, TransPath, and VCO are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the U.S. and certain other countries.

All other trademarks mentioned in this document or Web site are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (0304R)

Using Management Center for Cisco Security Agents V4.0
Copyright © 2003 Cisco Systems, Inc. All rights reserved.



Preface **xiii**

Objectives **xiii**

Intended Audience **xiii**

Typographical Conventions **xiv**

Obtaining Documentation **xv**

 Cisco.com **xv**

 Documentation CD-ROM **xv**

 Ordering Documentation **xvi**

 Documentation Feedback **xvi**

Obtaining Technical Assistance **xvi**

 Cisco.com **xvii**

 Technical Assistance Center **xvii**

 Cisco TAC Website **xviii**

 Cisco TAC Escalation Center **xix**

Obtaining Additional Publications and Information **xix**

CHAPTER 1

Product Overview **1-1**

 What the Cisco Security Agent Does **1-1**

 The Lifecycle of an Attack **1-2**

 How Cisco Security Agents Protect Against Attacks **1-3**

 Deployment Overview **1-4**

 Network Architecture **1-5**

 Cisco Security Agent Architecture **1-6**

 Preparing a Security Policy **1-7**

- Configuring Security Policies 1-8
- Communicating over Secure Channels 1-8
- Distributing Policy Updates 1-9
- Configuration Road Map 1-9

CHAPTER 2

Management Center for Cisco Security Agents Administration 2-1

- Overview 2-1
- Management Center for Cisco Security Agents Description 2-2
 - Browser Requirements 2-2
- About Management Center for Cisco Security Agents 2-3
- Accessing Management Center for Cisco Security Agents 2-4
 - Role-based Administration 2-5
 - Administration by Operating System 2-5
- Using Audit Trail 2-6
- Using Management Center for Cisco Security Agents 2-7
 - Creating, Saving, and Deleting Data 2-17
- Using the Correct Syntax 2-18

CHAPTER 3

Configuring Groups and Managing Hosts 3-1

- Overview 3-1
- Grouping Hosts Together 3-2
- Configuring Groups 3-3
 - No User Interaction Feature 3-7
- Building Agent Kits 3-8
 - Agent Kit Status 3-14
 - Agent Reboot vs. No Reboot 3-17
 - Agent Registration 3-18
 - Scripted Agent Installs and Uninstalls 3-19

- Registration Control 3-19
- Viewing Host Status 3-20
 - Changing Host Group Assignments 3-23
- Distributing Software Updates 3-26
 - Configuring Scheduled Software Updates 3-28

CHAPTER 4**Building Policies 4-1**

- Overview 4-1
- Developing a Security Policy 4-2
 - Providing Safe Access to Required Resources 4-3
- About Rules 4-5
 - Combining Policies 4-6
- Policy Components 4-6
 - Writing Rules: Allow vs. Deny 4-6
 - Writing Rules: Manipulating Precedence 4-8
 - Monitoring Access 4-9
 - Making a Policy Mandatory 4-9
- Querying the User 4-10
 - Caching Query Responses 4-12
- Configuring Policies 4-13
 - Adding Rules to a Policy 4-16
 - Filtering the Rules Display 4-17
 - Copying Rules between Policies 4-18
 - Comparing Configurations 4-19
 - Merging or Copying Policies 4-22
 - View Change History 4-22
 - Explanation of Rules 4-23
- Rules Common to Windows and UNIX 4-24
 - Agent Service Control 4-24

- Application Control 4-28
- Connection Rate Limit 4-33
- Data Access Control 4-35
- File Access Control 4-40
- File Monitor 4-45
- Network Access Control 4-49
- Windows Only Rules 4-54
 - COM Component Access Control 4-55
 - File Version Control 4-59
 - Kernel Protection 4-64
 - NT Event Log 4-67
 - Registry Access Control 4-70
 - Service Restart 4-74
 - Sniffer and Protocol Detection 4-77
- UNIX Only Rules 4-80
 - Network Interface Control 4-81
 - Resource Access Control 4-84
 - Rootkit / kernel Protection 4-87
 - Syslog Control 4-88
 - Syslog rule configuration examples 4-90
- Attaching Policies to Groups 4-92
- Using Test Mode 4-95
- Generating Rule Programs 4-97

CHAPTER 5

- Using System Correlation Rules 5-1**
 - Overview 5-1
 - Event Correlation and Heuristics 5-2
 - Network Worm Protection 5-3
 - Network Worm Event Correlation 5-4

Trojan Detection	5-4
Replicate Feature	5-9
Network Shield	5-9
Buffer Overflow	5-16
Global Events	5-18
Correlation	5-19

CHAPTER 6**Using Application Classes 6-1**

Overview	6-1
About Application Classes	6-2
Processes Created by Application Classes	6-2
Removing Processes from Application Classes	6-2
Shell Scripts and Application Classes	6-3
Included Application Classes	6-4
Preserving Application Process Classes	6-6
Configuring Static Application Classes	6-6
Configure Downloaded Content	6-9
Dynamic Application Classes	6-11
Configuring Dynamic Application Classes	6-13
Configure an Application-Builder Rule	6-15
Configure a Rule Using a Dynamic Application Class	6-19
Create New Application Classes from Rule Pages	6-20
Application Class Management	6-22

CHAPTER 7**Configuring Variables 7-1**

Overview	7-1
Where Variables are Used	7-2
Data Sets	7-3
File Sets	7-6

- Network Address Sets 7-11
- Network Services 7-14
- Registry Sets 7-17
 - Included Registry Sets 7-20
- COM Component Sets 7-21
 - COM Component Extract Utility 7-23

CHAPTER 8

Event Logging Alerts 8-1

- Overview 8-1
- The Event Log 8-2
 - Start date and End date 8-2
 - Minimum and Maximum Severity Settings 8-3
 - Host 8-3
 - Events / page 8-3
- Event Monitor 8-5
- Event Log Management 8-6
- How Logging Works 8-8
 - Verbose Logging 8-9
 - Logging and Query User Rules 8-9
- About the Event Management Wizard 8-9
 - Creating an Exception Rule 8-11
 - Creating a Logging Exception Rule 8-15
 - Creating an Analysis Job 8-17
- Event Sets 8-20
 - Third Party Access to Events 8-24
- Configuring Alerts 8-27
 - Generate an Alert Log File for Third Party Applications 8-34

CHAPTER 9**Generating Reports 9-1**

Overview 9-1

Types of Reports 9-2

Viewing Reports 9-2

Generating Reports 9-2

Events by Severity 9-2

Events by Group 9-4

Host Detail 9-5

Policy Detail 9-7

Group Detail 9-7

About the ActiveX Crystal Report Viewer 9-8

CHAPTER 10**Using Management Center for Cisco Security Agents Utilities 10-1**

Overview 10-1

Start and Stop Server Service 10-2

Start and Stop Agent Service 10-2

Backing Up Configurations 10-3

Restoring Backup Configurations 10-6

Free Up Disk Space on CSA MC (Insufficient Disk Space Event) 10-7

Using the COM Extract Utility 10-9

Manual Agent Data Filter Installation 10-10

Install Data Filter on Windows 10-10

Uninstall Data Filter on Windows 10-11

Install Data Filter on Solaris 10-12

Uninstall Data Filter on Solaris 10-12

Exporting and Importing Configurations 10-12

CHAPTER 11

Using Cisco Security Agent Profiler 11-1

- What is Profiler 11-1
- How Profiler Works 11-2
- The Analysis Process 11-3
 - Profiler Analysis Jobs 11-3
 - Creating, Saving, and Cancelling Job Data 11-4
- Configure an Analysis Job 11-6
 - Monitoring the Analysis Job 11-10
- Start Analysis 11-11
- Importing the Policy 11-12
- The Profiler Policy 11-14
 - Reviewing the Policy 11-15
 - Profiler Policy Methodology 11-15
 - Variable and Application Class Creation 11-17
- Profiler Reports 11-18
 - Report Components 11-19
 - File Event Reports 11-20
 - Registry Event Reports (Windows only) 11-21
 - COM Event Reports (Windows only) 11-21
 - Network Event Reports 11-22
 - Summary Reports 11-22
 - Working with Reports 11-23

CHAPTER 12

Policy Definition Guidelines 12-1

- Overview 12-1
- Analyzing Applications 12-2
- Configuring Policies—The Methodology 12-3
 - General Server Policy 12-5
 - Sample Web Server Policy 12-6

Combined General Server and Sample Web Server Policies	12-8
Reference	12-9

CHAPTER 13**Third Party Product Integration 13-1**

Overview	13-1
Cisco VPN Client Support	13-1
Cisco Security Monitor Integration Support	13-2
netForensics Integration Support	13-2
Check Point™ OPSEC™ Integration	13-2
Configuration Prerequisites	13-3
Integration Configuration	13-3

APPENDIX A**Cisco Security Agent Overview A-1**

Overview	A-1
Downloading and Installing	A-2
Network Shim Optional	A-4
The Agent User Interface	A-7
Responding to Pop-up Query Boxes	A-12
Suspend Agent Security	A-13
Installing Software Updates on Agents	A-14
Installing the UNIX Agent	A-16
Uninstall UNIX Agent	A-18
UNIX Agent csactl Utility	A-18

APPENDIX B**System Components B-1**

Overview	B-1
CSA MC Components	B-2
Agent Components	B-3

APPENDIX C

Third Party Copyright Notices C-1

Openssl license C-1

SSLEAY license C-2

Apache license C-4

TCL license C-5

Perl License: C-6

libwww License C-6

libpcap C-6

CMU-SNMP Libraries C-7

Open Market Inc., Fastcgi license C-8

CGIC License C-9

Mozilla 1.xx (libcurl) C-9

INDEX



Preface

Objectives

This manual describes how to configure the Management Center for Cisco Security Agents on Microsoft Windows 2000 operating systems and the Cisco Security Agent on Microsoft Windows XP, Microsoft Windows 2000, Microsoft Windows NT, and Solaris operating systems.

In addition to the information contained in this manual, consult the release notes for the latest information on the current release. Note that this manual does not provide tutorial information on the use of Windows and Solaris operating systems.

Intended Audience

This manual is intended for system managers or network administrators responsible for installing, configuring, and maintaining Management Center for Cisco Security Agents software. It is assumed that installers have a solid grounding in networking concepts and system management and have experience installing software on Windows operating systems.

Typographical Conventions

This manual uses the following conventions.

Convention	Purpose	Example
Bold text	User interface field names and menu options.	Click the Groups option. The Groups edit page appears.
<i>Italicized text</i>	Used to <i>emphasize</i> text.	You must <i>save</i> your configuration before you can deploy your rule sets.
Keys connected by the plus sign	Keys pressed simultaneously.	Ctrl+Alt+Delete
Keys not connected by plus signs	Keys pressed sequentially.	Esc 0 2 7
Monospaced font	Text displayed at the command line.	>ping www.example.com



Tip

Identifies information to help you get the most benefit from your product.



Note

Means *reader take note*. Notes identify important information that you should reflect upon before continuing, contain helpful suggestions, or provide references to materials not contained in the document.



Caution

Means *reader be careful*. In this situation, you might do something that could result in equipment damage, loss of data, or a potential breach in your network security.

Obtaining Documentation

Cisco provides several ways to obtain documentation, technical assistance, and other technical resources. These sections explain how to obtain technical information from Cisco Systems.

Cisco.com

You can access the most current Cisco documentation on the World Wide Web at this URL:

<http://www.cisco.com/univercd/home/home.htm>

You can access the Cisco website at this URL:

<http://www.cisco.com>

International Cisco websites can be accessed from this URL:

http://www.cisco.com/public/countries_languages.shtml

Documentation CD-ROM

Cisco documentation and additional literature are available in a Cisco Documentation CD-ROM package, which may have shipped with your product. The Documentation CD-ROM is updated regularly and may be more current than printed documentation. The CD-ROM package is available as a single unit or through an annual or quarterly subscription.

Registered Cisco.com users can order a single Documentation CD-ROM (product number DOC-CONDOCCD=) through the Cisco Ordering tool:

http://www.cisco.com/en/US/partner/ordering/ordering_place_order_ordering_tool_launch.html

All users can order monthly or quarterly subscriptions through the online Subscription Store:

<http://www.cisco.com/go/subscription>

Ordering Documentation

You can find instructions for ordering documentation at this URL:

http://www.cisco.com/univercd/cc/td/doc/es_inpk/pdi.htm

You can order Cisco documentation in these ways:

- Registered Cisco.com users (Cisco direct customers) can order Cisco product documentation from the Networking Products MarketPlace:

<http://www.cisco.com/en/US/partner/ordering/index.shtml>

- Nonregistered Cisco.com users can order documentation through a local account representative by calling Cisco Systems Corporate Headquarters (California, U.S.A.) at 408 526-7208 or, elsewhere in North America, by calling 800 553-NETS (6387).

Documentation Feedback

You can submit comments electronically on Cisco.com. On the Cisco Documentation home page, click **Feedback** at the top of the page.

You can e-mail your comments to bug-doc@cisco.com.

You can submit comments by using the response card (if present) behind the front cover of your document or by writing to the following address:

Cisco Systems
Attn: Customer Document Ordering
170 West Tasman Drive
San Jose, CA 95134-9883

We appreciate your comments.

Obtaining Technical Assistance

Cisco provides Cisco.com, which includes the Cisco Technical Assistance Center (TAC) website, as a starting point for all technical assistance. Customers and partners can obtain online documentation, troubleshooting tips, and sample

configurations from the Cisco TAC website. Cisco.com registered users have complete access to the technical support resources on the Cisco TAC website, including TAC tools and utilities.

Cisco.com

Cisco.com offers a suite of interactive, networked services that let you access Cisco information, networking solutions, services, programs, and resources at any time, from anywhere in the world.

Cisco.com provides a broad range of features and services to help you with these tasks:

- Streamline business processes and improve productivity
- Resolve technical issues with online support
- Download and test software packages
- Order Cisco learning materials and merchandise
- Register for online skill assessment, training, and certification programs

To obtain customized information and service, you can self-register on Cisco.com at this URL:

<http://tools.cisco.com/RPF/register/register.do>

Technical Assistance Center

The Cisco TAC is available to all customers who need technical assistance with a Cisco product, technology, or solution. Two types of support are available: the Cisco TAC website and the Cisco TAC Escalation Center. The type of support that you choose depends on the priority of the problem and the conditions stated in service contracts, when applicable.

We categorize Cisco TAC inquiries according to urgency:

- Priority level 4 (P4)—You need information or assistance concerning Cisco product capabilities, product installation, or basic product configuration. There is little or no impact to your business operations.
- Priority level 3 (P3)—Operational performance of the network is impaired, but most business operations remain functional. You and Cisco are willing to commit resources during normal business hours to restore service to satisfactory levels.
- Priority level 2 (P2)—Operation of an existing network is severely degraded, or significant aspects of your business operations are negatively impacted by inadequate performance of Cisco products. You and Cisco will commit full-time resources during normal business hours to resolve the situation.
- Priority level 1 (P1)—An existing network is “down,” or there is a critical impact to your business operations. You and Cisco will commit all necessary resources around the clock to resolve the situation.

Cisco TAC Website

The Cisco TAC website provides online documents and tools to help troubleshoot and resolve technical issues with Cisco products and technologies. To access the Cisco TAC website, go to this URL:

<http://www.cisco.com/tac>

All customers, partners, and resellers who have a valid Cisco service contract have complete access to the technical support resources on the Cisco TAC website. Some services on the Cisco TAC website require a Cisco.com login ID and password. If you have a valid service contract but do not have a login ID or password, go to this URL to register:

<http://tools.cisco.com/RPF/register/register.do>

If you are a Cisco.com registered user, and you cannot resolve your technical issues by using the Cisco TAC website, you can open a case online at this URL:

<http://www.cisco.com/tac/caseopen>

If you have Internet access, we recommend that you open P3 and P4 cases online so that you can fully describe the situation and attach any necessary files.

Cisco TAC Escalation Center

The Cisco TAC Escalation Center addresses priority level 1 or priority level 2 issues. These classifications are assigned when severe network degradation significantly impacts business operations. When you contact the TAC Escalation Center with a P1 or P2 problem, a Cisco TAC engineer automatically opens a case.

To obtain a directory of toll-free Cisco TAC telephone numbers for your country, go to this URL:

<http://www.cisco.com/warp/public/687/Directory/DirTAC.shtml>

Before calling, please check with your network operations center to determine the Cisco support services to which your company is entitled: for example, SMARTnet, SMARTnet Onsite, or Network Supported Accounts (NSA). When you call the center, please have available your service agreement number and your product serial number.

Obtaining Additional Publications and Information

Information about Cisco products, technologies, and network solutions is available from various online and printed sources.

- The *Cisco Product Catalog* describes the networking products offered by Cisco Systems, as well as ordering and customer support services. Access the *Cisco Product Catalog* at this URL:

http://www.cisco.com/en/US/products/products_catalog_links_launch.html

- Cisco Press publishes a wide range of networking publications. Cisco suggests these titles for new and experienced users: *Internetworking Terms and Acronyms Dictionary*, *Internetworking Technology Handbook*, *Internetworking Troubleshooting Guide*, and the *Internetworking Design Guide*. For current Cisco Press titles and other information, go to Cisco Press online at this URL:

<http://www.ciscopress.com>

- *Packet* magazine is the Cisco quarterly publication that provides the latest networking trends, technology breakthroughs, and Cisco products and solutions to help industry professionals get the most from their networking investment. Included are networking deployment and troubleshooting tips, configuration examples, customer case studies, tutorials and training, certification information, and links to numerous in-depth online resources. You can access *Packet* magazine at this URL:

<http://www.cisco.com/go/packet>

- iQ Magazine is the Cisco bimonthly publication that delivers the latest information about Internet business strategies for executives. You can access iQ Magazine at this URL:

<http://www.cisco.com/go/iqmagazine>

- Internet Protocol Journal is a quarterly journal published by Cisco Systems for engineering professionals involved in designing, developing, and operating public and private internets and intranets. You can access the Internet Protocol Journal at this URL:

http://www.cisco.com/en/US/about/ac123/ac147/about_cisco_the_internet_protocol_journal.html

- Training—Cisco offers world-class networking training. Current offerings in network training are listed at this URL:

http://www.cisco.com/en/US/learning/le31/learning_recommended_training_list.html



Product Overview

What the Cisco Security Agent Does

Cisco Security Agents provides intrinsic, distributed security to your enterprise by deploying agents that defend against the proliferation of attacks across networks and systems. These Cisco Security Agents enforce a set of policies provided by the Management Center for Cisco Security Agents and selectively applied to system nodes by the network administrator.

Operating under the direction of assigned policies, Cisco Security Agents provide strong system resource protection, tying together the auditing and control of network, file, and registry actions, as well as controlling COM component access.

This section contains the following topics.

- [How Cisco Security Agents Protect Against Attacks, page 1-3](#)
- [Deployment Overview, page 1-4](#)
- [Preparing a Security Policy, page 1-7](#)
- [Communicating over Secure Channels, page 1-8](#)
- [Distributing Policy Updates, page 1-9](#)
- [Configuration Road Map, page 1-9](#)

The Lifecycle of an Attack

When your network is targeted for attack, an assault is typically launched in a series of steps. Each step of an attack is often dependent upon the previous step being successful. [Table 1-1](#) displays the common evolution of an attack.

Table 1-1 *Lifecycle of an Attack*

Attack Action	Network Manifestation
Probe	<ul style="list-style-type: none"> • ping server IP addresses • run traceroute on IP addresses • sniff passwords • impersonate mail users
Penetrate	<ul style="list-style-type: none"> • email attachments • Java applets and ActiveX controls • buffer overflows • backdoors and trojans
Persist	<ul style="list-style-type: none"> • weaken security settings • install new services
Propagate	<ul style="list-style-type: none"> • email • Internet connections • IRC • FTP • infected file shares
Paralyze	<ul style="list-style-type: none"> • reformat disks • destroy or corrupt data • drill security holes • crash computers • consume work cycles • steal confidential data

How Cisco Security Agents Protect Against Attacks

Cisco Security Agents different from anti-virus and network firewall software in that it doesn't prevent users from accessing technologies they require. It assumes that users are going to put their systems at risk by making use of a wide range of Internet resources. Keeping this in mind, Cisco Security Agents install and work at the kernel level, controlling network actions, local file systems, and other system components, maintaining an inventory of what actions may be performed on the system itself. This way, malicious system actions are immediately detected and disabled while other actions are allowed. Both actions take place transparently, without any interruption to the user.

If an encrypted piece of malicious code finds its way onto a system via email, for example, as it attempts to unexpectedly execute or alter Cisco Security Agent-protected system resources, it is immediately neutralized and a notification is sent to the network administrator.

The Cisco Security Agents protect systems using policies which network administrators configure and deploy. These policies can allow or deny specific system actions. The Cisco Security Agents must check whether an action is allowed or denied before any system resources are accessed and acted upon.

Specifically, rule policies provide administrators with the ability to control access to system resources based on the following parameters:

- What resource is being accessed.
- What operation is being invoked.
- Which application is invoking the action.

The resources in question here may be either system resources or network resources such as mail servers.

When any system actions that are controlled by specific rules are attempted and allowed or denied accordingly, a system event is logged and sent to the administrator in the form of a configurable notification type such as email, pager or custom script.

Deployment Overview

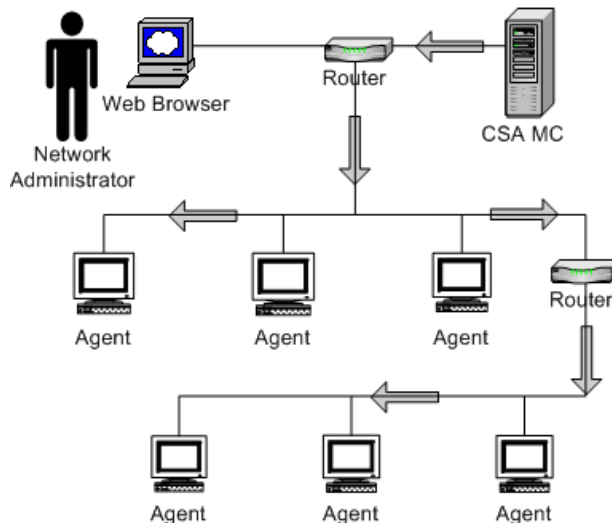
Management Center for Cisco Security Agents contains two components:

- CSA MC: installs on designated Windows 2000 systems and includes a configuration database server and a web-based user interface.
- Cisco Security Agent (the agent): installs on server and desktop Windows XP systems, Windows 2000 systems, Windows NT systems, and Solaris server systems across your enterprise network.

Using CSA MC, you assemble your network machines into specified groups and then attach security policies to those groups. All configuration is done through the web-based user interface and then deployed to the agents.

The network example shown in [Figure 1-1](#) illustrates a basic deployment scenario. CSA MC software is installed on a system which maintains all policy and host groups. The administration user interface is accessed securely using SSL (Secure Sockets Layer) from any machine on the network that can connect to the server and run a Web browser. Use the web-based interface to deploy your policies from CSA MC to the agents across your network.

Figure 1-1 Policy Deployment



Network Architecture

The CSA MC architecture model consists of a central management center which maintains a database of policies and system nodes all of which have Cisco Security Agent software installed on their desktops and servers.

Agents register with CSA MC. CSA MC checks its configuration database for a record of the system. When the system is found and authenticated, CSA MC deploys a configured policy for that particular system or grouping of systems.

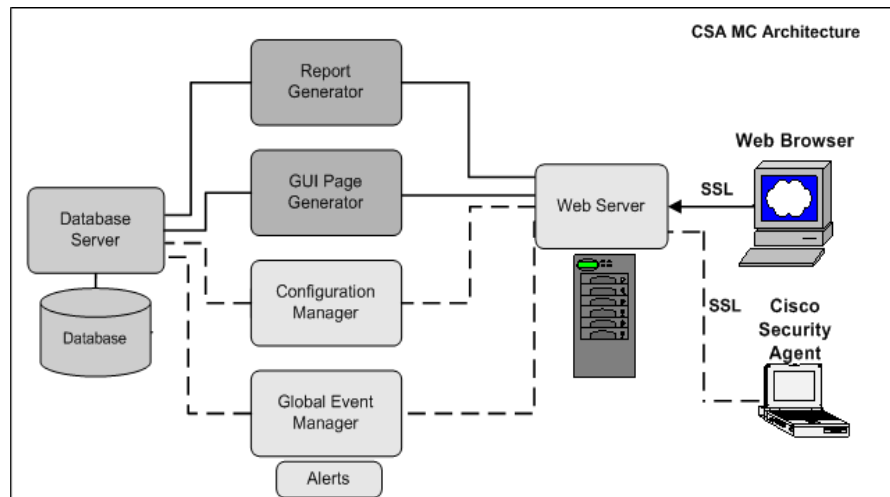
The Cisco Security Agent software now continually monitors local system activity and polls to the CSA MC at configurable intervals for policy updates. It also sends triggered event alerts to the CSA MC's global event manager. The global event manager examines system event logs and, based on that examination, may trigger an alert notification to the administrator or cause the agent to take a particular action.



Note

See [Appendix B, "System Components"](#) for detailed information on product architecture.

Figure 1-2 CSA MC Architecture



Cisco Security Agent Architecture

The Cisco Security Agent software installs locally on each system node and intercepts the operations of that system. A network application interceptor sits at the application level and intercepts all application operations. Other Cisco Security Agent mechanisms intercept network traffic, file actions, and system registry actions while the rule/event correlation engine controls all agent mechanisms watching for any events that trigger an agent policy.

Figure 1-3 Cisco Security Agent Software Architecture(Windows)

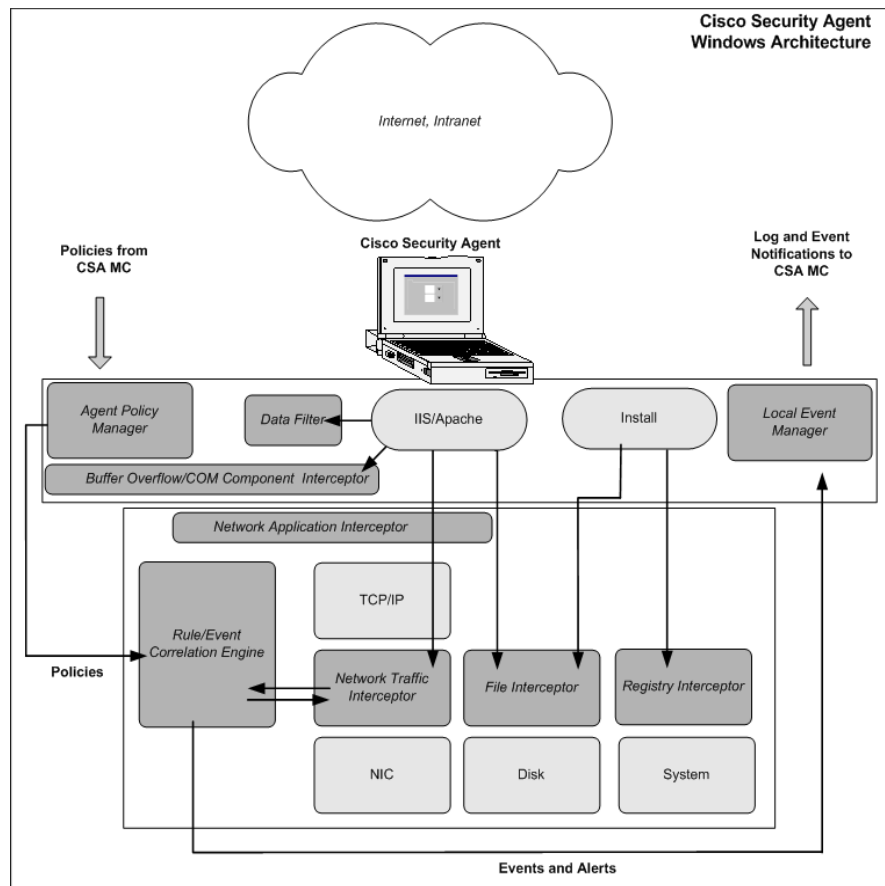
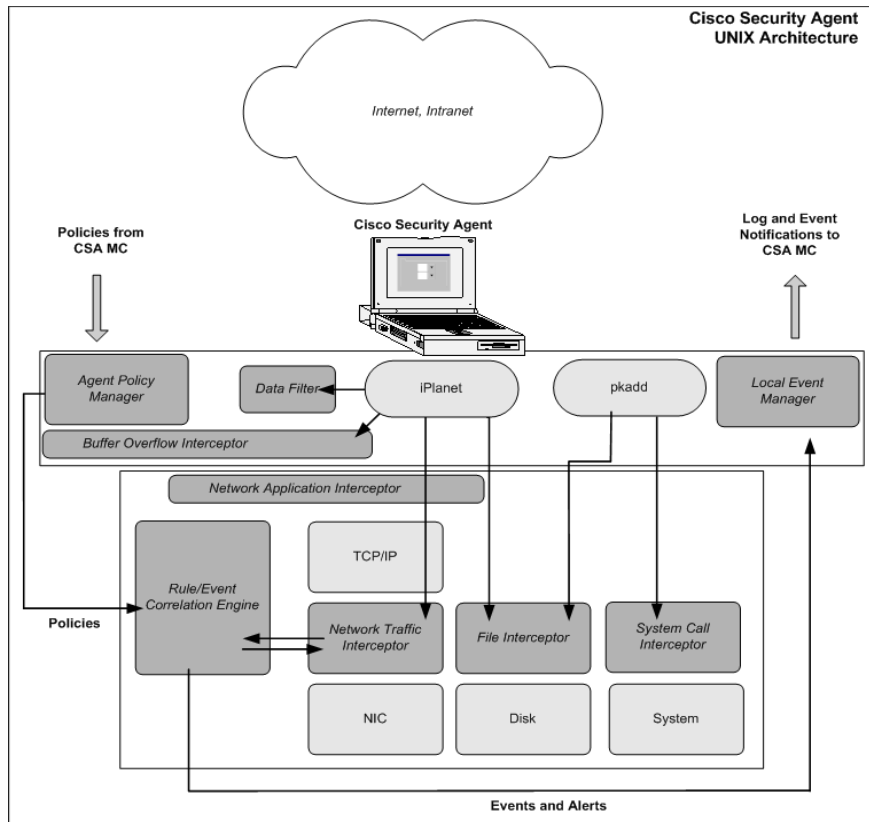


Figure 1-4 Cisco Security Agent Software Architecture (UNIX)



Preparing a Security Policy

It is crucial to have a carefully planned corporate security policy in place before you attempt to configure the Management Center for Cisco Security Agents. You must understand exactly what network resources and services you want to protect in order to adequately scale a set of policies that safeguard those valuable organizational resources. Keep in mind, a corporate security policy should allow the user community to easily access required resources, while protecting that community from the dangers those open resources can represent.

To help achieve this goal, CSA MC ships with a variety of rule templates and pre-configured policies. The policies you configure and deploy become the foundation of your security policy.

Configuring Security Policies

A policy is a collection of rules. The policy itself acts as the container for these rules and as the unit of attachment to groups. Machines with similar security needs are grouped together and assigned one or more policies that specifically target the needs of the group.

When you are creating rules for your policies, targeting the needs of machine groupings is central to your overall security plan. You can base these security needs on various criteria. For example, the concerns you have for your web servers may require you to group them separately from your mail servers based on the types of policies each set of servers require. Therefore, you could place your web servers into a common group, create rules that protect those servers from having their cgi files and html files written to (for example), and then attach the policy that contains these rules to the web servers group.

When first configuring and deploying policies, it is useful to put them into Test Mode. In Test Mode, the policies are not "live." The Cisco Security Agent will not deny any action or operation even if an associated policy says it should be denied. Instead, the agent will allow the action but log an event if a deny or query rule is triggered and log an event when an allow rule with logging enabled is triggered. This helps you to understand the impact of deploying a policy on a host before enforcing it.

Communicating over Secure Channels

All communications between the Management Center for Cisco Security Agents server system and systems accessing the browser-based user interface are protected using SSL (Secure Sockets Layer). Administrator authentication is also provided via the required entry of a username and password to authenticate and initiate each management session. Additionally, communications between the management server and the agents are passed over SSL.

See the Installation Guide for information on importing certificates and connecting securely over SSL.

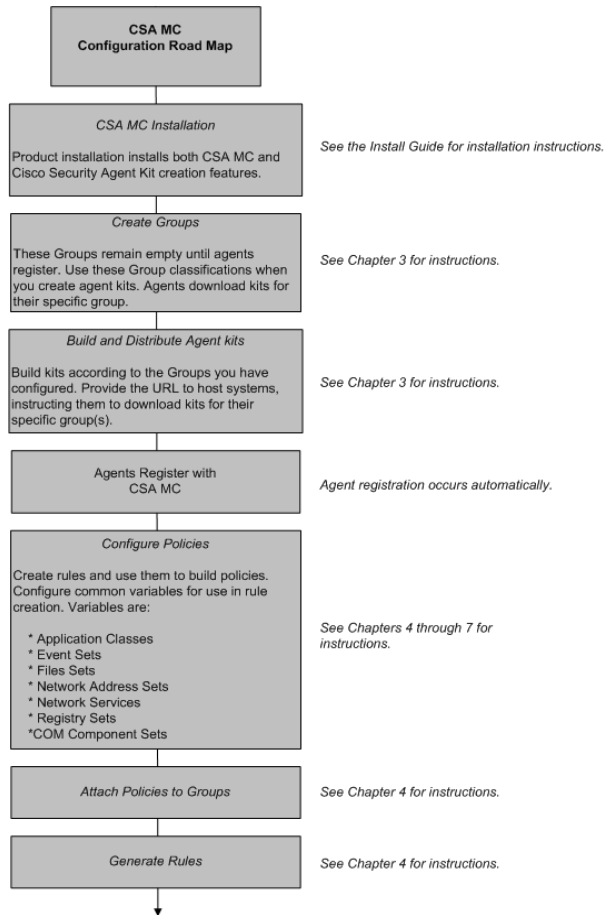
Distributing Policy Updates

At configurable time intervals (the default is every 10 minutes), Cisco Security Agents on the network poll in to CSA MC to check for updated rule sets. See [Chapter 3, “Configuring Groups and Managing Hosts”](#) for details.

When a rule is triggered on a system, the agent sends its event notifications to CSA MC. CSA MC identifies the agent, examines the event notifications presented by the agent and correlates this information.

Configuration Road Map

There are several elements you must configure to create the policies that are distributed to the agents. First, you must configure host groups and create Cisco Security Agent kits. Then once agents are installed on systems throughout your network, they register with CSA MC. Once this occurs they are automatically placed into their assigned groups. When you Generate rules, agents receive the policies intended for them. Refer to the following CSA MC configuration roadmap.

Figure 1-5 CSA MC Configuration Road Map



Management Center for Cisco Security Agents Administration

Overview

The Management Center for Cisco Security Agents supports editing of the database by multiple administrators. Administrators must identify themselves and authenticate to CiscoWorks before they can access any CSA MC configuration data.

CSA MC's web-based user interface provides secure access to the database from anywhere on the network. All changes to the database are logged. The logged information includes a summary description of the modification, the time the changes were made, and the identity of the administrator who made the changes.

This section contains the following topics.

- [Management Center for Cisco Security Agents Description, page 2-2](#)
- [About Management Center for Cisco Security Agents, page 2-3](#)
- [Accessing Management Center for Cisco Security Agents, page 2-4](#)
- [Using Audit Trail, page 2-6](#)
- [Using Management Center for Cisco Security Agents, page 2-7](#)
- [Using the Correct Syntax, page 2-18](#)

Management Center for Cisco Security Agents Description

The Management Center for Cisco Security Agents (CSA MC) is a web-based user interface which can be accessed from the CiscoWorks UI on any machine connected to the Internet running a web browser. Through CSA MC, administrators configure all aspects of the Cisco Security Agent product.

Refer to the following sections for CSA MC navigation and configuration details.

Browser Requirements

The browser you use to access CSA MC through CiscoWorks must meet the following requirements.

Internet Explorer:

- Version 5.5 or higher
- You must have cookies enabled. This means using a maximum setting of "medium" as your Internet security setting. Locate this feature from the following menu, Tools>Internet Options. Click the Security tab.
- JavaScript must be enabled.

Netscape:

- Version 6.2 or higher
- You must have cookies enabled. Locate this feature from the following menu, Edit>Preferences>Advanced.
- JavaScript must be enabled.



Note

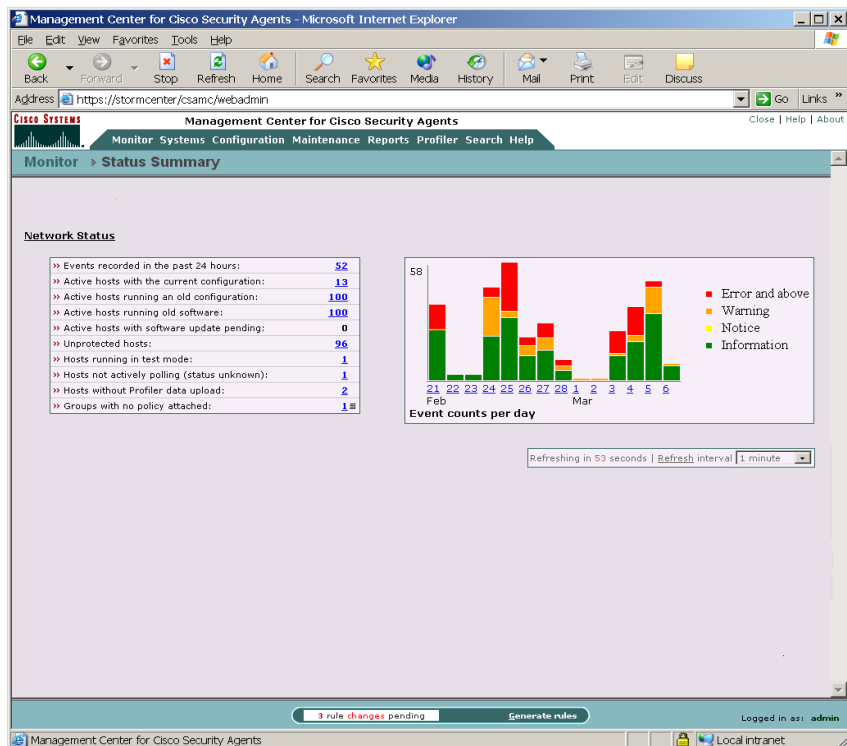
SSL must be enabled on the CiscoWorks UI to access CSA MC. Enable SSL in CiscoWorks from the **Server Configuration** drawer. Go to **Administration>Security Management>Enable/Disable SSL**. Click the **Enable** button in the right pane. You may have to restart both the CiscoWorks and CSA MC services for this change to take effect.

About Management Center for Cisco Security Agents

All Cisco Security Agent policies are configured and deployed through the CSA MC web-based user interface. CSA MC also provides a reporting tool, letting you generate reports with varying views of your network enterprise health and status. Providing an HTML web-based user interface allows an administrator to access CSA MC from any machine connected to the Internet running a web browser.

CSA MC provides a menu bar for easy navigation among the configurable administrator task items. Configurable items are displayed in drop-down menus that appear when you move the mouse over a category in the menu bar. When you click on an item, the properties and status for that item are displayed.

Figure 2-1 CSA MC Status View



Accessing Management Center for Cisco Security Agents

An administrator accesses CSA MC from the CiscoWorks UI. An initial administrator account was created as part of the CiscoWorks installation process. Once that administrator account is entered to login into CiscoWorks, it is not necessary to login again to CSA MC.

- To access CSA MC locally on the system hosting the CSA MC software, launch the CiscoWorks UI from **Start> Programs>CiscoWorks>CiscoWorks**. Log in to CiscoWorks.
- To access CSA MC from a remote location, launch a browser application and enter

```
http://<ciscoworks system hostname>:1741
```

For example, enter `http://stormcenter:1741`

From the CiscoWorks UI, the **Security Agents** item is located in the **VPN/Security Management Solution** “drawer.” Expand the **Management Center** or the **Administration>Management Center** folders.

**Note**

To launch CSA MC from CiscoWorks, the CiscoWorks UI must have SSL enabled. Refer to *Installing Management Center for Cisco Security Agents* for detailed information.

Role-based Administration

CSA MC can have multiple administrators, all with secure access to configuration data. Administrators can have different levels of CSA MC database access privileges. The initial administrator created by the CiscoWorks installation automatically has configuration privileges.

CiscoWorks/CSA MC Administrator Roles:

- **Configure**—If the CiscoWorks administrator has the Network Administrator or System Administrator option enabled, this provides full read and write access to the CSA MC database.
- **Deploy**—If the CiscoWorks administrator has only the Network Operations option enabled, this provides full read and partial write access to the CSA MC database. Administrators can manage hosts and groups, attach policies, create kits, schedule software updates, and perform all monitoring actions.
- **Monitor**— If the CiscoWorks administrator has none of the roles listed in the first two bullets enabled, this provides administrators with read access to the entire CSA MC database. Administrators can also create reports, alerts, and event sets.



Note

To view or edit your CiscoWorks administrator profile, in the CiscoWorks UI, go to **Server Configuration>Administration>Setup>Security>Modify My Profile**.

Administration by Operating System

Administrators can select to configure and deploy policies to only UNIX systems, to only Windows systems, or to All operating system types. If you leave the default of "All" for configuration items at the top of item list pages, you must then select an operating system type when you configure items such as policies, groups, and agent kits.

But if you do not select All and select UNIX, for example from the pulldown list at the top a page, then when you configure some CSA MC items, you will only see other UNIX configurations (if any) and you will only be able to select UNIX-specific configuration types. If you are only configuring or deploying

policies to UNIX system, this would be desirable as it is likely that you will only want to see those items. The same holds true for administrators deploying policies to only Windows systems.

**Caution**

Configuration items for UNIX and Windows cannot be combined.

Using Audit Trail

Accessible from the Maintenance drop-down list in the menu bar, the Audit Trail page displays a list of changes administrators have made to the CSA MC database. These changes are displayed according to the following information:

- The change itself.
- The "type" of change (configuration category: policies, file sets, groups, etc.).
- The date and time the change was made.
- The administrator who made the change.

Click the **Change Filter** link to edit the audit trail viewing parameters according to the following:

- Start date (enter date parameters using the same formats as in the Event Log).
- End date.
- The administrator who made the changes.
- The change type (configuration category: policies, file sets, groups, etc.).
- The number of changes to display per viewing page.

Using Management Center for Cisco Security Agents

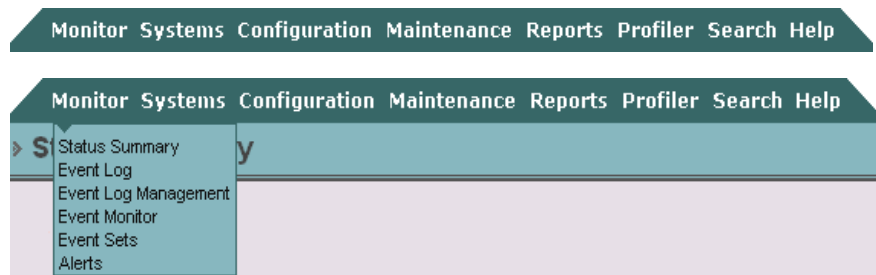
The following sections describe the various components you should understand in order to configure Cisco Security Agents using CSA MC.

Menu Bar

The menu bar at the top of CSA MC provides links to all configuration pages and list views. Arrows indicate that there are subcategories for which you can choose from those top-level main items (*see* [Figure 2-2](#)). These subcategories appear when you move the mouse over the main item itself.

When you select an item from the menu bar, the list view for that item appears.

Figure 2-2 Menu Bar



The configuration options available from each menu bar item are as follows.


- **Monitor:** The items available from the Monitor drop-down list provide tools for viewing system status and log files. You can also set alerts and alert parameters from here. (See [Chapter 8, “Event Logging Alerts.”](#))

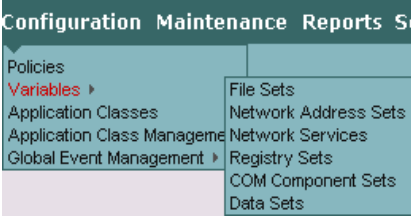



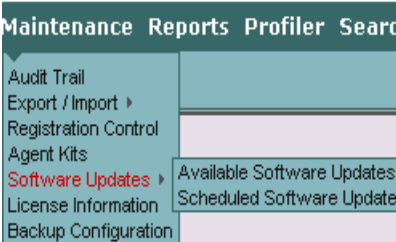
- **Systems:** The items available from the Systems drop-down list let you configure the groups which agent host systems are placed into when they register with CSA MC. (See [Chapter 3, “Configuring Groups and Managing Hosts.”](#))




- Configuration: It is from the Configuration drop-down list that you access most of the pages you need to configure your policies for agents. This list provides links to the rule pages you use to develop your policies, as well as links to application classes and variables. (See [Chapter 4, “Building Policies.”](#))


- Variables, such as file sets and network addresses, which are the building blocks for policies, are accessible from the cascading menu that appears when you move your mouse over the Variables item in the Configuration drop-down list. (See [Chapter 7, “Configuring Variables.”](#))


- Maintenance: The items available from the Maintenance drop-down list let you build agent kits, import and/or export configuration files, distribute available software updates, and backup your database configuration. When you move your mouse over the Export/Import and Software Updates items, you can select further options from the cascading menus that appear. (See this Chapter for administrator configuration details, also see [Chapter 3, “Configuring Groups and Managing Hosts.”](#))



- Reports: The items available from the Reports drop-down list let you generate reports by various categories such as event severity level, by the group(s) that generated the event and by individual host systems. (See [Chapter 9, “Generating Reports.”](#))



- Profiler: The item available from the Profiler drop-down list lets you configure analysis jobs for the purpose of analyzing applications and creating policies. (See [Chapter 11, “Using Cisco Security Agent Profiler.”](#))
- Search: Use the selections available from the Search drop-down list to search for a specific configuration item in the CSA MC database. You can limit your search to Hosts, Groups, Policies, Rules, Variables, Application Classes, or All by selecting one of those options from the available Search drop-down list. Each option has its own criteria by which you can search.



Using Search

Once you select a category to search on from the Search drop-down list, enter all or part of the name of the item for which you are searching in the Find field. (See [Figure 2-3.](#))

To further control your search, select one or more of the following checkboxes.

- Show references: Select this checkbox to also display configuration items which reference the name being searched for. Clicking on the referenced item in the right column lets you access the configuration(s) using the string value.
- Search on description: Select this checkbox to search for the string value in Description fields.
- Search all fields: Select this checkbox to search all database fields (including Description fields) for the string value.

You can limit Results per page by entering a value here (25 is the default). Click the **Find** button. Results are displayed as links. Click the item link to go to its configuration view.



Note

The search page does not search the event database.

(Use the Delete button to remove found items from the database. Once an item is deleted here, it cannot be recovered.)

- **Replace:** From the Search menu, Policies, Variables, and Application Classes allow you to perform a search and replace on items. Once you've selected a category from the Search menu, you can click the Replace link to access a pop-up box. In that box, you select references to an item and replace it where it appears with another item that you select. For example, search on Policies, click the Replace link and you can replace a policy that is attached to a group(s) with another policy that you select. Selecting the Preview checkbox allows to you see where all references will be replaced before you actually do the replacement.

The Hosts search page lets you search for hosts according to the time they last polled in to CSA MC. For time frame text, you can enter a relative time in the available field using any of the following terms: tomorrow, now, next, year, month, week, day, hour, minute, and second.

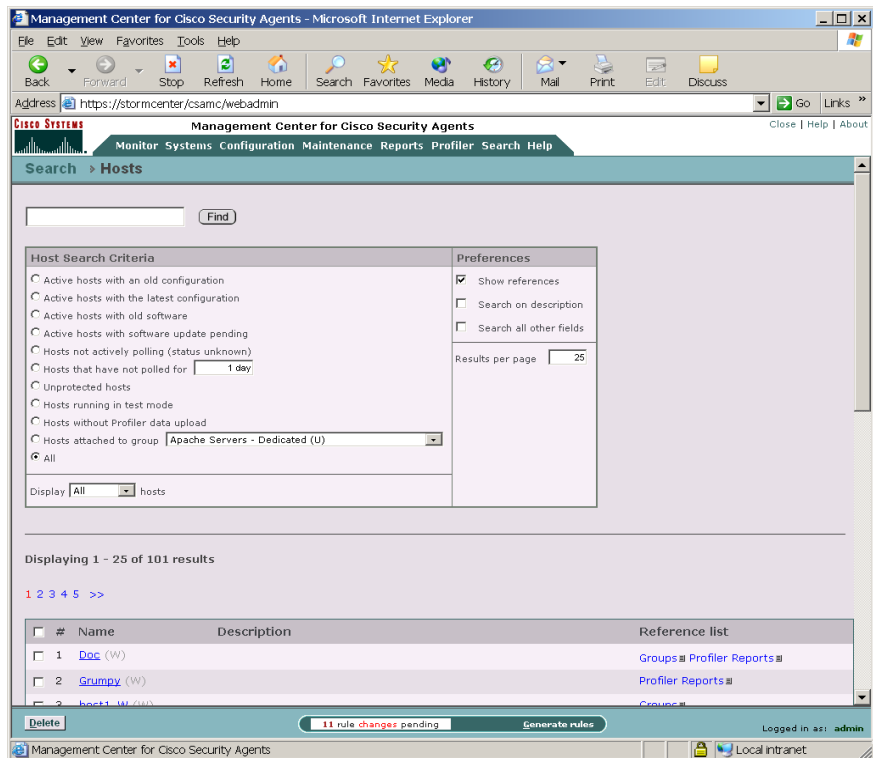
Enter a specific time using any of the follow time formats: hh:mm:ss. If no meridian (AM or PM) is specified, hh is interpreted on a 24 hour clock (0-23). Note that entering minutes and/or seconds is optional.

Enter a specific month and day with optional year in the formats: mm/dd/?/yy?, monthname dd ?, yy? (Note that the question marks here indicate that the information inside the question marks is optional.) The default year is the current year.

You can also search for hosts according to those with "old rule sets", "the latest rule set", "old (outdated) software", "those with pending software updates", "hosts not actively polling" see Not active hosts for details on this item, "hosts that have not polled in for a specified time", and "unprotected hosts." Unprotected hosts are hosts which are not members of any group or are members of a group which has no policies.

See the [“Viewing Host Status” section on page 3-20](#) for more information on Hosts.

Figure 2-3 Search Feature

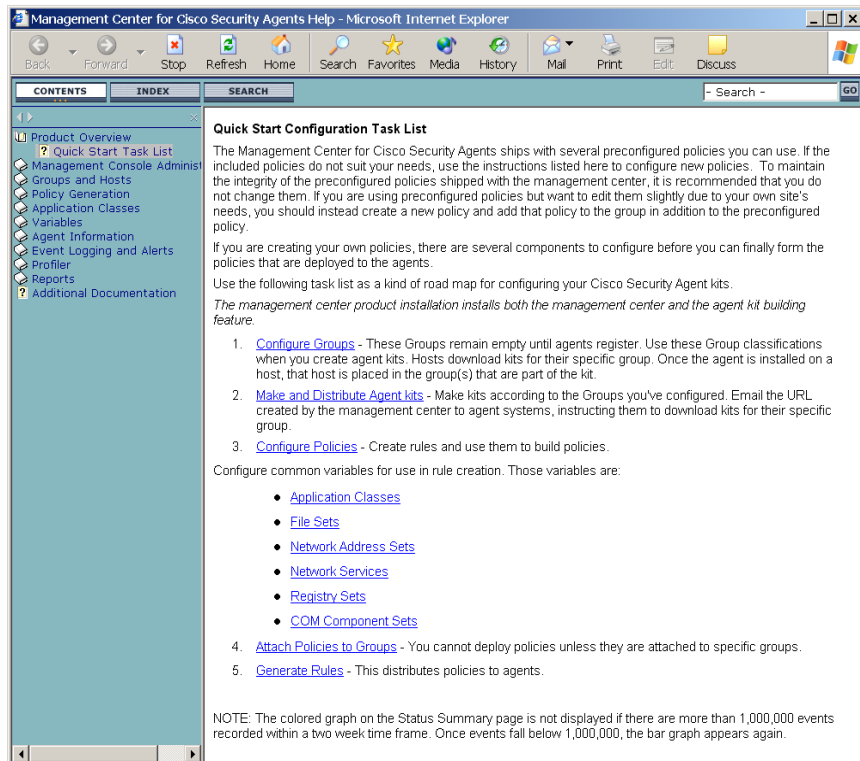


- **Help:** In addition to this configuration guide, CSA MC provides online help. When you click Help in the far right of the menu bar, you can select Online Help or you can click a link for the Technical Support web site. When you select Online help, a new browser window is launched. This window contains help information on the configuration item from which you have accessed the help. To view help on other topics, click the corresponding topic link in the Contents frame of the help window.



You can also access Quick Help for fields that provide question marks beside them. Quick Help provides syntactical information for specific text fields.

Figure 2-4 Main Help System

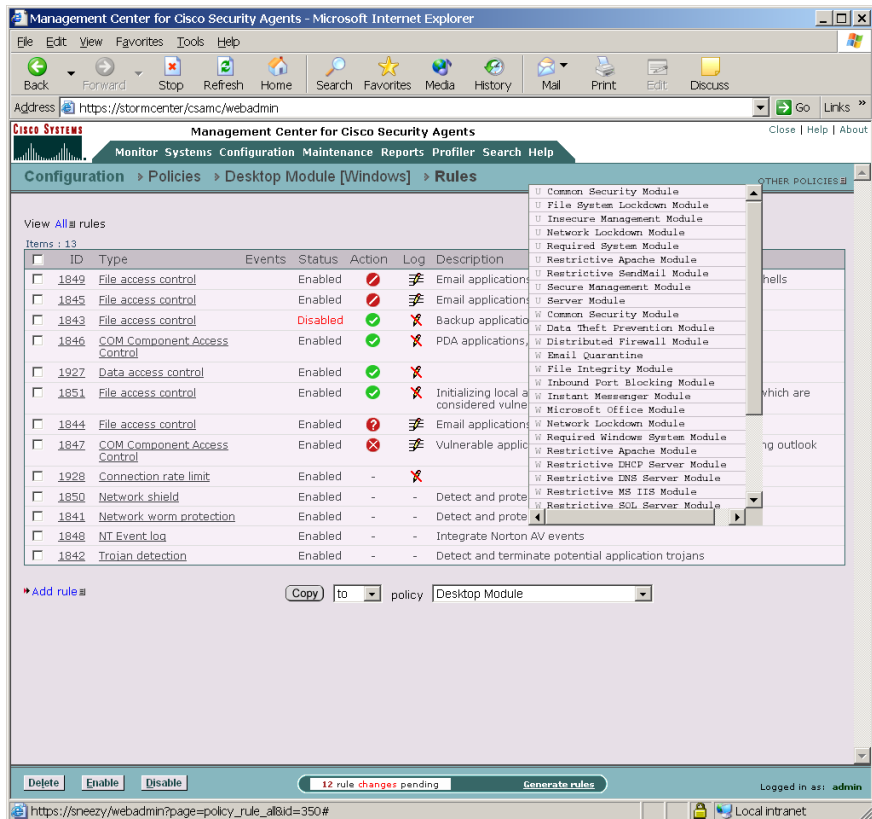


- Other [configuration item]: In configuration views, an Other [Policies, Applications, Files Sets, etc] link appears on the right side of the user interface below the menu bar. Click this link to view a drop-down list containing the names of other configurations within the category you are currently working (see Figure 2-5). Click on one of these names to go to the configuration page for that item.

**Note**

If you jump to another configuration page without saving the page you are currently working in, the information in the current page is lost.

Figure 2-5 Other Policies Link



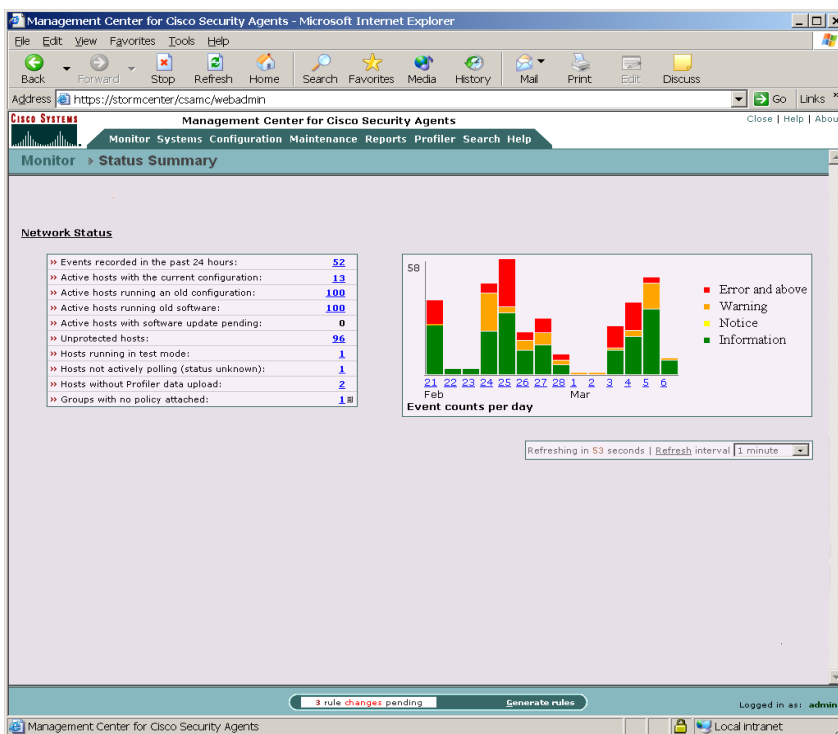
- **Status Summary:** When you first login, the Status Summary view appears (see [Figure 2-6](#)). This page supplies overall system summary information including recorded events and agent rule versions. You can access this page at any time by selecting it from the Monitor category in the menu bar. A *colored graph* on the right side of the page displays the event log according to severity level. Click on a color in the graph to view logged events of that severity level.

**Note**

The colored graph on the Status Summary page is not displayed if there are more than 1,000,000 events recorded within a two week time frame. Once events fall below 1,000,000, the bar graph appears again.

Refer to [Chapter 8, “Event Logging Alerts”](#) for detailed information on the Status Summary view and event logging features.

Figure 2-6 Status Summary View



- **List View:** Each CSA MC configuration category has a top level list view. This list view displays a list of links, each of which represent a configured item for that category. It is from this list view that you create new configurations and delete existing configurations. Buttons for New, Clone and Delete actions are present on list view pages. From the list view, you click on an item link to access the configuration page for that item.
- **Configuration View:** Access the configuration view for an item by clicking on that item in the list view. Configuration views may contain edit fields, radio buttons, checkboxes, and/or listboxes depending on the configuration requirements. Enter the necessary information and click the Save button to store data in the CSA MC database. Configuration views contain Save and Delete buttons. See [Creating, Saving, and Deleting Data, page 2-17](#) for further details.
- **Navigation Tools:** The heading link (below the menu bar, see [Figure 2-7](#)) contains hierarchal links for the item you're configuring. Use these header links to switch between top level list views and subcategory configuration views. For example, in [Figure 2-7](#), the header bar contains links to the top level Policies list view and the MS IIS Server policy. Note that leaving a configuration view without clicking the Save button causes any newly entered data to be lost.
- **Show reference list:** Configuration items that are used in other configurations have a "Show reference list" link on their pages. Clicking this link displays all the configurations where the current item is used. This display also links to the items that are shown.

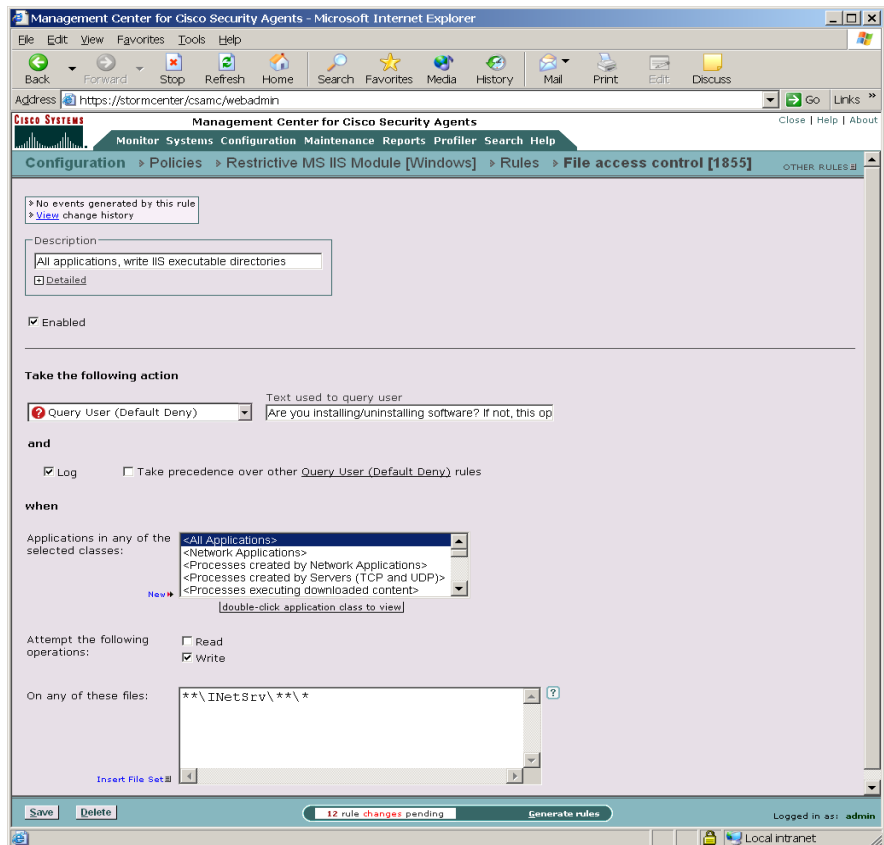
Note that you can click the gray arrow on the right side of the Policy list page to go directly to the rules contained in the policy.

- **Configuration Shortcuts:** Rule pages allow you to insert pre-configured variables such as file sets and COM components into your rules. If there is no pre-configured variable that you wish to use and you want to create a new one, you can do it without leaving the rule page. When you click the Insert link beside any edit box in the rule page, there is a New item in the list that appears first. Selecting the "New" item pops up a configuration page for that variable type. You can then configure a new variable and use it in the rule without having to leave the rule page to access the variable page.

Application classes also have a shortcut you can use to create a new item. Clicking the New link beside the list of application classes in each rule configuration page lets you create a new application for your rule.

Additionally, you can right click your mouse on any page in the server user interface to select from a list of menu items that are also available from the menu bar.

Figure 2-7 Configuration View



Creating, Saving, and Deleting Data

CSA MC Button Frame

All CSA MC action items appear in a frame at the bottom of CSA MC. The buttons in this frame change in accordance with the actions available for the page you're viewing. Available CSA MC buttons and links are as follows.

- **Generate rules (pending changes):** When you are ready to deploy your configuration (policies, rules, variables, etc.) to Cisco Security Agent systems, you must click this link in the bottom frame to view all pending database changes and then to generate them. (See [Chapter 4, "Building Policies."](#))

In most list view pages in CSA MC, there are New, Clone, and Delete buttons (Clone is not present in all list views as you can only clone certain configurations)

- **New:** Use the New button to create a new configuration item within the list view you have selected. Click the New button and a new item appears in the list view. Click the new item link to access the configuration view for that item.
- **Clone:** Use the Clone button in conjunction with the checkboxes beside each list view item. To clone a particular configuration, select its checkbox and click the Clone button. You can clone one item at a time. New links to the cloned configurations appear in the list view.



Note

When you clone an item, such as a policy, that contains variable items like file sets or network services, the cloned rule uses the same variables used in the original rule. The variables themselves are not cloned.

- **Delete:** Use the Delete button in conjunction with the checkboxes beside each list view item. To delete a configuration, select its checkbox (you can select several at once) and click the Delete button. All checked items are deleted. To quickly select all checkboxes, click the very top checkbox in the list view heading bar. Clicking the Delete button then deletes all items.
- **Save:** When you enter configuration information, whether you are entering new data or editing existing data, you must click the Save button once you are finished to save your configuration in the CSA MC database. If you do not click Save before moving to another page in CSA MC, your data is lost.

**Note**

Although your information is stored in the database when you click Save, it is not distributed to the agents across your network until you generate rules. See the [“Generating Rule Programs” section on page 4-97](#) for further information.

- Compare: Policies, Variables, and Application Classes provide a Compare button in their item list views. When you select the checkbox next to 2 items (you cannot compare more than 2 configurations at a time) and click the Compare button, CSA MC displays the configurations side by side and highlights the differences in red. Once you’ve examined how the configurations compare, you can select to merge them.

The purpose of this compare tool is to assist you after you’ve imported configurations or upgraded CSA MC. These processes can cause you to have duplicate or very similar configuration items. Comparing and merging configurations can help you to more easily consolidate duplicate items.

See [Chapter 4, “Building Policies”](#) for details on using Compare to merge configurations.

**Note**

Right-clicking your mouse on a CSA MC page displays a shortcut menu for performing the tasks provided by buttons on that page and for additional configuration tasks not as easily accessible from the current page you’re viewing.

Using the Correct Syntax

CSA MC contains text fields that require you to enter information using a specific syntax. Most of the text fields in these pages are similar and require similar syntax. The text fields are categorized and listed below with the required syntax.

When using configuration variables in rules, application classes, and alerts you must enter the variable name preceded by a dollar sign. The insert links beside each text field automatically insert variables using the correct syntax.

For example, if you have a file set variable named Web Browsers, clicking the Insert File Set link lets you select Web Browsers. It then places \$Web Browsers in the corresponding field using the correct syntax. The dollar sign tells CSA MC that this is a variable value.

When entering a Name for any item you configure, use the following syntax:

Names of items must be unique per operating system. Items may only have the same name if they have different operating system designations. (Host names, however, do not have to be unique.) All names are case insensitive, must start with an alphabetic character, can be up to 64 characters long and can include alphanumeric characters, spaces, hyphens - and underscores _ . (Note one exception, agent kits do not accept spaces in names.)

File entry text boxes require:

In a list of items, each item must appear on a single line. Do not specify multiple items on a single line.

Leading and trailing spaces are removed from each line. Other spaces, such as the one located in "Program Files" are recognized. To indicate leading or trailing spaces you must use special characters. The following special characters are recognized. (Note that the need to use the characters listed below should occur very rarely.)

- 'b Leading/Trailing Space
- 't Tab
- 'n Line feed
- 'r Carriage return



Note

If you want to use a single quote(') in a file name, you must enter two single quotes (') for CSA MC to recognize the syntax correctly. Two single quotes are seen as one quote.

Local system files are entered using full path and disk drive.

Windows:

```
c:\Program Files\Outlook\msimn.exe  
c:\winnt\regedit.exe
```

You can also use **@fixed** to indicate all local system drives without having to indicate the drive letters.

For example, @fixed:\Program Files\Outlook\msimn.exe

UNIX:

```
/etc/passwd
```

Local system files are entered using full path and disk drive (Windows) with optional wildcard notations.

Windows

```
c:\Program Files\Outlook\*.exe
```

UNIX

```
/usr/bin/*
```



Note

Windows peripherals, such as floppy and CD drives, can be referenced by their drive letter.

Use the wildcard notation (* or ?) to indicate files within directories and whether directories and subdirectories are recursive.

Table 2-1 Wildcard Operators

Example	Translation
*	One wildcard entry indicates a single directory level or all files in a specified directory.
**	Two wildcards entered in this manner indicate a recursive directory path (including all directories, passing down as many levels as exist in the path).
?	Use the question mark wildcard to represent a single character. For example, ????.doc. This indicates a file name containing only three characters with a .doc extension.

For example:

The following entry indicates all files one directory level down in the winnt directory. It does not include files in the winnt directory itself.

```
c:\winnt\*\*
```

The following entry indicates all files in the winnt directory and all subfolders recursively (digging down as deep as necessary) and files.

```
c:\winnt\**\*
```

The following entry indicates all files in the winnt directory and all subfolders recursively (digging down as deep as necessary) and files that contain exactly two characters in their name and have any extension.

```
c:\winnt\**\??.*
```

If you do not specify a drive path in a file text field, CSA MC always appends the variable **\ to the named file. For example, if you enter foo.doc into a text field, it is saved as **\foo.doc.



Note

You can use the same wildcard notations for indicating UNIX files and directories.

**Caution**

File access control rules apply to files, not directories. You must make some file specification when you are entering literal paths. A wildcard is acceptable to specify all files in a named directory. CSA MC always assumes the last entry in a literal directory path (not a variable) is a file.

You can use the following "short hand" entries in File Sets and in File access control and File monitor rules to indicate common system directories. The @symbol must appear at the start of the short hand name. These entries resolve to the Windows directory on each agent system.

Table 2-2 Sample Short Hand File Tokens (Windows only)

Example	Translation
@windows	Use @windows to indicate the directory pointed to by the %SystemRoot% environment variable When using @windows, for example, in the File access control rule Files field, it is interpreted as @windows* to indicate the files within the directory.
@system	Use @system to indicate %SystemRoot%system32
@program_files	Use @program_files to indicate %SystemDrive%Program Files

Table 2-2 Sample Short Hand File Tokens (Windows only) (continued)

Example	Translation
@(regpath Registry key/value pair default=default directory)	Use @(regpath Registry key/value pair default=default directory) to localize the directory structure of an application or other resource. This is useful to indicate software regardless of the directory to which it has been installed. Note that the default= field is optional but recommended, For example: <pre> @(regpath HKLM\CCS\foo\instdir default=**\Program Files\foo\bin) </pre>
@dynamic	Use @dynamic in the File set text field to indicate all files that have been quarantined by CSA MC as a result of correlated email worm events or correlated virus scanner log messages. Files are quarantined by CSA MC for up to one hour. This list updates automatically (dynamically) as logged quarantined files are received. To view the files that are added to the dynamically quarantined files list, click the numbered link beside Quarantined File Events . This link is located beside the last checkbox on the Global Event Correlation page. It takes you to the pertinent event log messages. Read the messages there to locate the names of quarantined files. See the “Global Events” section on page 5-18 .

You can use the following "short hand" entries in File Sets and in File access control and File monitor rules to indicate removable media. The @ symbol must appear at the start of the name. These entries resolve as follows:

Table 2-3 Removable Media Token Syntax (Windows only)

Example	Translation
@removable	This indicates all removable media. That includes, floppies, CDs, zip drives, etc. Note that if you want to indicate all removable media except floppies, for example, you'd have to configure a file set that explicitly excludes floppies from all removable media.
@floppy	This indicates all floppy drives. You can specify particular file paths on floppy media using the following syntax: @floppy:\<specify wildcards or paths>. Note that @floppy:\ means only the top level files on the floppy media. @floppy or @floppy:** means all files on the floppy media.
@CD	This indicates all CD-ROM drives(including DVD). You can specify particular file paths on CD media using the following syntax: @CD:\<specify wildcards or paths>. Note that @CD:\ means only the top level files on the floppy media. @CD or @CD:** means all files on the floppy media.

**Note**

USB connected drives are removable media.

UNIX Directory Path Protection

You can protect directory paths as well as files on UNIX systems. (On Windows platforms, File access control rules apply to files, not directories. You cannot protect Windows directory structures.)

When protecting UNIX directory paths, you should use a file literal rather than a file set variable. Although, you can use a file set, the configuration is more cumbersome. In File access control rules applied to UNIX systems, you can control writes to directory structures, but the syntax required to do this does not include the protection of directory content. That must be entered separately.

In the following examples, `/usr/admn/sg/` is the directory and `x`, `y`, and `z` are files in the `sg` directory.

The following entry protects files `x`, `y`, and `z` in the `sg` directory. It does not protect the directory structure.

```
/usr/admn/sg/*
```

The following entry protects the `usr` directory. (No files in the directory are protected with this entry alone. The same applies to the subsequent directory entries in this example).

```
/usr
```

The following entry protects the `adm` directory.

```
/usr/admn
```

The following entry protects the `sg` directory.

```
/usr/admn/sg
```

In order to protect the entire `usr/admn/sg` directory structure and files `x`, `y`, and `z`, you must enter all four examples above in the file access control files field (each entry on a separate line in the files edit field).

**Caution**

Symbolic Links—For UNIX, if you create a File access control rule to protect a symbolic link, **ONLY** that symbolic link is protected. The underlying resource, unless also specified, is **NOT** protected. For example, a File access control rule written for `/etc/hosts` does not protect `/etc/inet/hosts`. Similarly, a File access control rule written for `/etc/inet/hosts` does not protect `/etc/hosts`. If you want to protect a symbolic link and its underlying resource, both must be specified in the rule.

Network system paths (Windows only) entered using the following syntax:

```
\\<machine name>\<share>\<path>\<filename>  
\\Backup_Server\finance\records\database.db
```

You can also use **@network** to indicate all network shares. For example,
`@network:\finance\records\database.db`

**Caution**

Do NOT enter a drive letter for network share paths.

Network address text boxes require addresses entered in any of the following formats:

Enter single fully qualified addresses.

a.b.c.d

Enter address ranges.

a.b.c.d-y.z

This address range indicates all addresses from a.b.c.d-a.b.y.z

You can use the following "short hand" entry in Network Address Sets and in Network Access Control rules to indicate all local addresses on the agent system in question. The @ symbol must appear at the start of the short hand name.

Use **@local** to indicate all local addresses on the agent system.

**Caution**

On UNIX platforms, IPV6 addresses are not officially supported; however, an IPV6 connection will work as the applied rules dictate if the address in question is covered by the "all" addresses range (0.0.0.0-255.255.255.255 includes IPV6 addresses) or by @local. Local addresses on the agent system (indicated by @local) also include IPV6 addresses.

See the [“Network Access Control” section on page 4-49](#) for more information.

Network service text boxes require protocols and port ranges entered in the following formats:

Format all entries as:

```
protocol/port or port range
TCP/80
UDP/53
TCP/1024-65535
```

The protocol here is either "TCP" or "UDP".

Port ranges are designated in the range 0-65535.

Designating ephemeral ports—In some cases, an application may want to offer a temporary service port for callback data connections. An ephemeral port is a temporary system-assigned port for this purpose.

For example, an ephemeral port would be the likely data connection for active FTP. If you do not specify an ephemeral port range for accepting an active FTP connection, you would have to allow clients to listen on a wide range of ports to accept this connection type. This would unnecessarily open a wide range of data channels and possibly create a vulnerability that could be exploited by a Trojan.

You can specify an ephemeral port range for a Network service as follows:

```
TCP/ephemeral
UDP/ephemeral
```

**Note**

It only makes sense to use ephemeral ports on systems accepting connections. Also note that Deny log messages triggered by a rule using an ephemeral port range appear in the event log containing the real port number.

**Caution**

Ports that are ephemeral allocated are only matched against an explicit ephemeral class. Ephemeral ports are treated as "port 0" for rule comparisons. For example, ephemeral port 2000 matches port 0, not port 2000.



Configuring Groups and Managing Hosts

Overview

The system hosts across your network, including mobile systems in the field, must download Cisco Security Agent software and register with the Management Center for Cisco Security Agents to receive the security policies configured for them. When you are ready to apply policies to the hosts running agents, having those hosts placed into common groups streamlines the process of assigning policies to several hosts at once. In order to place hosts into groups, you must first analyze the security needs of each host system and map out a security plan. Hosts with similar requirements can then be grouped together.

Management Center for Cisco Security Agents ships with several pre-configured groups you can use. If the included groups do not suit your needs, use the instructions in this chapter to configure new groups or to edit existing ones.

This section contains the following topics.

- [Grouping Hosts Together, page 3-2](#)
- [Configuring Groups, page 3-3](#)
- [Building Agent Kits, page 3-8](#)
- [Viewing Host Status, page 3-20](#)
- [Distributing Software Updates, page 3-26](#)

Grouping Hosts Together

Host groups reduce the administrative burden of managing a large number of agents. All hosts across your network, including mobile systems in the field, must exist as registered host entries in the Management Center for Cisco Security Agents for policy configurations to be assigned to them.

Grouping individual host systems together provides the following advantages:

- It lets you consistently apply the same set of policies across multiple host systems.
- It lets you apply Alert mechanisms and Event Set parameters based on group configurations.
- It lets you use Test Mode to try out policies on groups of hosts before you actively enforce those policies.

You can group hosts together based on any criteria that best fits your enterprise. For example:

- Group hosts according to system function, such as Web servers. Then you would create a policy that corresponds specifically to the needs of your Web servers and distribute it to that group.
- Group hosts according to business groups, such as finance, operations, and marketing. Distribute policies based on each business group's individual needs.
- Group hosts according to geographical or topological location. For example, group hosts based on their subnet designation for reporting purposes.
- Group hosts according to their importance to your organization. Place mission-critical systems into a common group to apply critical alert level configurations to them.

**Note**

Hosts may belong to multiple groups and automatically receive policies that are attached to every group to which they belong. You can add or remove hosts from a group at any time. However, the policy configuration of a host that is moved to another group will not take affect until you generate your rule programs and distribute them.

Configuring Groups

Host groups reduce the administrative burden of managing a large number of agents. Grouping hosts together also lets you apply the same policy to a number of hosts. A group is the only element required to build agent kits.

You do not configure hosts with CSA MC as you do other CSA MC elements. When hosts across your network download and install agent kits, they automatically and transparently register with CSA MC. Hosts inherit membership to the groups that were associated with the agent kit they installed. Successfully registered hosts appear in a linked list when you select Hosts from the Systems category in the menu bar. At registration time, hosts are also automatically put into their assigned group. You can change host groupings at any time.

**Note**

Management Center for Cisco Security Agents ships with preconfigured groups you can use if they meet your initial needs. If you use a preconfigured group, you do not have to create your own group as detailed in the following pages.

To configure a group, do the following.

-
- Step 1** Move the mouse over **Systems** in the menu bar and select **Groups** from the drop-down list that appears. The list of existing Groups is displayed in the right pane. Management Center for Cisco Security Agents ships with several pre-configured groups.
- Step 2** Click the **New** button to create a new group entry. (This group is empty until hosts install agents and register.)

**Note**

If you have "All" designated as the operating system type for your administrator session, you are prompted to select whether this is a Windows or a UNIX group. See the [“Administration by Operating System” section on page 2-5](#) for details. (You cannot combine UNIX and Windows hosts in the same group.)

Step 3 In the available group fields, enter the following information:

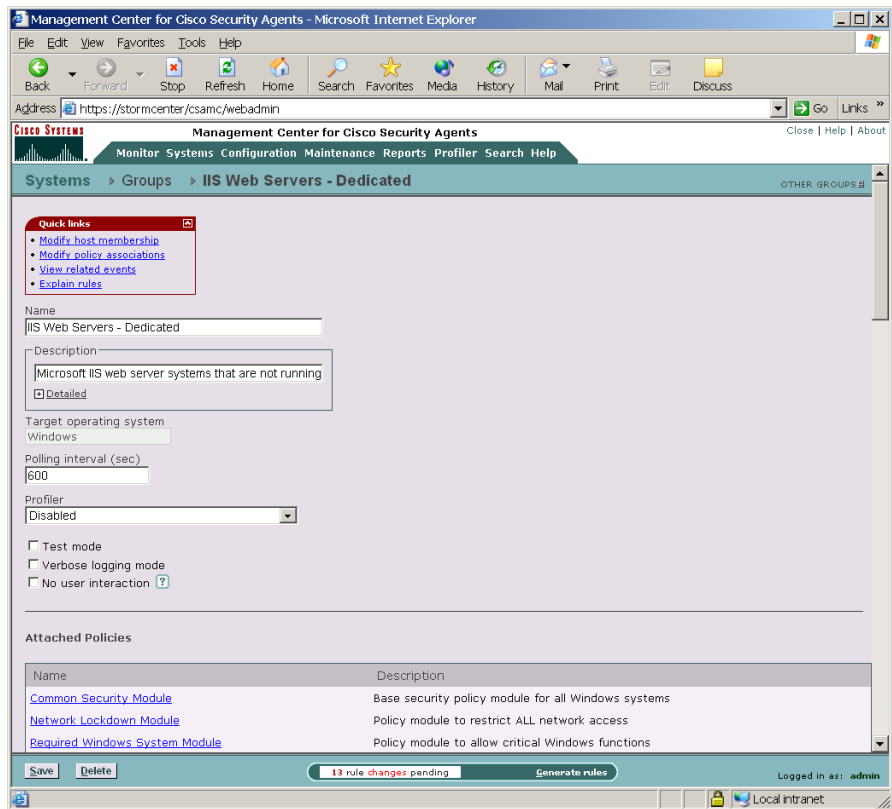
- **Name** This is a unique name for this group of hosts. Names are case insensitive, must start with an alphabetic character, can be up to 64 characters long and can include alphanumeric characters, spaces, hyphens -, and underscores _ . Generally, it's a good idea to adopt a naming convention that lets you quickly recognize groups in the CSA MC group list view.
- **Description** This is a useful line of text that is displayed in the list view and helps you to identify this particular group. Optionally, expand the **+Detailed** field to enter a longer description.



Tip

You can use the Tab key to navigate between edit fields.

Figure 3-1 Group Configuration Page



Step 4 Optionally, you can select the **Test Mode** checkbox for this group.



Caution

In Test Mode, the Cisco Security Agent will not deny any action even if an associated policy says it should be denied. Instead, the agent will allow the action but log an event (if logging is selected for the rule). This helps you to understand the impact of deploying a policy on a host before enforcing it. For further information, see [Chapter 4, “Building Policies.”](#)

Step 5 Optionally, enable **Verbose Logging Mode** to change the event log timer to log all reoccurring events rather than suppressing duplicates. See [Chapter 8, “Event Logging Alerts”](#) for more information on the event log.

- Step 6** Optionally, enable **No user interaction** (available on Windows groups only) to have no agent user interface or query pop-ups appear on end user systems. You may wish to do this if you do not want end users to have any interactions with CSA MC using a local agent UI (i.e. clearing the cache, polling, and self-protection and rule queries).

See [page 3-7](#) for more details.

**Caution**

When there is no agent UI present, there are no query user pop-up boxes displayed. The default is immediately taken on all query user rules and heuristics that are present in the assigned policies. This means that the default of Query User Allow or Query User Deny is taken on all query user "access control rules" and the default of Terminate or No is taken on all heuristics (Trojan detection, Network worm, etc.) unless specific application class exceptions are made for heuristic rules.

If an end user system already has an agent UI installed, when you select this No user interaction checkbox for the agent's group and generate rules, the agent UI disappears when the new rules are downloaded.

**Note**

To fully restrict end users from interacting with the agent, you could use the "No user interaction" capability in combination with the Agent service control rule (see the [“Agent Service Control”](#) section on page 4-24) and the Quiet software update capability (see the [“Building Agent Kits”](#) section on page 3-8).

- Step 7** Optionally, you can change the default **Polling interval** from 600 seconds (10 minutes) to any value between 10 seconds and 86400 seconds. This controls how often agents in this group poll into CSA MC for policy updates. Shortening the polling time can be useful when you are trying out new policies. Otherwise, the default value is recommended. (If you have the same hosts in multiple groups, the group containing the shortest polling interval setting takes precedence for the hosts in question.)



Note If you change a group's polling interval, that new interval time will not take effect until the host polls in again for new rules. Therefore, it may take as long as the previous polling interval setting before hosts begin polling in using the new setting.

Step 8 When all required information is entered, click the **Save** button to enter and save your group in the CSA MC database.

Once you attach (associate) policies to specific groups, the configuration view for the group displays a table listing all the rules, in order of precedence, that are applied to that group. From this table, you can navigate to those rules and policies.

No User Interaction Feature

Enabling the "No user interaction" checkbox for a group has the following affects.

Software updates

- **Not automatic:** Pop-up box prompts still appear to prompt user to install update. (This box warns that after install completes, system automatically reboots after 2 minutes. User cannot stop this reboot once installation begins.) User must click the OK button in the pop-up box to begin update. Pop-up box remains on screen until user performs the update. When install is complete, a 2 minute automatic reboot warning message appears.
- **Automatic:** User is not prompted before update installation begins. When the install is complete, a 2 minute automatic reboot warning message appears. User cannot stop this reboot and has 2 minutes to save any open documents. Regardless if the user is present or not, if the machine is running, both the install and the automatic reboot take place.

See [Table 3-2 on page 3-32](#) for more Software update details.

Queries

- When there is no agent UI present, there are no query user pop-up boxes displayed. The default is immediately taken on all query user rules and heuristics that are present in the assigned policies. This means that the default of *Allow* or *Deny* is taken on all query user "access control rules" and the default of *Terminate* or *No* is taken on all heuristics (Trojan detection, Network worm, etc.) unless specific application class exceptions are made for heuristic rules.

Unavailable end user features

- No messages to inform user that actions have been denied and why.
- No ability to clear cache or re-enable logging.
- No fast polling ability.
- No end user contact information can be sent to CSA MC.

No user interaction feature notes

- If a host belongs to multiple groups, having a visible agent UI setting, if present in any group for which the host is a member, takes precedence over a no user interaction agent UI setting.
- Whether or not an end user system is going to have a visible agent UI or a hidden one, the end user (or administrator) must download and install the agent kit on the system. The initial installation of an agent kit cannot be done automatically (unless you have written your own script to do so, see [Scripted Agent Installs and Uninstalls, page 3-19](#)).

Building Agent Kits

CSA MC allows for the creation of custom Cisco Security Agent installation kits that greatly reduce the administrative burden of deploying Cisco Security Agents on new systems.

At the time of creation of the Cisco Security Agent kit, it may optionally be associated with one or more groups. The particular agent kit a host installs determines what group(s) it is initially placed into. You can create as many kits as necessary to distribute your policies to targeted hosts.

After a kit is installed on a host, the agent running on that host registers itself with CSA MC. CSA MC then automatically places the host in the groups that were associated with the installed kit.

**Note**

CSA MC ships with preconfigured Agent kits you can use if they meet your initial needs. There are prebuilt kits for desktops, servers, and many more. These kits place hosts in the corresponding groups and enforce the associated policies of each group. (If you use a preconfigured agent kit, you do not have to build your own kit as detailed in the following pages.)

To create agent kits, do the following.

Step 1 Move the mouse over **Maintenance** in the menu bar and select **Agent Kits** from the drop-down menu that appears. Existing agent kits are displayed.

Step 2 Click the **New** button to create a new agent kit.

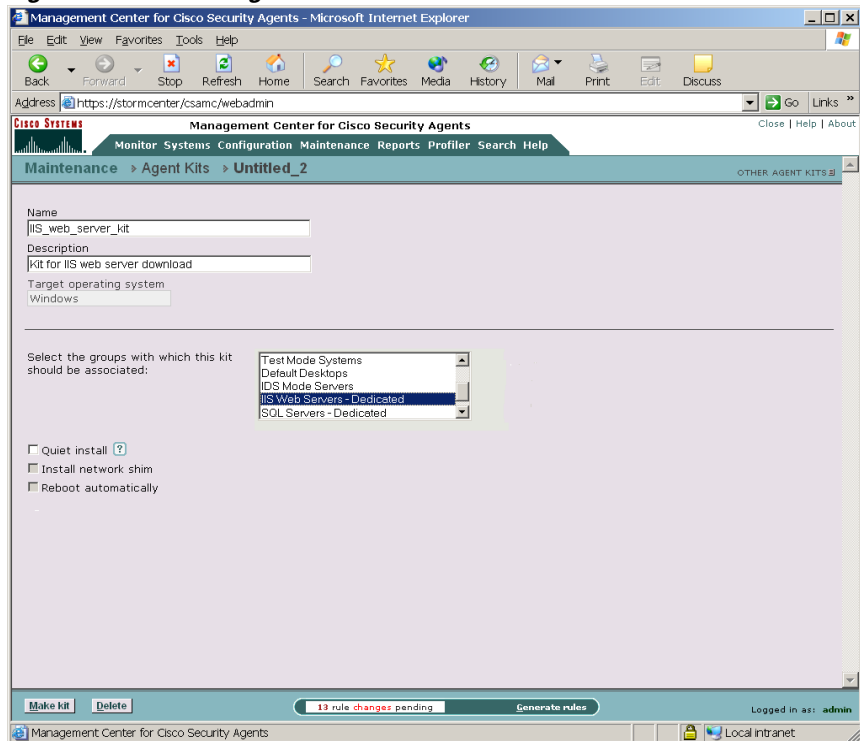
**Note**

If you have "All" designated as the operating system type for your administrator session, you are prompted to select whether this is a Windows or a UNIX group. See the [“Administration by Operating System” section on page 2-5](#) for details. (You cannot select a UNIX group for an agent kit that you have configured for Windows systems.)

Step 3 In the agent kit configuration view (see [Figure 3-2](#)), enter a **Name** for this kit. This is a unique name (Agent kit names are not allowed to have spaces). Generally, it's a good idea to adopt a naming convention that lets you and the systems that will be downloading the kit, recognize it easily.

Step 4 Enter a **Description**. This is an optional line of text that is displayed in the agent kit list view and helps you to identify this particular kit.

Figure 3-2 Create Agent Kit



- Step 5** From the available list box, select the group or groups that will download and install this kit. To select multiple items in a list box, hold down the **Ctrl** key as you select each item. To unselect a single item, hold down the **Ctrl** key when you click on the item in question. Press and hold the **Shift** key when you click on an item to select multiple successive items.
- Step 6** Select whether or not to have agents install "quietly" on end-user systems (Windows only). A **Quiet install** requires users to download the self-extracting executable as does the non-quiet install. The difference is, no prompts appear and the user is not required to enter any information or select any options. A non-quiet install prompts the user for installation options, such as enabling the network shim, in addition to the reboot prompt. See [on page 3-15](#) for details on Agent kit creation combination options.

- Step 7** For Windows agent kits, if you select Quiet install, you can also select whether the **Network shim** is installed or not during the installation.

**Caution**

In some circumstances, you may not want users to enable the network shim on their systems as part of the agent installation. For example, if users have VPN software or a personal firewall installed on their systems, the network shim's Portscan detection, SYN flood protection, and malformed packet detection capabilities may not be needed. To allow users to enable it, you would create kits as non-quiet installations. (Do not select the Quiet install checkbox.) This way, users are prompted to enable the network shim during the agent installation. For more information, see the “[Network Shim Optional](#)” section on page A-4. This chapter (Appendix A) also provides information on the agent user interface.

**Note**

Not enabling the network shim does not mean that Network access control rules won't work. It only means that the system hardening features (configured in the Network Shield rule page) mentioned in the previous paragraph are not enabled.

- Step 8** If you select Quiet install, you can also select whether the system is **automatically rebooted** once the install is complete. (Even if an end user is present when the installation is finished, this reboot cannot be stopped.)

**Note**

In some cases, you may not want a system to reboot after the installation completes. If a reboot does not occur after the agent installation, partial security is enforced immediately. Full security is enforced after the first reboot. See [Agent Install Complete Prompt for Automatic Reboot](#), page 3-16 for details. (Note that Windows NT systems must be rebooted after an agent installation.)

Step 9 Click the **Make Kit** button.

Once you click the Make Kit button, CSA MC produces a bundled kit for distribution. It displays a URL for this particular kit (see [Figure 3-3](#)). You may distribute this URL, via email for example, to the host systems the kit is designated for. They access the URL to download and then install the kit. This is the recommended method of agent kit distribution.

But you may also point users to a URL for the CiscoWorks system. This URL will allow them to see all kits that are available. That URL is:

```
https://<ciscoworks system name>/csamc/kits
```

If you are pointing users to the “kits” URL and you have multiple agent kits listed here, be sure to tell users which kits to download.

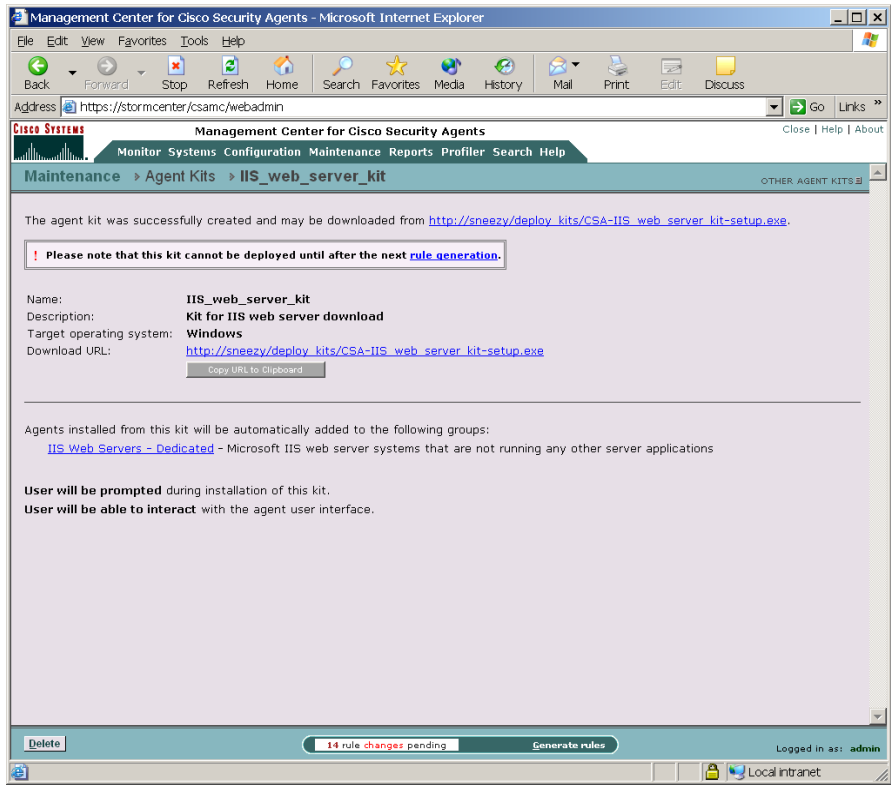


Note Note that the Registration Control feature also applies to the <ciscoworks system name>/csamc/kits URL. If the Registration Control feature (see [Registration Control, page 3-19](#) for details on the feature) prevents your IP address from registering, it also prevents you from viewing this kits URL.



Note You must regenerate your rule program after agent kits are created. See [Agent Kit Status, page 3-14](#) for details on when a kit is ready for download.

Figure 3-3 Agent Kit Download URL

**Note**

If you installed Management Center for Cisco Security Agents to the default directory, all agent kits are placed in the %Program Files%\CSCOpX\CSAMC\bin\webserver\htdocs\deploy_kits directory.


Agent Kit Status

When you create an agent kit, it is given one of three status levels based on how far into the configuration you've progressed. Those status levels are as follows:

- **Ready:** This means the agent kit is ready for download to host systems.
- **Needs rule generation:** This means that all agent kit configuration parameters are complete, but you must generate rules before the kit can be downloaded.
- **Incomplete:** This means that you have not configured all the necessary parameters for this agent kit. You must complete the configuration and then generate rules before the kit can be downloaded.

Table 3-1 Agent Kit Creation Combinations

Windows Options/ Installation Types	Non-Quiet Installation No configuration required at kit creation time. End user receives various prompts during install.	Quiet Installation Select the Quiet install checkbox at kit creation time. End user receives no prompts during install.
Network Shim	No administrator configuration. The end user is prompted to install or not install the network shim.	Administrator decides if the network shim is installed by selecting both the Quiet install and Network shim checkboxes at kit creation time.
Automatic Reboot	<p>This feature is not available for a non-quiet installation. The user is prompted to make a reboot decision at the end of the installation. See Figure 3-4.</p> <p>*If Yes is selected to reboot, full security is enforced after the reboot.</p> <p>*If No is selected to not reboot, partial security is enforced immediately.</p> <p>(In all cases, full security is enforced after the next reboot.)</p>	<p>Administrator decides if an automatic reboot will occur by selecting both the Quiet install and Automatic reboot checkboxes at kit creation time.</p> <p>*If Automatic reboot is selected, the system will reboot when the installation is complete. A 2 minute warning is given before the reboot occurs. This reboot cannot be stopped. See Figure 3-5.</p> <p>*If Automatic reboot is not selected, partial security is enforced immediately when the installation completes. No reboot occurs. An "Installation complete" message appears. (This message disappears after 10 seconds.) Full security is enforced the next time the user chooses to reboot the system.</p> <p>(In all cases, full security is enforced after the next reboot.)</p>

Figure 3-4 Agent Install Complete Prompt for Optional Not-Automatic Reboot**Figure 3-5** Agent Install Complete Prompt for Automatic Reboot

Agent Reboot vs. No Reboot

If a system is not rebooted following the Cisco Security Agent installation, the following functionality is not immediately available. (This functionality becomes available the next time the system is rebooted.)

Windows agents

- Network Shield rules are not applied until the system is rebooted.
- Buffer overflow protection (located on the Trojan page for Windows) is only enforced for new processes.
- Data access control rules are not applied until the web server service is restarted.
- COM component access control rules are not applied until the system is rebooted.

UNIX agents, when no reboot occurs after install, the following caveats exist:

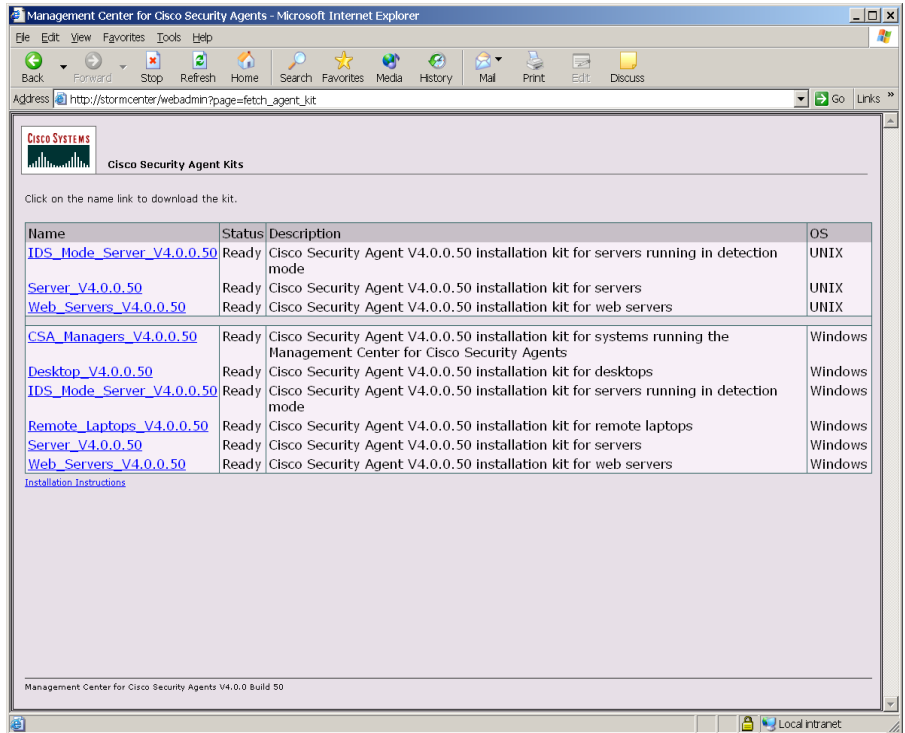
- Buffer overflow protection is only enforced for new processes
- Network access control rules only apply to new socket connections
- File access control rules only apply to newly opened files.



Caution

Windows NT systems must be rebooted after the agent installation completes. Windows NT systems will not receive a reboot optional prompt at the end of an agent installation (even if that option is part of the agent kit installation).

Figure 3-6 Download Agent Kits



Agent Registration

When an agent kit is ready for distribution, you can notify end users to download and install the kit from the URL produced by CSA MC when the kit is made. Once the kit installation is complete, each individual host's agent automatically and transparently registers with CSA MC.



Note

Each kit is created for particular groups based on the policies that will be attached to those groups. Policies are described in [Chapter 4, "Building Policies."](#)

Scripted Agent Installs and Uninstalls

You can use scripts to silently install and uninstall Windows Cisco Security Agents on end user systems.

- **Scripted install:** The agent kit is a self-extracting executable placed in the following directory on the server: `%Program Files%\CSCOpX\CSAMC\bin\webserver\htdocs\deploy_kits`. (Retrieve the kit from this directory or download it from the server.) You can then use a script to copy and silently install agent kits on systems. Note that you must select the **Quiet install** checkbox when you build the kit if you are planning to install it via a script.
- **Scripted uninstall:** The agent installation places a bat file in the system32 directory. Administrators may use a script to remotely and silently uninstall the agent by invoking the `CSA_uninstall.bat` file in the system32 directory. You must also pass a parameter to the file for the agent to uninstall silently regardless of whether the original agent kit was a Quiet install. Enter the following: `CSA_uninstall.bat 3`

**Note**

Before silently uninstalling the agent via a script, you must disable any agent service control rules that deny or query administrators before stopping the agent service.

Registration Control

This feature is accessible from the Maintenance item in the menu bar. Access the Registration Control page to enter a range of addresses which restricts agent hosts attempting to successfully register with CSA MC to those with addresses listed here.

This feature prevents unauthorized hosts from downloading agent kits and receiving rules. (Note that any user who is logged in to CSA MC, can download a kit.)

The default entry here is <all> (0.0.0.0-255.255.255.255) which applies no address registration restrictions. An example entry of restricted registration addresses is as follows. (Only those addresses within the range listed can register. This range is inclusive):

```
192.168.10.0-192.168.10.255
172.16.20.0-172.16.20.255
```

Viewing Host Status

You can see which hosts have successfully registered by moving your mouse over Systems in the menu bar and clicking Hosts in the drop-down list. This takes you to the Hosts list page. Use the pulldown options on the right side of the window to view an abbreviated host status in the following categories.

- **Active** : A host is active if it polls into the management server at regular intervals. When you select this viewing option, a "Yes" for Active or a "No" for Not Active appears in the column.

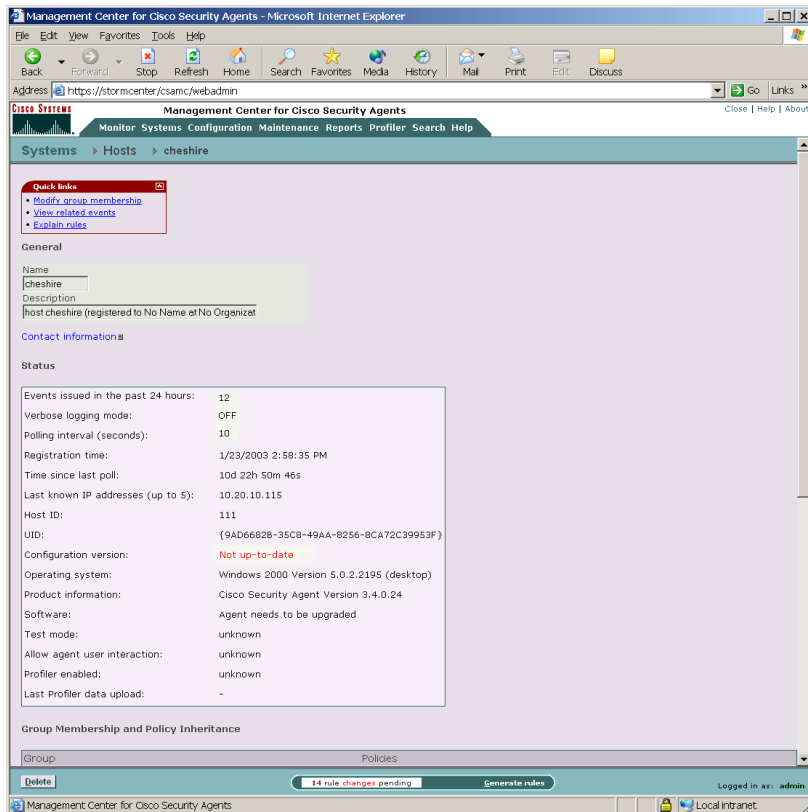
Note that a "Not active host" is a host that has missed three polling intervals or has not polled into the server for at least one hour.

- **Protected**: When you select this viewing option, a "Yes" for Protected or a "No" for Not Protected appears in the column. A system is not protected if it does not belong to a group or if it belongs to a group that has no policies attached.
- **Latest software**: When you select this viewing option, a "Yes" for Latest Software or a "No" for Not Latest Software appears in the column. If an agent is not running the latest software, you will want to deploy a software update.
- **Test Mode**: When you select this viewing option, a "Yes" for running in Test Mode or a "No" for Not Running in Test Mode appears in the column.
- **Last Poll**: When you select this viewing option, the time and date of the most recent poll for the host is displayed.

Note that by default, agents poll the management server every 10 minutes for updated policies.

Click on the host link itself for detailed host information. In the Host Detail page (see [Figure 3-7](#)), the additional options and information are available.

Figure 3-7 Host Detail View



- Click the **Modify group membership** link in the Host detail page (see Figure 3-7) to add or remove this host from a group.
- CSA MC provides an explanation, in paragraph form, of the policies attached to each host. Clicking the **Explain rules** link takes you to this paragraph explanation.
- Once hosts are registered, they automatically receive policies from CSA MC.

When host agents register with CSA MC, the CSA MC database receives the following information on each host.

- **Name and Description:** These fields are populated with information received from the agent system when it registers. This is the name that identifies this host system on the network. This name does *not* have to be unique. CSA MC assigns each registering host a unique ID number by which the database identifies it.
- **Contact Information:** Click this link to view the contact information provided to the agent by the user. (The available fields for the user are: first name, last name, email, telephone, location.)
- **Events issued in the past 24 hours:** This is the number of events (rule triggers) that have occurred on the host system in the given time frame.
- **Verbose logging mode:** This field can read as either OFF or ON, indicating whether this feature is enabled for this host. This feature is configurable through the Groups page.
- **Polling interval (seconds):** The value shown here indicates the time interval in which this system polls in to CSA MC. This feature is configurable through the Groups page.
- **Registration time:** This is the time that the agent registered with CSA MC.
- **Time since last poll:** This is the interval since the host system's last polling request.
- **Last known IP address:** This is the IP address of the host. If DHCP addressing is used, this is the last known address of the host. (Up to 5 IP addresses can be listed.)
- **Host ID:** CSA MC assigns each registering host a unique ID number by which the database identifies it.
- **UID:** This is a globally unique ID for your agent. It is obtained from the agent kit. Different kits present different IDs. Every host that installs a particular kit will have the same registration ID. Once registered, however, each host receives a unique global ID.
- **Configuration version:** This field reads Up-to-date or Not up-to-date, indicating whether the agent has the latest policy configuration from CSA MC.

Note that by default, agents poll to CSA MC every 10 minutes for updated policies.

- **Operating System:** This is the operating system installed on this particular machine.
- **Product Information:** This is the agent version for this particular machine.
- **Software:** This is the version of agent software the system is running. If there is a software update available for this host, this field provides that information. If an update for a host is scheduled but not yet installed, this field provides that information as well. See [Configuring Scheduled Software Updates, page 3-28](#) for further information.
- **Test Mode:** If this host is part of a group operating in "test mode," that information is displayed here. For more information on test mode, see the ["Using Test Mode" section on page 4-95](#).
- **Allow agent user interaction:** This indicates whether the end user has an agent UI.
- **Profiler enabled:** This item appears if Profiler is enabled on the end user system.
- **Last Profiler data upload:** If Profiler is enabled on the end user system, this indicates the time of the most recent upload of analysis logging data.

Optionally, you can enter contact information such as user name, location, email, and telephone number for each host system. If an agent is generating alerts, having this contact information readily available could expedite troubleshooting measures.

The host view also displays a table listing all the rules and policies that are applied to that host. From this table, you can link to those rules and policies.

Changing Host Group Assignments

When a host registers with CSA MC, it is automatically placed into the group(s) you designate for it. There is no need to add a host to a group initially. You only need to add hosts to groups when you are changing their group designation after they've registered.

Hosts may belong to multiple groups and receive policies that are attached to every group to which they belong.

**Caution**

You can add or remove hosts from a group at any time. If you do change host group assignments, the policy configuration of a host that has been moved to another group will not take affect until you generate your rule programs and distribute them.

**Note**

See [Viewing Host Status, page 3-20](#) for details on hosts.

There are several ways to add a host to a group.

- To add a host to multiple groups, use the Hosts>Modify group membership link.
- To add multiple hosts to a single group, use the Group>Modify host membership link.
- To move or copy all hosts in one group to another group, use the Bulk Transfer feature accessible from the Group>Modify host membership link.

To add one or more hosts to a single group, do the following.

-
- Step 1** Add hosts to a particular group by accessing that group's edit view. Move the mouse over **Systems** in the menu bar and select **Groups** from the drop-down menu that appears.
- Step 2** From the group list view, click the link for the group to which you want to add hosts. This brings you to that group's edit view.
- Step 3** From the edit view, click the **Modify host membership** link. This takes you to a swap box page containing a list of host systems (if any) in the right box that are in this group. Hosts listed in the left box are not in the group.
- Step 4** To add a host to this group, select the host in the left box and click the **Add** button to move it to the right box (see [Figure 3-8](#)). It is now a part of the group.

To select multiple nonsuccessive items in a swap box, hold down the **Ctrl** key as you select each item. To unselect a single item, hold down the Ctrl key while you click on the item in question. Click **Select all** to select all items in the swap box. When you click the Add button, all selected items are added.

To remove a host from a group, from the same view referred to in step 4, select the host in the right swap box that you wish to remove. Click the **Remove** button. The moves to the left (unattached) box.

Bulk Transfer

Use the bulk transfer feature to easily move or copy all hosts from the group you select in the available pulldown field, into the Group you are currently viewing. When you click the OK button beside the group selection pulldown, all hosts are moved or copied.

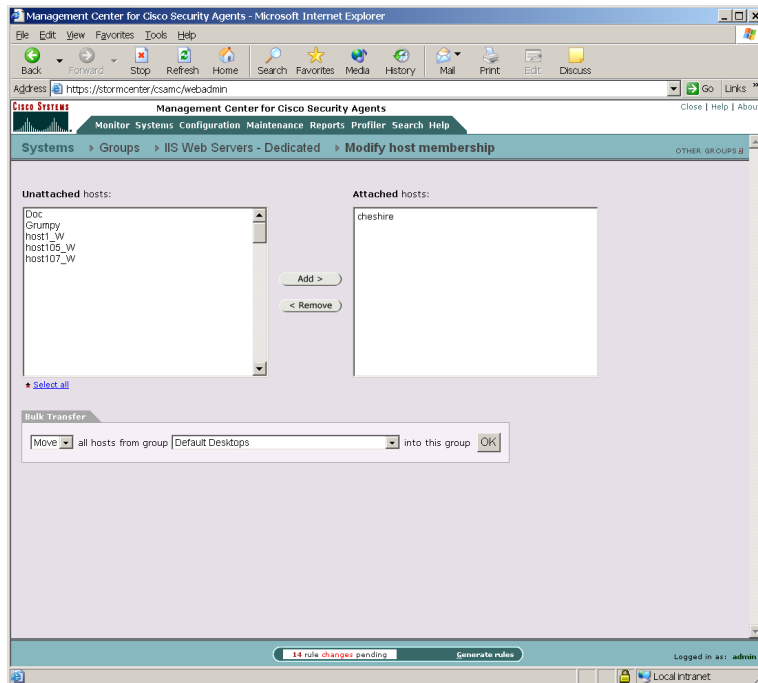
When you next click the **Generate** button, policies associated with this group will no longer be applied to the removed hosts. (The host is not deleted from the database, it is just no longer part of the group.)



Caution

When you configure new groups and policies or make changes to existing configurations, they are saved in the database when you click the Save button, but they are not yet distributed to the agents across your network. Once your configuration changes are complete, you must click the Generate configuration link in the menu bar to first view all new and edited configurations and then distribute them to the agents. This process is detailed in the [“Generating Rule Programs” section on page 4-97](#).

Figure 3-8 Add Hosts to Group



Distributing Software Updates

Cisco provides software updates via its web site (www.cisco.com) for both CSA MC and the agent. You can download these updates, install them on CSA MC, and then distribute them to agent systems across your network as easily as you deploy new rule programs. When you download a self-extracting executable update and install it on the server system, the agent software update files get placed under **Available Software Updates** in CSA MC (accessible from **Maintenance>Software Updates** in the menu bar).

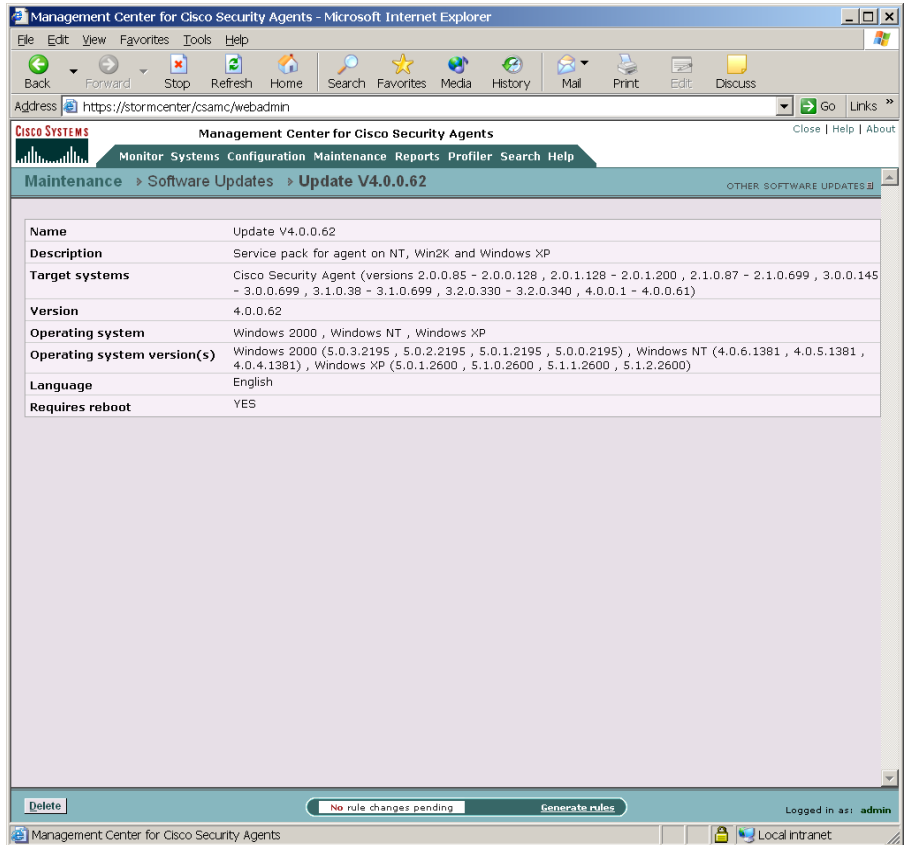
From the list of available updates that is created in the Available Software Updates page, you can make the appropriate updates available to agents through the Scheduled Software Updates page. Creating Scheduled Software Updates allows you to distribute updates to designated groups of agent systems. See [Configuring Scheduled Software Updates, page 3-28](#) for details.

The next time agent systems poll in to the server, the agent GUI prompts the user that there is a software update available (Windows only, UNIX agents receive no automatic prompt). Users can either install the update at that time or postpone the installation. If an automatic installation is an available option for a particular update, the update is automatically installed on designated agents systems the next time they poll in to the server.

From the Available Software Updates page, you can click on a particular update and view the following information (see [Figure 3-9](#)):

- Name of the software update, for example SP 3.3.0.58
- Description of the software update, for example Service pack for agent on NT and Win2K
- File, a link to the software update file itself on the server system
- Target system, a description of the system type for which the update is issued (agent and/or server)
- Version, this is the version of the software update
- Operating system, the operating system for which the update is issued
- Operating system version(s), the exact OS version numbers for which the update is issued
- Language, for example English
- Allow user interaction, tells you if configuring an optional "no user interaction" installation is possible

Figure 3-9 Available Software Updates Page



Configuring Scheduled Software Updates

Create Scheduled Software Updates to distribute an update or updates you have available in Available Software Updates to a selected group or groups.

To create Scheduled Software Update for distribution to agent systems, do the following.

-
- Step 1** From the menu bar **Maintenance** drop-down list, move the mouse over **Software Updates**. A cascading menu with further selections appears. Select **Scheduled Software Updates** (see [Figure 3-10](#)).
 - Step 2** Click the **New** button to create a new entry. This takes you to the update configuration page.
 - Step 3** Enter a **Name** for the update that makes it easily identifiable.
 - Step 4** Enter a **Description**. This is a useful line of text that is displayed in the list view and helps you to identify this particular configuration.
 - Step 5** Select the **Target operating system** for the update you're distributing (UNIX or Windows). When you select an OS, the available updates and selectable groups change accordingly.
 - Step 6** From the **Software update** pulldown list, select the Windows or UNIX update you want to distribute. Generally, it's called something like Update V4.0.0.52.
 - Step 7** **Enable update for hosts in selected groups** From the available list of groups, select one or more to distribute this update to.
 - Step 8** To select multiple items in a list box, hold down the Ctrl key as you select each item. To unselect a single item, hold down the Ctrl key when you click on the item in question. Press the Shift key to select multiple successive items.
 - Step 9** **Update time** Enter a time frame during which agent systems can receive and install updates. By default, the time frame is set to "any time" or for 24 hours. This way, users can update at any time. If you put a time limit on the update, for example enter 10:00 to 11:00 (this would be AM), then after 11:00, if the user misses this hour window, the update would not be available again until the same time the next day.
 - Step 10** If the update in question allows for an automatic installation (an install that occurs begins automatically without any user prompt) an **Automatic update** checkbox will be available on this page. Enable this checkbox for automatic software updates to take place on agents systems. In this case, the update takes place automatically during the time frame specified.

**Caution**

All updates (both automatic and not automatic) reboot systems within 2 minutes of the installation completing. This reboot cannot be stopped by the end user. Keep in mind, if the update is automatic, users are not prompted to begin the installation. Therefore, regardless if the end user is present or not, if the machine is running and an automatic update is received, both the install and the automatic reboot take place within the time frame specified in the update.

At this time, all software updates require systems to reboot after installation. (Summary information is included in the software update summary page, viewable when you click on the File link in the Software Updates page.)

If the update is not automatic (and the end user has an agent UI) a pop-up window which gives the end user the ability to postpone the update appears. (See the [“Suspend Agent Security” section on page A-13](#) for a description of the agent UI.)

See [Table 3-2 on page 3-32](#) for details on Automatic update and No user interaction behaviors on end user systems.

**Note**

To fully restrict end users from interacting with the agent installed on their system, you could use the Automatic update feature in combination with the "No user interaction feature" on the Group page (see the [“Configuring Groups” section on page 3-3](#)) and the "Agent service control" rule (see the [“Agent Service Control” section on page 4-24](#)).

Step 11 Click the **Save** button.

You must Generate rules to deploy software updates to agents.

Figure 3-10 Scheduled Software Updates Page

The screenshot shows a web browser window titled "Management Center for Cisco Security Agents - Microsoft Internet Explorer". The address bar shows the URL "https://stormcenter/csamc/webadmin". The page header includes the Cisco Systems logo and navigation tabs: "Monitor", "Systems", "Configuration", "Maintenance", "Reports", "Profiler", "Search", and "Help". The current page is "Maintenance" > "Scheduled Software Updates" > "Untitled_1".

The main content area contains the following configuration fields:

- Name: Update for 4.0 servers
- Description: This is a software update for windows servers
- Target operating system: Windows
- Software update: Update V4.0.0.62
- Enable update for hosts in selected groups: A dropdown menu is open, showing options: Default Servers, DHCP and DNS servers - Dedicated, IDS Mode Servers, IIS Web Servers - Dedicated, and Mission Critical Systems.
- Update time (hh:mm): from 00:00 to 23:59
- Automatic update ?

At the bottom of the page, there are buttons for "Save" and "Delete", a status bar indicating "14 rule changes pending", and a "Generate rules" button. The user is logged in as "admin". The browser's status bar shows "Local intranet".

Table 3-2 Software Update/Usr Interaction Config Combinations

Software Updates	Agent UI present	No user interaction (no agent UI)
Not Automatic	<p><i>Start install behavior: Visible</i></p> <p>Pop-up window prompts user to install update. (This window warns that after install completes, system automatically reboots after 2 minutes. User cannot stop this reboot once installation begins.)</p>	<p><i>Start install behavior: Visible</i></p> <p>Pop-up box prompts user to install update. (This box warns that after install completes, system automatically reboots after 2 minutes. User cannot stop this reboot once installation begins.)</p>
	<p><i>Delay install option: Click Postpone button</i></p> <p>"Postpone" button on pop-up window allows user to delay installation for 1-10 days. An "Update Available" button appears on agent UI for updating later.</p>	<p><i>Delay install option: Do not click OK button.</i></p> <p>User must click OK button on pop-up box to begin update. Pop-up box remains on screen until user performs the update.</p> <p>*Although pop-up box does not disappear, install does not take place unless the user clicks OK.</p>

Table 3-2 Software Update/Usr Interaction Config Combinations (continued)

Software Updates	Agent UI present	No user interaction (no agent UI)
Automatic	<i>Start install behavior: Not visible</i> User is not prompted before update installation begins.	<i>Start install behavior: Not visible</i> User is not prompted before update installation begins.
	<i>Delay install option: None</i> When installation is complete, a 2 minute automatic reboot warning message appears. User cannot stop this reboot and has 2 minutes to save any open documents. *Regardless if the user is present or not, if the machine is running, both the install and the automatic reboot take place.	<i>Delay install option: None</i> When install is complete, a 2 minute automatic reboot warning message appears. User cannot stop this reboot and has 2 minutes to save any open documents. *Regardless if the user is present or not, if the machine is running, both the install and the automatic reboot take place.

The next time agents poll in to CSA MC, they receive a prompt (if update is not automatic, see [Figure A-8 on page A-15](#)) letting them know there is an update available. They can choose to install it immediately or select to postpone the installation for a later time. On UNIX agent systems, use the `csact1` utility to check for software updates and to install them. See [Appendix A, “Cisco Security Agent Overview”](#) for details.

The agent online help system (Windows only) provides instructions to users on the software update download and installation process. You may want to refer them to it.



Building Policies

Overview

The policies you create on the Management Center for Cisco Security Agents should reflect a well-planned, enterprise-wide security policy. If your corporate enterprise already has security guidelines in place, the rules that form your policies should reflect those guidelines.

It is important that you spend time charting out your security needs in advance rather than attempting to backfill holes as they are discovered. Because both networks and network security are dynamic entities, it is expected that you will need to adjust policies to meet the changing and growing needs of your enterprise. A well thought-out security plan is certain to save you time in the end.

This section contains the following topics.

- [Developing a Security Policy, page 4-2](#)
- [About Rules, page 4-5](#)
- [Policy Components, page 4-6](#)
- [Querying the User, page 4-10](#)
- [Configuring Policies, page 4-13](#)
- [Rules Common to Windows and UNIX, page 4-24](#)
- [Windows Only Rules, page 4-54](#)
- [UNIX Only Rules, page 4-80](#)
- [Attaching Policies to Groups, page 4-92](#)

- [Using Test Mode, page 4-95](#)
- [Generating Rule Programs, page 4-97](#)

Developing a Security Policy

If you are crafting your own policies, please refer to [Chapter 12, “Policy Definition Guidelines”](#) for information.



Caution

To maintain the integrity of the preconfigured policies shipped with CSA MC, it is recommended that you do not change them. If you are using preconfigured policies but want to edit them slightly due to your own site’s needs, you should instead create a new policy and add that policy to the group in addition to the preconfigured policy.

Note that each pre-configured rule, policy, and group page has data in the expandable **+Detailed** description field explaining the item in question. Read the information in these fields to learn about the items described and to determine if the item in question meets your needs for usage.

A corporate security policy should temper business concerns with security concerns. It should allow the user community to access required resources, while protecting that community from the dangers those resources can introduce. To achieve this goal, it is crucial to have a carefully planned network security policy in place to safeguard valuable organizational resources and information.

Before configuring your policies, it is important to understand exactly what network resources and services you want to protect and what threats you are most concerned about. The first step in planning a security policy is identifying the resources your user community requires to do business. That could include specific applications, protocols, network servers and web servers. Collect this information and use it to design the main features of your policy.

Providing Safe Access to Required Resources

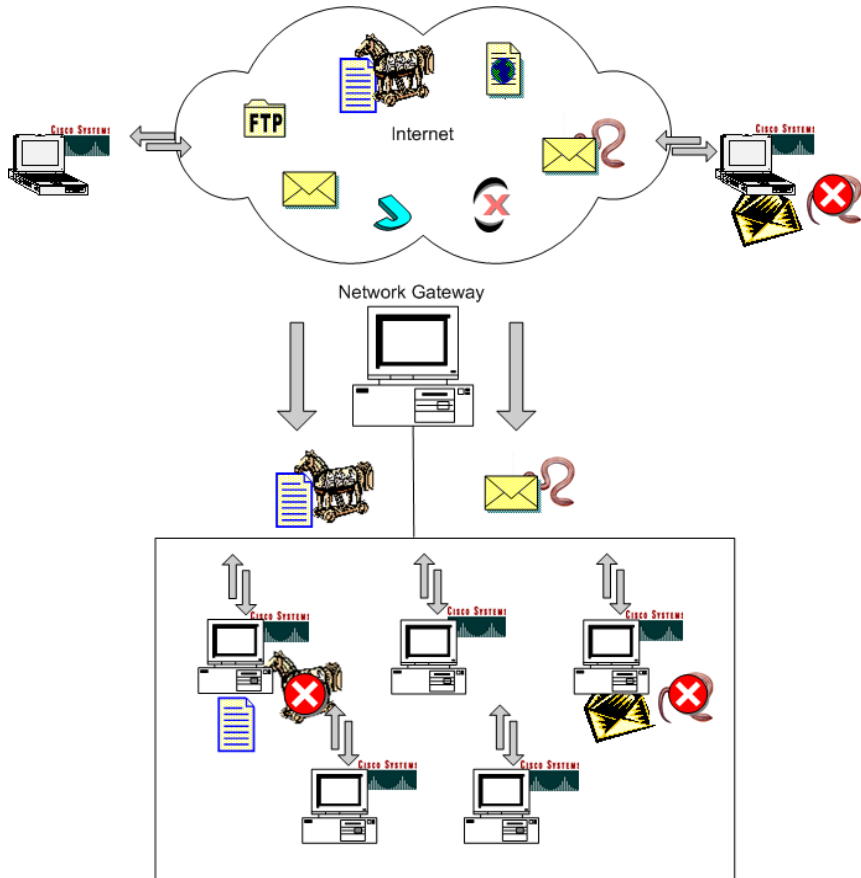
As you determine the network resources that are required by your user community, you can identify some of the threats posed against those resources. For example, while putting together a security plan, you might find it beneficial to limit access to some resources based on various parameters such as traffic direction and allowed file types.

Upon examining past breaches of security, you could determine that email attachments and Internet file downloads pose the greatest threat to your network. In this case, you would want to develop policies to diminish the danger of accessing these particular resources. Your security plan should then incorporate policies for commonly used services such as HTTP, POP3, IMAP (email), and FTP (file transfer).

You could take a couple of approaches to enforcing your security plan depending upon the immediacy of any perceived threats and your basic corporate philosophy toward security. Both approaches are equally valid. On the one hand, you might choose to enforce known good behaviors and selectively add targeted restrictions. This would be a more permissive security model. This approach facilitates uptime, but may be less secure. Conversely, you could decide to shut everything down and then slowly add targeted permissions. This approach is far more restrictive and some legitimate requests could be rejected, but this may be suitable for highly secured environments. You could use both approaches for different groups.

As your security plan evolves, you can refine your policies, making them more or less granular to keep pace with your user community's needs. Your network system security depends on your implementing security policies carefully, and checking to see that they work as intended.

Figure 4-1 Protecting Information



Formulate a policy to protect systems from common email worms e.g. the ILOVEYOU worm and Trojans (Back Orifice). Once these attacks infiltrate your network and propagate to the user community, a well-defined policy can identify errant system actions and stop an attack before it can damage mission-critical information.

About Rules

Rules are the foundation of your security policies. CSA MC lets you create several rule types. Each rule type requires you to enter varying combinations of information using a specific syntax. For example, the following basic access control rules require information as follows:

Use file access control rules to allow or deny what operations(read, write) selected applications can perform on files according to:

- the action you are allowing or denying
- the application attempting to access the file
- the operation (read, write) attempting to act on the file

Use network access control rules to control access to specified network services according to:

- the action you are allowing or denying
- the application attempting to access the service or address
- the direction (client, server) of the communication
- the service a system is attempting to use
- the address a system is attempting to communicate with

Use registry access control rules (Windows only) to allow or deny selected applications from writing to specified registry keys according to:

- the action you are allowing or denying
- the application attempting to write to the registry keys and values

Use COM component access control rules (Windows only) to allow or deny selected applications from accessing specified COM components according to:

- the action you are allowing or denying
- the application accessing the COM component

Other types of policies shipped with CSA MC provide event correlation and heuristic features which can be enabled on a per group basis, like portscan detection, SYN flood protection, the prevention of predictable TCP sequence numbers, and the blocking of malformed IP packets. (These features are located on the Network shield rule page.) This is especially useful for network servers. (See [Chapter 5, “Using System Correlation Rules”](#) for more information.)

Operating under the direction of assigned policies, agents provide overall system protection, tying together the controlling of various system components.

Combining Policies

You can attach multiple policies to a group. Moreover, a host can belong to multiple groups and inherit policies from all of them. For example, a desktop can belong to the Finance group and inherit the Accounting policy. It can also belong to the All group through which it receives the Corporate Mail Policy.

When more than one policy is associated with a host, the rules in the individual policies are merged as though they were all defined within a single policy. In particular, the rules are ordered in the same sequence as they would be within a single policy. See the section on [Writing Rules: Allow vs. Deny](#), page 4-6 for priority order information.



Note

You can view merged policy rules at both the group and host levels.

Policy Components

The following sections describe the various components you must configure as part of the access control rules that will form your policies.

Writing Rules: Allow vs. Deny

When you write rules, you create them as allow or deny actions. When you add your rules to policies, CSA MC orders them in the following manner within each policy.

Priority 1	Add process to application class
Priority 2	High Priority Deny
Priority 3	Allow
Priority 4	Query User (Default Allow)

Priority 5	Query User (Default Deny)
Priority 6	Deny
Priority 7	Default Action (Allow)

The priority listings beside each bullet item indicate the manner in which CSA MC processes rules. All priority 1 rules (Add to application class) are checked first and priority 6 rules (deny) are checked last and that is only if no other higher priority rules have already been triggered by a system action.

**Note**

An **Add process to application class** rule type takes precedence over all other types. But you should note that the only action of this rule is to build a dynamic application class for any rules that make use of it. The application-builder rule does NOT override other rules as allow, deny, and query rules do when triggered. See [Chapter 6, “Using Application Classes”](#) for details on dynamic application class building.

For every policy you configure, the default action of that policy is Allow. All policies allow all system actions until you write a rule denying a specific action. Following that logic, it is unlikely that you would write allow rules unless they are to make exceptions to deny rules you are writing within a policy or for monitoring purposes (see [Monitoring Access, page 4-9](#)). If you do write a stand-alone allow rule, because the default action is allow, the allow rule itself is then essentially irrelevant.

**Note**

Generally, if you’re going to write a high priority deny rule, you’ll want to have it stand alone in a policy. You can then attach that high priority deny policy to several groups. The high priority deny can control access to resources that you want to deny across the board to either your entire network or to specific groups within your network.

A good model for configuring rules within policies would be to take the priority levels into account and work from the bottom up, lowest priority to highest priority. Before you even add a single parameter to a rule, by default (priority 7), it allows all system actions. First, write a deny rule (priority 6) and then if you want to make any exceptions to that particular deny, write an allow (priority 3)

rule. Next consider using query rules for access controls that allowing the user to decide if an action should be allowed or denied. Lastly write any high priority deny rules (priority 2) you might need.

Writing Rules: Manipulating Precedence

In addition to using the selected "action" type to order rules within a policy, CSA MC uses the selected logging type as a way to suborder similar rules within a policy. Logging automatically takes precedence over disabled logging if the action type is the same for multiple rules in a policy. Therefore, for rules of a given priority, e.g. Allow, a Log rule will be evaluated before a No Log rule.

For most policies, this automatic ordering and subordering of rules provides the desired effect when policies are combined and deployed. However, there are cases when the CSA MC ordering scheme causes policies to behave in an undesired manner. For this reason, most rule types provide a checkbox that allows you to manipulate how similar rule types are subordered within a policy. This checkbox, called **Take precedence over other <similar action> rules**, is located in the rule configuration page. A rule with this precedence checkbox selected is evaluated before similar rules that do not have this checkbox selected.

Here is an example of two rules within the same policy which do not behave as expected due to automatic rule ordering. There are two Network access control rules in the same policy as follows:

- Log, Deny, All applications, acting as a server, for TCP/1-65000
- No Log, Deny, All applications, acting as a server, for TCP/1900

The rule that involves connections on TCP/1900 would be denied and logged despite the fact that logging is not selected for that rule. This is because the rule involving connections on TCP/1-65000 would be evaluated within the policy first and connections made on TCP/1900 would go to the event log even though the rule did not have logging selected.

In this example, using the **Take precedence over other <action> rules** checkbox in the TCP/1900 rule would allow you to designate its precedence as higher than other deny rules in the policy, giving you the ability to suppress log messages for actions you want to be denied but for which you do not want to be continually notified due to another rule within the policy.

**Caution**

The **Take precedence over other <action> rules** checkbox is a rule ordering tool you should rarely need. In most cases, the CSA MC automatic ordering of rules is sufficient. But if you are using this checkbox to manipulate rule ordering, you should understand the following rule order scheme. Within a given policy, rules are sorted using this criteria:

- * Action type
- * Precedence checkbox On/Off
- * Log checkbox On/Off

**Note**

For a given policy, if you have multiple rules of the same action type, the same logging type, and the same "take precedence" type, the ordering of these rules is inconsequential within the policy because there is no differential criteria by which to order them.

Monitoring Access

As stated in the previous sections, Cisco Security Agents allows all actions by default. Therefore, it is unlikely that you would write an allow rule unless it's to make an exception to a deny rule. However, in some instances, you may want to write an allow rule with logging enabled or a query user (default allow) rule with logging enabled so that CSA MC receives an event log when a specific event occurs on a system. This way, the event is allowed, but it is also logged, letting you monitor when it occurs.

For file access monitoring capabilities, see [File Monitor, page 4-45](#).

Making a Policy Mandatory

Normally, when you create a policy, you have to "associate it" with a specific group to have it apply to that group. (See [Attaching Policies to Groups, page 4-92](#).) But if you designate a policy as "Mandatory" via the checkbox available in the Policy configuration page, it is automatically attached to ALL groups. Instructions for configuring this mandatory policy are provided in [Configuring Policies, page 4-13](#).

You might want to make a policy mandatory and have it automatically applied to all groups in order to prevent some critical service from being inadvertently banned. For example, you could create a mandatory policy to prevent DNS or DHCP from being disabled by an overly restrictive rule.

Querying the User

When you create access control rules, beyond simply allowing or denying a specific action, you can select to query the user when an action triggers the rule in question. The user can then decide to allow the action or deny it at that time. When you select to query the user, you are also configuring the rule to allow or deny the action by default if the query is not answered within 5 minutes. If the user is not logged in to the system, the default action is taken immediately.



Note

See the [“Network Worm Protection” section on page 5-3](#) and the [“Trojan Detection” section on page 5-4](#) for information on Query User pop-ups that include a **Terminate** button.



Caution

For UNIX rules, Query user options are not available. For Windows agents, if "No user interaction" is selected for the group, there are no query user pop-up boxes displayed. The default is immediately taken on all query user rules and heuristics that are present in the assigned policies.

When an action is attempted on a system where a query user rule is triggered, a pop-up box appears on the system where resource is located.

Figure 4-2 Query User Pop-up Box

In the Query pop-up box, the user reads the information given on the attempted action and selects one of the following:

- **Yes:** Allows the application access to the resource in question.
- **Yes to All:** Allows the application access to all related query user protected resources, with no further queries appearing.
- **No:** Denies the application access to the resource in question.
- **No to All:** Denies the application access to all related query user protected resources, with no further queries appearing.

Query user Yes to All and No to All responses apply to both file and registry resources equally. For example, if the user answers Yes to All for a query user pop-up triggered by a file access rule, the application name in question is now also allowed to access query user controlled registry resources as well.

Query user Yes to All and No to All responses for network resources apply to the application name in question attempting to access a network service. For example, if the user answers Yes to All for a query user pop-up triggered by a network access rule, the application name in question is now also allowed to access all query user controlled network resources.

**Note**

When you configure your Query User rule, the text you type into the rule's **Text used to query user** field is the same text that will appear in the Query User pop-up box to explain what is occurring on the system to the user. Therefore, making this information descriptive of the system action that triggered the pop-up is important.

**Caution**

With file access control rules, the query user pop-up box appears on the system where the file or files in question are located. If a user is attempting to remotely access restricted files, the pop-up box appears on the remote machine where the files are located, not on the user's machine. That being the case, you would likely not want to place "query user" file access restrictions on files that are kept on an unattended system.

**Note**

Note that it takes several seconds for buttons on the query user pop-up box to become active (selectable).

Caching Query Responses

When a user responds to a Query User box (by pressing Yes, Yes to All, No, No to All), the agent remembers the response and caches it. This way, if the same rule is triggered again, the action is allowed or denied based on what the user answered previously with no pop-up query box appearing again. Query user pop-up boxes have an hour inactivity timer. If the action in question is not attempted again within an hours time, the user is queried the next time the action is attempted.

For Query User rule logging information, see [Chapter 8, "Event Logging Alerts."](#)

**Note**

The Advanced tab on the agent user interface (Windows only) lets the user clear the query user cache. Clicking the Clear button in the agent Advanced tab tells the system to forget all query responses and display a Query User pop-up box when the event in question occurs again. Use this feature for troubleshooting or if the

user answered incorrectly and wants to change the response. See [Appendix A, “Cisco Security Agent Overview”](#) for further information on the agent user interface.

Configuring Policies

When you configure a policy, you are combining your configured rules under a common name. That policy name is then attached to a group of hosts and it uses the rules that comprise the policy to control the actions that are allowed and denied on those hosts. You can have several access control and monitoring rules in one policy.

See [Chapter 5, “Using System Correlation Rules”](#) for information on event correlation and heuristic rules such as Trojan detection and network worm protection.



Note

Carefully read the section on writing allow and deny rules ([“Writing Rules: Allow vs. Deny” section on page 4-6](#)) so that you will understand how rule precedence works once policies are deployed. You should also refer to the chapter on configuration Variables ([Chapter 7, “Configuring Variables”](#)) to help you understand the information required by the rule text fields.



Caution

To maintain the integrity of the preconfigured policies shipped with CSA MC, it is recommended that you do not change them. If you are using preconfigured policies but want to edit them slightly due to your own site’s needs, you should instead create a new policy (you can do this by cloning an existing policy) and add that policy to the group in addition to the preconfigured policy.

To configure a policy, do the following.

Step 1 Move the mouse over **Configuration** in the menu bar and select **Policies** from the drop-down list that appears. The list of existing Policies is displayed. CSA MC ships with several pre-configured sample policies.

Step 2 Click the **New** button to create a new policy. You may be prompted to select whether this is a Windows or a UNIX policy. Click the appropriate button to go to the policy configuration view.



Tip You can click the <#>rules link on the policy list page to go directly to the rules contained in the policy.

Step 3 In the policy configuration view, enter a unique **Name** for your policy. Names are case insensitive, must start with an alphabetic character, can be up to 64 characters long and can include hyphens and underscores _ . Spaces are also allowed in names. Use a descriptive name that you can easily recognize in the group policy listbox when you are attaching policies to groups.

Step 4 Enter a **Description** of your policy. This description is visible in the policy list view. Optionally, expand the **+Detailed** field to enter a longer description.

- Step 5** Optionally, select the **Mandatory policy for all groups** checkbox to have this policy automatically attached to ALL groups. Before selecting this checkbox, you should read the Caution note below and refer to [page 4-9](#) for further explanation of why you would want to make a policy mandatory.

**Caution**

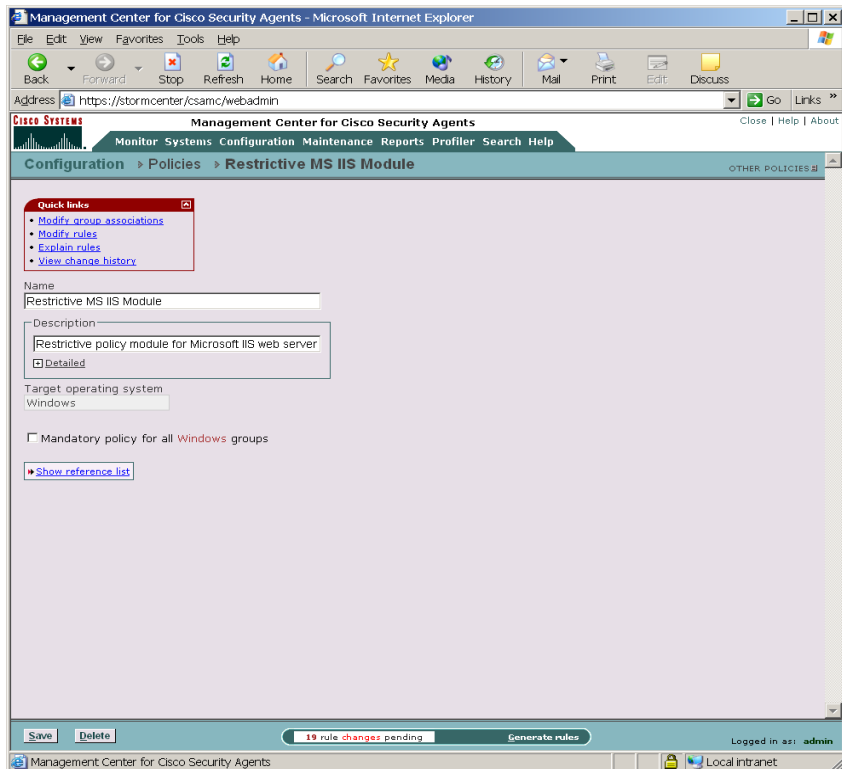
Selecting the **Mandatory policy for all groups** checkbox attaches this policy to ALL groups for that operating system. Unselecting the checkbox does NOT remove the policy from all groups. To remove a mandatory policy, you must click the **Modify group associations** link at the top of the Policy configuration page and manually remove groups from the policy.

Some additional configuration notes on the **Mandatory policy for all groups** checkbox are:

- Any new groups you create are automatically associated with mandatory policies.
- When importing a mandatory policy, the mandatory setting is NOT maintained. Designating a policy as mandatory is an explicit action that must be taken by the administrator.

- Step 6** Click the **Save** button.
- Step 7** Now you add rules to your policy. Click the **Modify rules** link at the top of the page.
- Refer to the following sections for details on adding, copying, and configuring rules.

Figure 4-3 Policy Creation View



Adding Rules to a Policy

First, click the **Modify rules** link at the top of the Policy page to go to the Rules page.

To add rules to this policy, click the **Add rule** link in the Rules page. A menu list of the available rule types appears. Click on one to select it. This takes you to the configuration view for this rule type. Note that this rule contains no parameters until you create them.

**Note**

Refer to the following sections for details on configuring particular types of rules.

Use the **Enable** and **Disable** buttons in the Policy configuration view to enable or disable rules within a policy without having to navigate to the configuration view for that particular rule. Select the checkbox for the rule you want to enable or disable and click the corresponding button.

The **ID** column in the Rules section is the rule ID number assigned to the particular rule in question. This number increments each time a new rule is created. It is only used as an identifier for the rule. This ID is referenced in Event Log messages and can help you refer back to a particular rule.

The **Events** column in the Rules section (see [Figure 4-4](#)) displays the number of events generated by the rule in the last 24 hours. Clicking this number link takes you to a list of the events themselves.

Filtering the Rules Display

The Groups configuration page and the Rules configuration page each display a table listing either the rules attached to the group or the rules included in the policy (see [Figure 4-4](#)). On both pages, there is a **View All rules** item above the table. Clicking the **All** link here lets you filter your view of this rule list by selected rule type. When you click **All**, a pop-up appears listing the rule types present in the policy or policies. Select a rule type from the pop-up, and that is now the only rule type displayed in the table. You can also select to view only enabled rules by selecting the **show enabled rules only** checkbox and then select the rule type you wish to view.

**Note**

When you filter the rules display, other rules are NOT removed from the policy. It is only your view of the policy that changes. You can revert back to the entire summary view by selecting **All** from the same pop-up menu.

This filtering feature is useful when lists of rules grow extensive and you want to pare down your view to specific rule types.

Copying Rules between Policies

Use the **Copy** button in conjunction with the pulldown lists at the bottom of the Policy page to copy selected rules to another policy that you designate. Copying rules across policies works similar to the way cloning configurations works. (You can also clone rules within policies using the Copy button that will be described in this section.)

To copy selected rules from one policy to another policy, do the following:

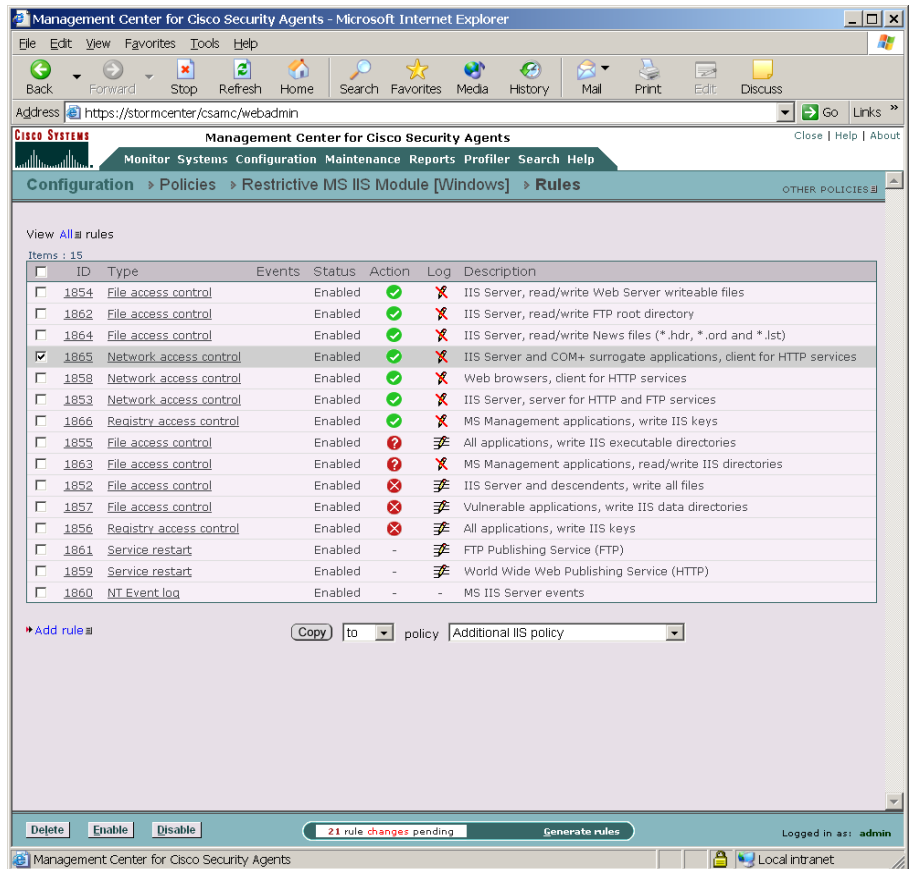
-
- Step 1** From the Policy page (see [Figure 4-4](#)), select the checkbox for the rule or rules you want to copy to another policy.
 - Step 2** Beside the Copy button, **to** is the default selection in the pulldown menu. (Do not change this for copying individual rules between policies.) From the **policy** pulldown list, select the name of the policy to which you want to copy the selected rule or rules.
 - Step 3** Click the **Copy** button.

All checked rules are copied to the selected policy.

To clone rules within a policy, repeat step 1 above. Then, rather than selecting another policy in the policy pulldown list, select the current policy you are in from that same pulldown. Selected rules are cloned within the same policy when you click the Copy button.

Select **from** in the pulldown menu beside the **Copy** button to copy ALL the rules from the selected policy (in the policy pulldown list) to the current policy.

Figure 4-4 Copy Rules between Policies



Comparing Configurations

When you select the checkbox next to 2 items (you cannot compare more than 2 configurations at a time) and click the **Compare** button, CSA MC displays the configurations side by side and highlights the differences in red (see Figure 4-5). Once you've examined how the configurations compare, you can select to merge

specific rules, to copy rules to another policy, or to copy rules to a new policy. (You can compare application classes and variables, but you can only copy and merge rules from the compare page.)

The purpose of this compare tool is to assist you after you've imported configurations or upgraded CSA MC. These processes can cause you to have duplicate or very similar configuration items. Comparing and merging configurations can help you to more easily consolidate duplicate items. This Compare utility is also available for Application Classes and Variables.

Feature notes:

- When you compare policies, the similar rules within those policies are displayed side by side with the differences highlighted in red. If there are no differences, rule description text appears in black.
- If there is a rule in one policy and no corresponding similar rule in the second policy, there is nothing displayed beside that rule in the comparison.
- If you have rules in your policy comparison that have the same description, application class and other configuration items, they will not appear side by side if they have different logging options selected or different Allow/Deny actions. Logging and allow/deny actions change the priority of the rule within the policy. If the priority is not the same for each rule, they are not displayed side by side.

Figure 4-5 Compare Policies

The screenshot shows the Management Center for Cisco Security Agents web interface in Microsoft Internet Explorer. The browser address bar shows the URL: <https://stormcenter/csamc/webadmin>. The page title is "Management Center for Cisco Security Agents". The navigation menu includes "Monitor", "Systems", "Configuration", "Maintenance", "Reports", "Profiler", and "Search Help". The current page is "Configuration > Policies > Compare Copy of Desktop Module and Desktop Module".

The comparison table shows the following details:

Name	Copy of Desktop Module	Desktop Module
Description	Base policy module for desktops	Base policy module for desktops
Detailed Description		
Platform	Windows	Windows
Mandatory	No	No
Rules	14 items	13 items

Use the checkboxes to merge or copy rules to a new policy or to existing policies.

File access control

Enabled	Copy of Desktop Module	Desktop Module
	<input type="checkbox"/>	<input checked="" type="checkbox"/>
Description	1970 Email applications, read/write dynamically quarantined files	1845 Email applications, read/write dynamically quarantined files
Detailed Description		
Action	High Priority Deny	High Priority Deny
User Query		
Log	Yes	Yes
Application Classes	Email applications	Email applications
Dynamic Application Classes		
File Operations	Read Write	Read Write
Files	***@*(dynamic)	***@*(dynamic)

File access control

Enabled	Copy of Desktop Module	Desktop Module
	<input type="checkbox"/>	<input type="checkbox"/>
Description	1974 Email applications, user invoked applications and command shells	1849 Email applications, user invoked applications and command shells
Detailed Description		
Action	High Priority Deny	High Priority Deny
User Query		
Log	Yes	Yes
Application Classes	Email applications	Email applications
Dynamic Application Classes		
File Operations	Read	Read

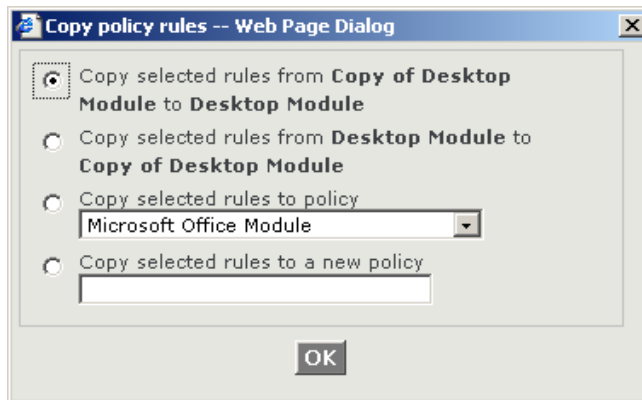
At the bottom of the interface, there are buttons for "Copy" and "Delete", a status bar indicating "39 rule changes pending", and a "Generate rules" button. The user is logged in as "admin".

Merging or Copying Policies

Merge or copy rules by selecting the available checkbox above the rule or rules in question. When you click the Copy button in the bottom frame, a pop-up window appears. From this window, you select to do one of the following:

- Copy the selected rules from one policy in the comparison to the other policy in the comparison
- Copy the selected rules to another policy you select (not part of the current comparison)
- Copy the selected rules to a new policy which you create at this time by entering its name in the available field

Figure 4-6 Copy Policy Pop-up Box



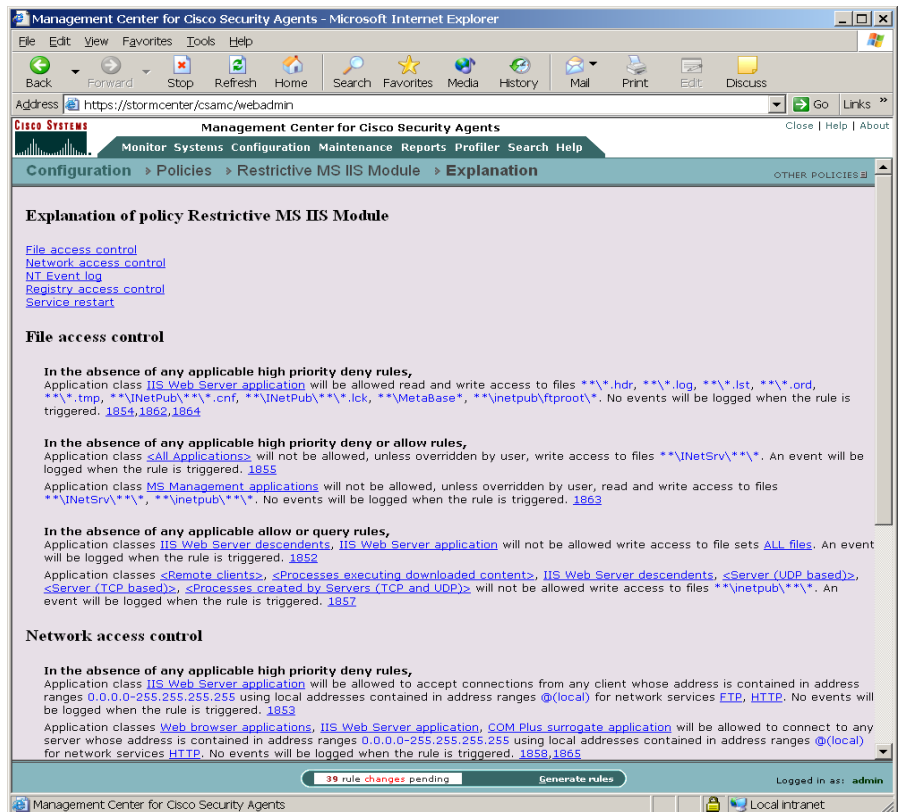
View Change History

At the top of each rule page, there is a View change history link. Click this link to go to a page which lists all the changes that have been made to this rule. This View change history link is also available for Application classes, Variables, and Policies.

Explanation of Rules

CSA MC provides an explanation, in paragraph form, of the policy in question, describing each rule and its role in the policy. Clicking the **Explain rules** link in the Groups, Host, or Policy page, takes you to this paragraph explanation. See [Figure 4-7](#).

Figure 4-7 Rule Explanation Page



Rules Common to Windows and UNIX

The following rules are available for both Windows and UNIX policies.

Agent Service Control

Use the Agent service control rule to control whether users (including non-administrator users) are allowed to suspend agent security. Suspending agent security disables all rules until security is manually resumed or the system is rebooted.

This rule can also restrict administrators from stopping and starting the Cisco Security Agent service (This is via a net stop command on Windows or via /etc/init.d/csa stop on UNIX. See [Chapter 10, “Using Management Center for Cisco Security Agents Utilities”](#) for details).

If you use this rule to deny service stop and start actions, the agent service cannot be stopped on the system in question and agents cannot be uninstalled.

**Note**

Although agents cannot be uninstalled by administrative users if this rule denies the stopping of the agent service, this rule does not prevent agent software updates from occurring.

**Note**

This rule also controls logging for built-in agent self-protection rules. See [page 4-27](#) for details.

These instructions are a continuation of [Configuring Policies, page 4-13](#).

- Step 1** To add rules to your policy, click the **Add rule** link at the bottom of the rule list. A pop-up list of the available rule types appears.
- Step 2** Select the **Agent service control** rule. This takes you to the configuration view for this rule type (see [Figure 4-8](#)).

- Step 3** In the Agent service control rule configuration view, enter the following information:
- **Description** Enter a description of this rule. This description appears in the list view for the policy. Optionally, expand the **+Detailed** field to enter a longer description.
 - **Enabled** Use this checkbox to enable this rule within the policy. (It is enabled by default.) By not selecting this checkbox, you can save this rule, but it will not be active in the policy and it will not be distributed to groups.
- Step 4** **Take the following action** (Select an action type from the pulldown list. For further details on using allow versus deny rules, see [Writing Rules: Allow vs. Deny, page 4-6](#).)
- **High Priority Deny** Select this action type to create a deny rule that takes precedence over all other allow, deny, and query rules. For example, if you configure an allow rule that conflicts with this high priority deny, the high priority deny always takes precedence. Under the same circumstances, if the deny is not high priority and a conflicting allow rule exists in the same policy, the allow takes precedence.
 - **Allow** Select this action type to create an agent service control rule that allows users/admins to enable and disable the agent service. Because the default action of all policies is "allow", generally, you'll only want to configure allow rules as exceptions to existing deny rules within policies.
 - **Query User (Default Allow)** Select this action type to prompt the user when the action you indicate occurs. The user then has 5 minutes to answer Yes, No, Yes to all, or No to all. By default, it will be allowed unless the user decides otherwise. See Query User for more details. (Query User options are not available for UNIX rules.)
 - **Query User (Default Deny)** Select this action type to prompt the user when the action you indicate occurs. The user then has 5 minutes to answer Yes, No, Yes to all, or No to all. By default, it will be denied unless the user decides otherwise. See Query User for more details. (Query User options are not available for UNIX rules.)
- Text used to query user** If you are configuring a Query User rule, the text you type into this field is the same text that will appear in the Query User pop-up box to explain what is occurring on the system to the user.
- **Deny** Select this action type to create an agent service control rule that denies users/admins the ability to enable and disable the agent service.

Step 5 and

- **Log** Enable this checkbox to turn logging on for this rule. Generally, you will want to turn logging on for all deny rules. This means that the denied system action in question is logged and sent to the server at regular time intervals.
- **Take precedence over similar <action type> rules** Enable this checkbox to manipulate policy rule precedence so that this rule is evaluated before other similar rules. You should generally NOT require this checkbox. Do not use it without understanding how it works. See [Writing Rules: Manipulating Precedence, page 4-8](#) for details.

Step 6 when (The functionality differs on Windows and UNIX rule pages.)

- Windows:

any user attempts to suspend security from agent UI

This checkbox controls the Suspend security feature available from the agent UI. Allowing this action permits all users (including non-administrative users) to disable all rules on the agent until they are re-enabled or until the next reboot. Users cannot make use of this UI feature if this action is denied by this rule. (Also note that if there is no agent UI present, agent security cannot be suspended.)

any Windows administrator user attempts to stop the agent service

This checkbox controls whether users with administrator privileges can stop the agent service from the Service Control Manager or by running `net stop "Cisco Security Agent"` from a command prompt.

any application attempts to modify the agent configuration (Built-in Agent Self-Protection)

The Cisco Security Agent has built-in (non-configurable) security policies which protect agent binaries and data. (Note that this protection is only offered when the agent service is running and is not suspended.) While you cannot turn these built-in rules off while the agent is active, by turning off logging for the Agent service control rule, you can suppress log messages when these event types are triggered. You can also **Select any application classes to be excluded from logging**.

Likely, you should not see these message types very often and you would only want to exclude an application class from these built-in rules if you're running Anti-Virus software on the agent system. AV software can trigger self-protection rules. (However, you can exclude any application you choose.)



Note If you select the **modify agent configuration** checkbox, the rule must be high priority deny. This is because there may be rule precedence issues when you combine these rules types. To work as expected in policies, this must be a high priority deny. If you attempt to select any other action type, when you click Save, the action is automatically changed to High priority deny for you.

- UNIX:

Applications in any of the selected classes

Select *one or more* preconfigured application classes here.

Note that the entry, <All Applications>, is selected by default. You can use this default or you can unselect it and create your own application classes. For application class configuration details, see [Chapter 6, “Using Application Classes.”](#)

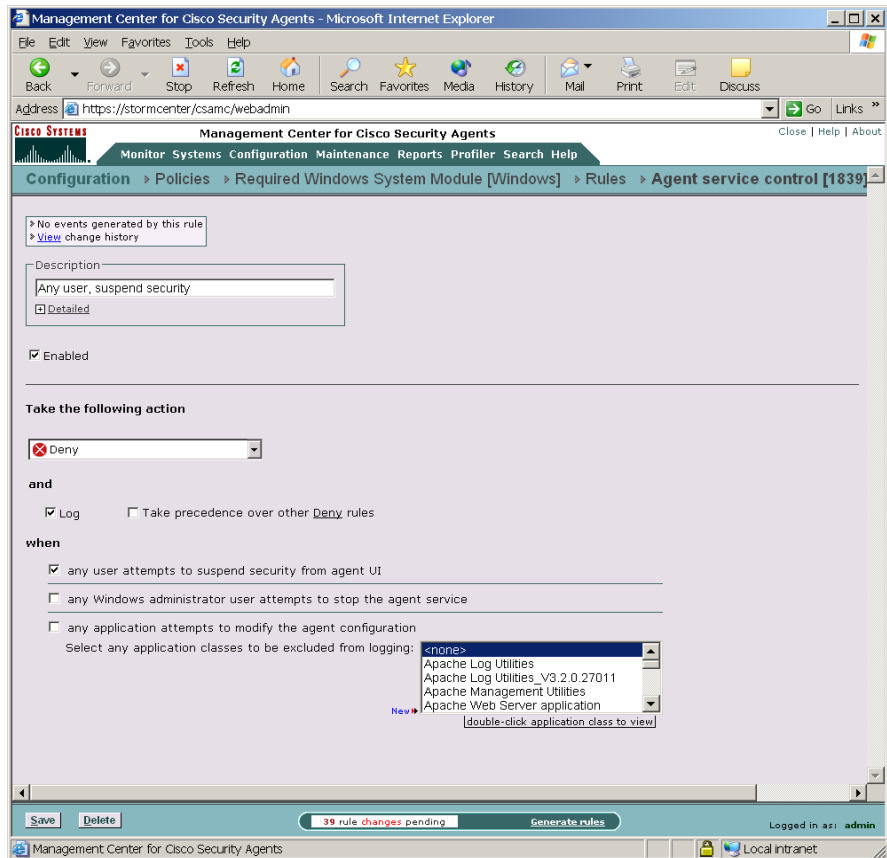
try to stop the agent service.



Note On UNIX systems, anyone with root access can stop the agent service. To prevent this, while still allowing administrators to stop the agent service, you would configure an Agent service control rule to Deny <All Applications> from stopping the service. Then configure another Agent service control rule which Allows only a UNIX Secured Management application class to stop the service.

Step 7 Click the **Save** button.

Figure 4-8 Agent Service Control Rule (Windows)



Application Control

Use Application control rules to control what applications can run on designated agent systems. This rule type does not control what application can access what resources as do other access control rules. This rule type can stop selected applications from running on systems. If you deny an application class (in total) in this rule, users cannot use any application in that class.

With this rule, you can also prevent an application from running only if that application was invoked by another application you specify. This way, you could prevent a command prompt from running on a system if it is invoked by an application that has downloaded content from the network.

**Note**

These instructions are a continuation of [Configuring Policies, page 4-13](#).

- Step 1** To add rules to your policy, click the **Add rule** link at the bottom of the rule list. A pop-up list of the available rule types appears.
- Step 2** Select the **Application Control** rule. This takes you to the configuration view for this rule type (see [Figure 4-9](#)).
- Step 3** In the Application control rule configuration view, enter the following information:
- **Description** Enter a description of this rule. This description appears in the list view for the policy. Optionally, expand the **+Detailed** field to enter a longer description.
 - **Enabled** Use this checkbox to enable this rule within the policy. (It is enabled by default.) By not selecting this checkbox, you can save this rule, but it will not be active in the policy and it will not be distributed to groups.
- Step 4** **Take the following action** (Select an action type from the pulldown list. For further details on using allow versus deny rules, see [Writing Rules: Allow vs. Deny, page 4-6](#).)
- **Add current/new process to application class** Use this for defining dynamic application classes. See [Chapter 6, “Using Application Classes”](#) for details.

**Note**

Creating dynamic application classes from the Application control rule is a bit different than creating them from other rule types. Because this rule has two application class fields, you can choose to add the current application to the dynamic class or choose to add the new application that is invoked by the first application to the dynamic class.

- **High Priority Deny** Select this action type to create a deny rule that takes precedence over all other allow, deny, and query rules. For example, if you configure an allow rule that conflicts with this high priority deny, the high priority deny always takes precedence. Under the same circumstances, if the deny is not high priority and a conflicting allow rule exists in the same policy, the allow takes precedence.
- **Allow** Select this action type to create an application control rule that allows the applications you specify to run. Because the default action of all policies is "allow", generally, you'll only want to configure allow rules as exceptions to existing deny rules within policies.
- **Query User (Default Allow)** Select this action type to prompt the user when the action you indicate occurs. The user then has 5 minutes to answer Yes, No, Yes to all, or No to all. By default, it will be allowed unless the user decides otherwise. See Query User for more details. (Query User options are not available for UNIX rules.)
- **Query User (Default Deny)** Select this action type to prompt the user when the action you indicate occurs. The user then has 5 minutes to answer Yes, No, Yes to all, or No to all. By default, it will be denied unless the user decides otherwise. See Query User for more details. (Query User options are not available for UNIX rules.)

Text used to query user If you are configuring a Query User rule or a deny rule, the text you type into this field is the same text that will appear in the Query User pop-up box to explain what is occurring on the system to the user.

- **Deny** Select this action type to create an application control rule that stops the application you specify from running on systems. (When you select Deny for this rule, if the user attempts to run the application in question, he/she is notified with a pop-up box explaining that the application is forbidden to run.

Step 5 and

- **Log** Enable this checkbox to turn logging on for this rule. Generally, you will want to turn logging on for all deny rules. This means that the denied system action in question is logged and sent to the server at regular time intervals.
- **Take precedence over similar <action type> rules** Enable this checkbox to manipulate policy rule precedence so that this rule is evaluated before other similar rules. You should generally NOT require this checkbox. Do not use it without understanding how it works. See [Writing Rules: Manipulating Precedence, page 4-8](#) for details.

Step 6 when

Current applications in any of the selected classes If you want to control an application (allow or deny) running on a system no matter how it is invoked, allow "All Applications" to remain selected by default. (Then you will select the application you want to control from the second Application class list.)

If you want to control which application(s) can invoke other applications, select one or more preconfigured application classes here to indicate the application that is doing the invoking (such as Network Applications).

(When your rule is configured, currently selected application classes appear at the top of the list. See the "[Configuring Static Application Classes](#)" section on page 6-6 for configuration details.)

attempt to run

New applications in any of the selected classes If you are controlling which applications can invoke other applications, this second field indicates the application class that you do not want to run when invoked by the application you chose in the top field.

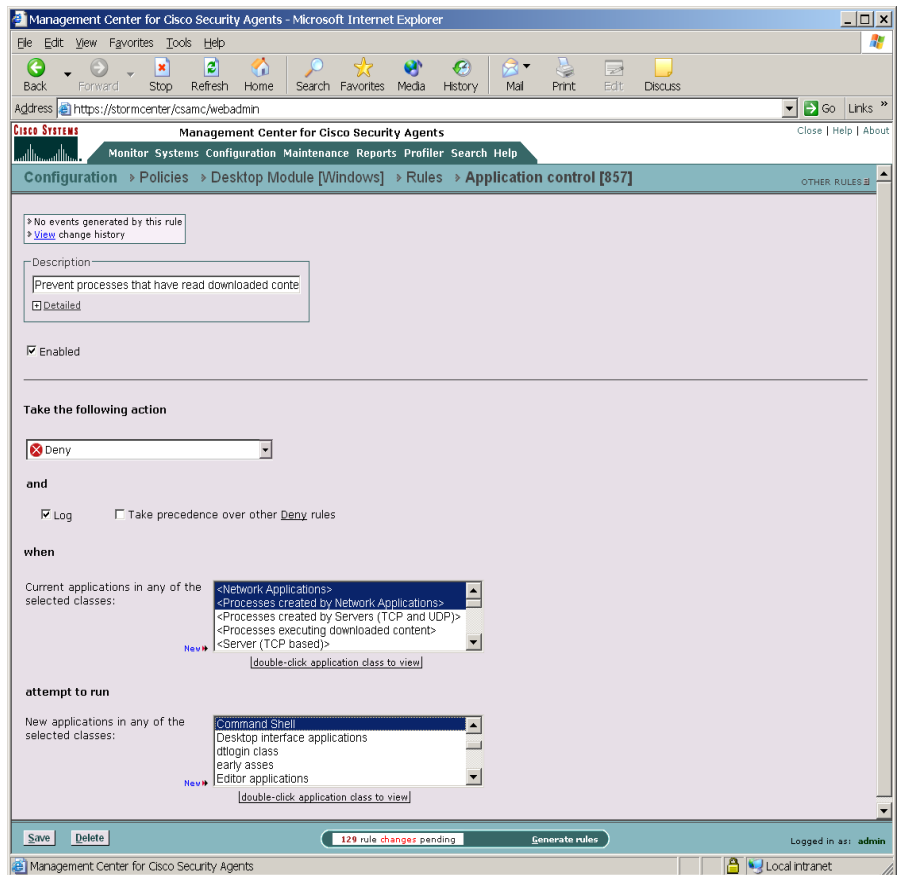
If you selected "All Applications" in the top application field, you cannot select All Applications in this second field. If you did so, all applications would be completely prevented from running on systems if this is a deny rule.



Note Most dynamic application classes are not available in this second application class field.

- Step 7** When you are finished configuring your Application control rule, click the **Save** button. This rule is now part of your new policy. It takes effect when the policy is attached to a group and then deployed to an agent on the network.

Figure 4-9 Application Control Rule



Connection Rate Limit

Use the connection rate limit rule to control the number of network connections that can be sent or received by systems within a specified time frame. This is useful in preventing attacks aimed at bringing down system services, e.g. denial of service attacks (server connection rating limiting). This is also useful in preventing the propagation of denial of service attacks (client connection rate limiting).

**Note**

These instructions are a continuation of [Configuring Policies, page 4-13](#).

Click the **Modify rules** link at the top of the Policy page to go to the Rules page.

- Step 1** To add rules to your policy, click the **Add rule** link at the bottom of the rule list. A pop-up list of the available rule types appears.
- Step 2** Select the **Connection rate limit** rule. This takes you to the configuration view for this rule type.
- Step 3** Enter the following information for the rule:
- **Description** Enter a description of this rule. This description appears in the list view for the policy. Optionally, expand the **+Detailed** field to enter a longer description.
 - **Enabled** Use this checkbox to enable this rule within the policy. (It is enabled by default.)
By not selecting this checkbox, you can save this rule, but it will not be active in the policy and it will not be distributed to groups.
 - **Log** Use this checkbox to enable logging within the policy.
- Step 4** **Limit to <100> network connections**
in <5> minutes

Reasonable values are entered into these fields by default. They define the number of connections that can normally be expected during a time frame from either specific hosts or all hosts. If the limit is exceeded in this time frame (abnormal amount of connections that could represent an attack of the system), subsequent connection requests are dropped. (The dropped connections can be those received to/from individual “specific” hosts or to/from all hosts. This setting is configured at the bottom of the page.)

Step 5 When Applications in any of the selected classes

Select *one or more* preconfigured application classes here to indicate the application(s) whose connection rate access you want to exercise control over.

Note that the entry, <All Applications>, is selected by default. You can use this default or you can unselect it and create your own application classes. For application class configuration details, see [Chapter 6, “Using Application Classes.”](#)



Note When your rule is configured, currently selected application classes appear at the top of the list.

Step 6 Attempt to act as a—Select server or client

From the pulldown menu, select **server** or **client** depending on the *direction* of the connection you are controlling.

If you are limiting a server’s connection limit, select server here. If you are limiting a client connection, select client here.

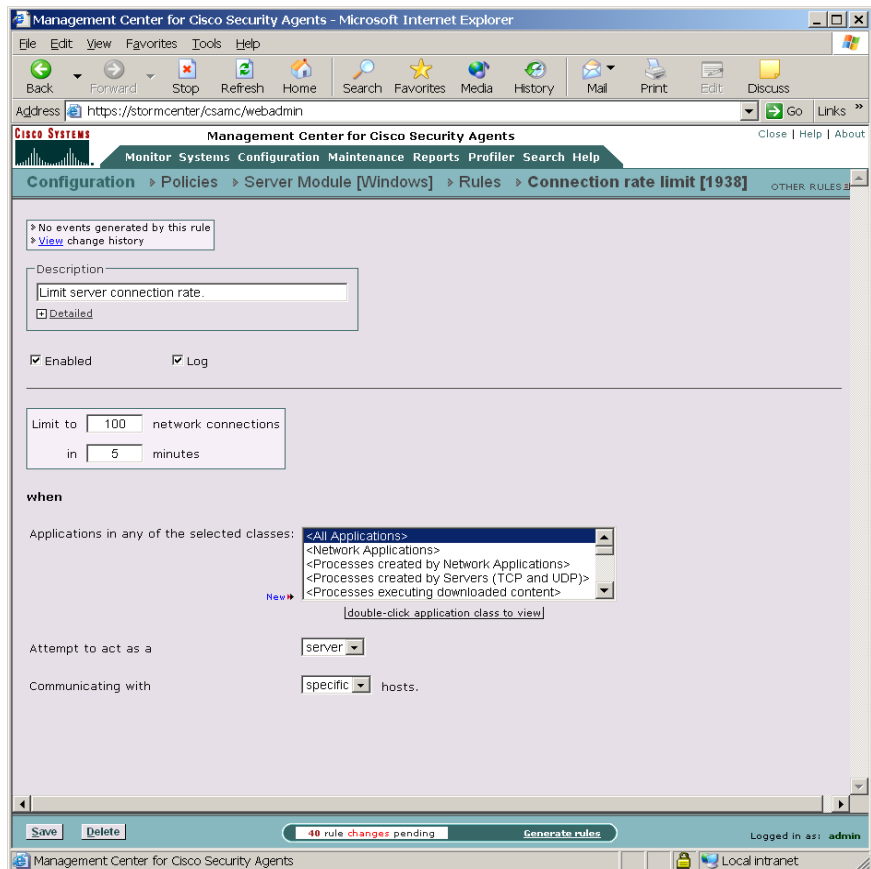
Step 7 Communicating with—Select specific hosts or all hosts

When the rate limit set here is reached, you can determine whether all subsequent service requests are dropped or only those received or sent by a specific host. If you select a “specific” host, this indicates that the host in question exceeded the rate limit. If you select all hosts, this indicates that the sum total of to and from all hosts exceeds the limit and all hosts are blocked.

Step 8 When you are finished configuring your connection rate limit rule, click the **Save** button.

This rule is now part of your new policy. It takes effect when the policy is attached to a group and then downloaded by an agent on the network.

Figure 4-10 Connection Rate Limit Rule



Data Access Control

Use data access control rules on Web servers to detect clients making malformed web server requests where such requests could crash or hang the server. A malformed request could also be an attempt by an outside client to retrieve configuration information from the web server or to run exploited code on the server. This rule detects and stops such web server attacks by examining the URI portion of the HTTP request.

An HTTP request consists of:

- the request method (a “get” or a “post”)
- the request URI (Uniform Resource Identifier—This includes the URL and related request parameters and arguments)
- the HTTP version (for example, HTTP/1.0)
- the HTTP header

The Data access control rule examines patterns in the URI portion of the HTTP request. The pre-configured Data sets (see “[Data Sets](#)” section on page 7-3) group patterns to match based upon

- functional associations of meta-characters (e.g. "(" and ")")
- examples of known classes of attacks
- Web server specific exploits

Use the data access control rule to allow or deny specified underlying network data requests for the following web servers and platforms:

- Microsoft IIS (Windows platforms)
- Apache (Windows and UNIX platforms, versions 1.3, 2.0)
- IPlanet (UNIX platforms, version 6.0)

**Caution**

On Windows platforms, if you install Web server software (IIS or Apache) after you install the Cisco Security Agent on the server system, or if you have installed the Web server in a directory other than the default, you must manually install the Cisco Security Agent data filter in order to use Data access control rules on the system in question.

If your Web server software is already installed (in its default directory) when you install the agent on Windows, the server software is detected by the agent and the data filter capability is automatically installed with the agent.

On Solaris, in order to use Data access control rules (on Apache or IPlanet servers) you must install the data filter manually after you install the Cisco Security Agent. Unlike Windows, the Solaris installation does not detect Web server software and does not install the data filter with the agent. You must always manually install it.

See the [“Manual Agent Data Filter Installation” section on page 10-10](#) for instructions.

**Note**

These instructions are a continuation of [Configuring Policies, page 4-13](#).

Click the **Modify rules** link at the top of the Policy page to go to the Rules page.

- Step 1** To add rules to your policy, click the **Add rule** link at the bottom of the rule list. A pop-up list of the available rule types appears.
- Step 2** Select the **Data access control** rule. This takes you to the configuration view for this rule type.
- Step 3** Enter the following information for the rule:
- **Description** Enter a description of this rule.
This description appears in the list view for the policy. Optionally, expand the **+Detailed** field to enter a longer description.
 - **Enabled** Use this checkbox to enable this rule within the policy. (It is enabled by default.)
By not selecting this checkbox, you can save this rule, but it will not be active in the policy and it will not be distributed to groups.

- Step 4 Take the following action**—Select an action type from the pulldown list. For further details on using allow versus deny rules, see [Writing Rules: Allow vs. Deny](#), page 4-6.
- **Add process to application class** Use this for defining dynamic application classes. See [Chapter 6, “Using Application Classes”](#) for details.
 - **High Priority Deny** Select this action type to create a deny rule that takes precedence over all other allow, deny, and query rules. For example, if you configure an allow rule that conflicts with this high priority deny, the high priority deny always takes precedence. Under the same circumstances, if the deny is not high priority and a conflicting allow rule exists in the same policy, the allow takes precedence.
 - **Allow** Select this action type to create a data access control rule that allows the applications you specify to perform the selected operation on the files you indicate. Because the default action of all policies is "allow", generally, you'll only want to configure allow rules as exceptions to existing deny rules within policies.
 - **Query User (Default Allow)** Select this action type to prompt the user when the action you indicate occurs. The user then has 5 minutes to answer Yes, No, Yes to all, or No to all. By default, it will be allowed unless the user decides otherwise. See [page 4-10](#) for details. (Query User options are not available for UNIX rules.)
 - **Query User (Default Deny)** Select this action type to prompt the user when the action you indicate occurs. The user then has 5 minutes to answer Yes, No, Yes to all, or No to all. By default, it will be denied unless the user decides otherwise. See [page 4-10](#) for details. (Query User options are not available for UNIX rules.)

Text used to query user If you are configuring a Query User rule, the text you type into this field is the same text that will appear in the Query User pop-up box to explain what is occurring on the system to the user.
 - **Deny** Select this action type to create a data access control rule that stops the application you specify from performing the selected operation on the files you indicate.

Step 5 and

- **Log** Enable this checkbox to turn logging on for this rule. Generally, you will want to turn logging on for all deny rules. This means that the denied system action in question is logged and sent to the server at regular time intervals.
- **Take precedence over similar <action type> rules** Enable this checkbox to manipulate policy rule precedence so that this rule is evaluated before other similar rules. You should generally NOT require this checkbox. Do not use it without understanding how it works. See [Writing Rules: Manipulating Precedence, page 4-8](#) for details.

Step 6 **When—Applications in any of the selected classes**

Select one or more preconfigured application classes here to indicate the application(s) whose access to the listed data sets you want to exercise control over.

Note that the entry, <All Applications>, is selected by default. You can use this default or you can unselect it and create your own application classes. For application class configuration details, see [Chapter 6, “Using Application Classes.”](#)



Note When your rule is configured, currently selected application classes appear at the top of the list.

Step 7 **Attempt to access these data sets**

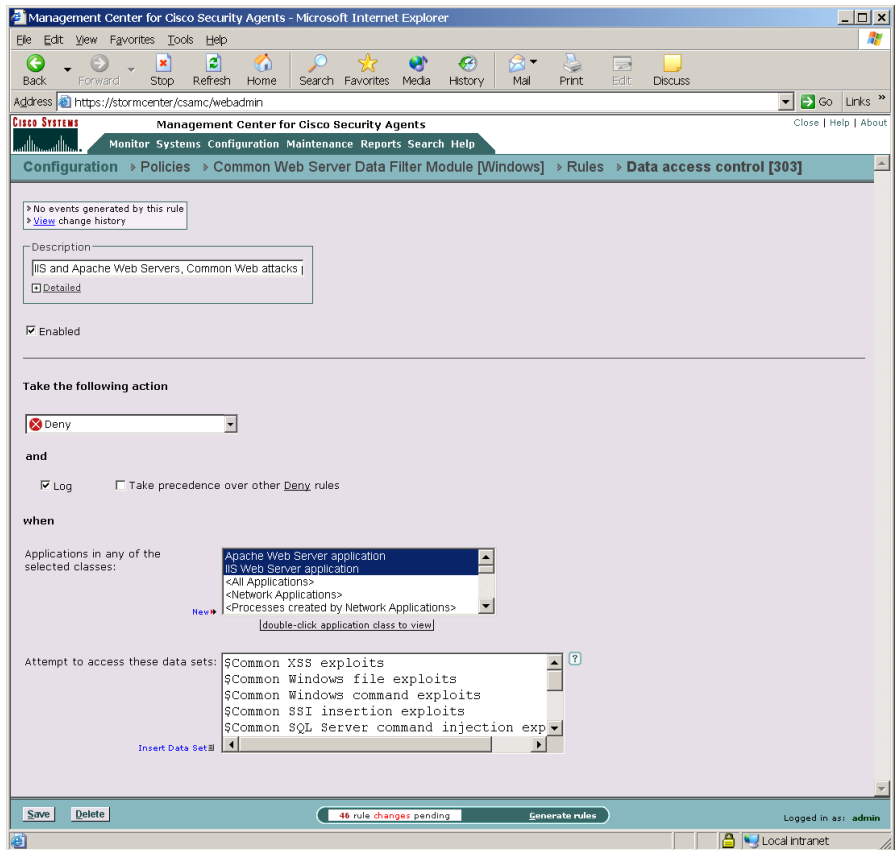
Click the **Insert Data Set** link to enter a pre-configured data set here. When you click this link, a list of the Data Sets you've configured appears here, allowing you to select one or more. Instead of data sets, you can list the literal data strings you want to protect. You can use a wildcard designation.

For information on configuring Data Sets, see the [“Data Sets” section on page 7-3](#).

Step 8 When you are finished configuring your Data access control rule, click the **Save** button.

This rule is now part of your new policy. It takes effect when the policy is attached to a group and then downloaded by an agent on the network.

Figure 4-11 Data Access Control Rule



File Access Control

Use file access control rules to allow or deny what operations (read, write) selected applications can perform on files.



Note

These instructions are a continuation of [Configuring Policies, page 4-13](#).

Click the **Modify rules** link at the top of the Policy page to go to the Rules page.

Step 1 To add rules to your policy, click the **Add rule** link at the bottom of the rule list. A pop-up list of the available rule types appears.

Step 2 Select the **File access control** rule. This takes you to the configuration view for this rule type.

Step 3 Enter the following information for the rule:

- **Description** Enter a description of this rule. This description appears in the list view for the policy. Optionally, expand the **+Detailed** field to enter a longer description.
- **Enabled** Use this checkbox to enable this rule within the policy. (It is enabled by default.)
By not selecting this checkbox, you can save this rule, but it will not be active in the policy and it will not be distributed to groups.

Step 4 **Take the following action**—Select an action type from the pulldown list. For further details on using allow versus deny rules, see [Writing Rules: Allow vs. Deny, page 4-6](#).

- **Add process to application class** Use this for defining dynamic application classes. See [Chapter 6, “Using Application Classes”](#) for details.
- **High Priority Deny** Select this action type to create a deny rule that takes precedence over all other allow, deny, and query rules.
For example, if you configure an allow rule that conflicts with this high priority deny, the high priority deny always takes precedence. Under the same circumstances, if the deny is not high priority and a conflicting allow rule exists in the same policy, the allow takes precedence.
- **Allow** Select this action type to create a file access control rule that allows the applications you specify to perform the selected operation on the files you indicate.
Because the default action of all policies is "allow", generally, you'll only want to configure allow rules as exceptions to existing deny rules within policies.
- **Query User (Default Allow)** Select this action type to prompt the user when the action you indicate occurs. The user then has 5 minutes to answer Yes, No, Yes to all, or No to all.
By default, it will be allowed unless the user decides otherwise. See [page 4-10](#) for details. (Query User options are not available for UNIX rules.)

- **Query User (Default Deny)** Select this action type to prompt the user when the action you indicate occurs. The user then has 5 minutes to answer Yes, No, Yes to all, or No to all.

By default, it will be denied unless the user decides otherwise. See [page 4-10](#) for details. (Query User options are not available for UNIX rules.)

Text used to query user If you are configuring a Query User rule, the text you type into this field is the same text that will appear in the Query User pop-up box to explain what is occurring on the system to the user.

- **Deny** Select this action type to create a file access control rule that stops the application you specify from performing the selected operation on the files you indicate.

Step 5 and

- **Log** Enable this checkbox to turn logging on for this rule. Generally, you will want to turn logging on for all deny rules. This means that the denied system action in question is logged and sent to the server at regular time intervals.
- **Take precedence over similar <action type> rules** Enable this checkbox to manipulate policy rule precedence so that this rule is evaluated before other similar rules. You should generally NOT require this checkbox. Do not use it without understanding how it works. See [Writing Rules: Manipulating Precedence, page 4-8](#) for details.

Step 6 When—Applications in any of the selected classes

Select one or more preconfigured application classes here to indicate the application(s) whose access to the listed files you want to exercise control over.

Note that the entry, <All Applications>, is selected by default. You can use this default or you can unselect it and create your own application classes. For application class configuration details, see [Chapter 6, “Using Application Classes.”](#)



Note When your rule is configured, currently selected application classes appear at the top of the list.

Step 7 Attempt the following operations

Select the operations **Read** and/or **Write** you are allowing/denying on the files named in the Files field.

Step 8 On any of these files

Click the **Insert File Set** link to enter a pre-configured file set here. When you click this link, a list of the File Sets you've configured appears here, allowing you to select one or more. Instead of file sets, you can list the literal files you want to protect, using the file paths (including wildcards).

For local system paths, you must specify the disk drive.

You can use a wildcard designation.

For information on entering file path literals here rather than using pre-configured File Sets, see the [“Using the Correct Syntax” section on page 2-18](#).

File access control rules apply to files, not directories. You must make some file specification. A wildcard is acceptable to specify all files in a named directory. For example:

Windows:

```
*:\Program Files\winnt\*
or @system\** (this indicates all files below the system directory)
```

UNIX:

```
/etc/passwd
```



Note You can protect directory paths as well as files on UNIX systems (You cannot do this on Windows.) See the [“Using the Correct Syntax” section on page 2-18](#) for details.

For network machines (Windows only), enter

```
\\<machine name>\<share>\<path>\<filename>
```

For example: `\\Backup_Server\finance\records\database.db`

You can enter more than one file path, but each entry must appear on its own line. For File Set configuration details, see the [“File Sets” section on page 7-6](#).

**Caution**

Symbolic Links: For UNIX, if you create a File access control rule to protect a symbolic link, **ONLY** that symbolic link is protected. The underlying resource, unless also specified, is **NOT** protected. For example, a File access control rule written for `/etc/hosts` does not protect `/etc/inet/hosts`. Similarly, a File access control rule written for `/etc/inet/hosts` does not protect `/etc/hosts`. If you want to protect a symbolic link and its underlying resource, both must be specified in the rule.

**Note**

Use **@dynamic** in the File set text field to indicate all files that have been quarantined by CSA MC as a result of correlated email worm events or correlated virus scanner log messages. Files are quarantined by CSA MC for up to one hour. This list updates automatically as logged quarantined files are received.

Files of type `.doc`, `.vbs`, `.vba`, and `.js` are a special case. To prevent denial of service attacks, these files can be added to the `@dynamic` list whether or not they are seen as downloaded content. Other file types are only quarantined if they have just been downloaded. See the [“Global Events” section on page 5-18](#).

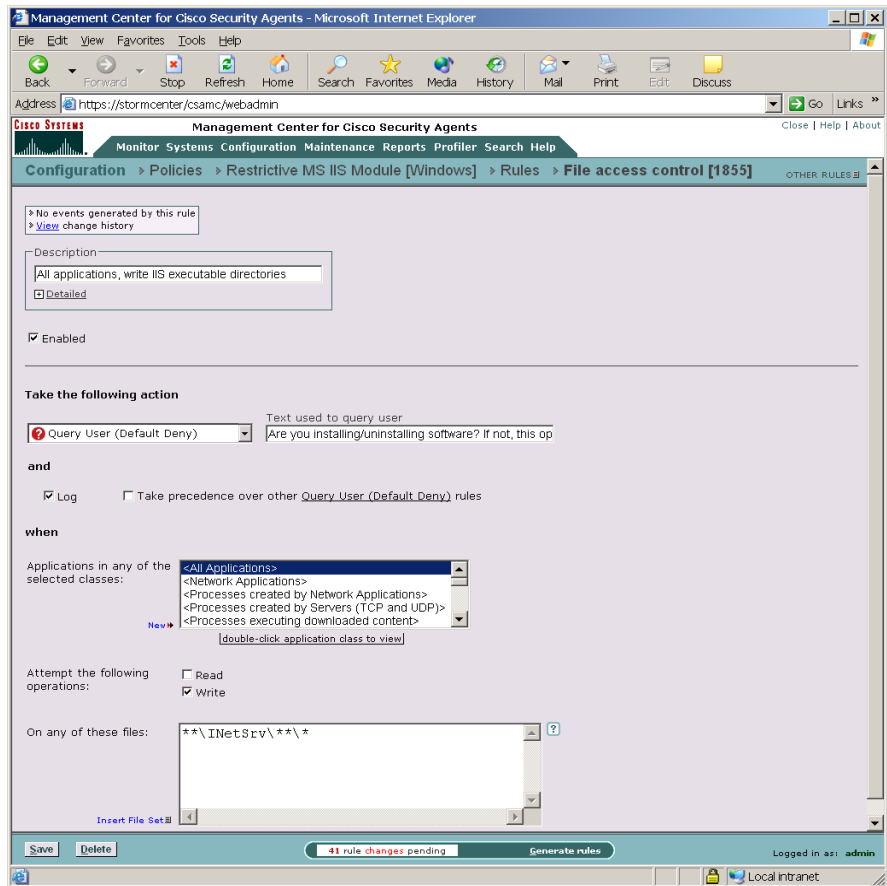
Step 9 When you are finished configuring your File access control rule, click the **Save** button.

This rule is now part of your new policy. It takes effect when the policy is attached to a group and then downloaded by an agent on the network.

**Caution**

In order to distribute rules to the correct hosts, you must associate policies with groups and then Generate rules. See [page 4-92](#) and [page 4-97](#) for instructions.

Figure 4-12 File Access Control Rule



File Monitor

Use the file monitor rule to track read and/or write access to a specified file or files. Unlike a file access control rule, a file monitor rule will simply log an event when a user attempts read or write access of a file. With this rule, you are not allowing or denying access to this file. All access to files specified as part of a file monitor rule is allowed (unless denied by another rule). But this rule lets you monitor that access by generating log entries.

You do not have this capability by configuring an "allow-log" file access control rule. With file access control rules, if you deny reading, you are also denying writing and if you allow writing, you are also allowing reading. You cannot distinguish between these two actions in some cases. A file monitor rule permits and logs this distinction.

**Note**

These instructions are a continuation of [Configuring Policies, page 4-13](#).

- Step 1** To add rules to your policy, click the **Add rule** link at the bottom of the rule list. A pop-up list of the available rule types appears.
- Step 2** Select the **File monitor** rule. This takes you to the configuration view for this rule type.
- Step 3** Enter the following information:
- **Description** Enter a description of this rule. Optionally, expand the **+Detailed** field to enter a longer description.
 - **Enabled** Use this checkbox to enable this rule within the policy. It is enabled by default.
By not selecting this checkbox, you can save this rule, but it will not be active in the policy and it will not be distributed to groups.
- Step 4** **Log an event When—Applications in any of the selected classes:**
- Select one or more preconfigured application classes here to indicate the application(s) whose access to the listed files you want to monitor. Note that the entry, <All Applications>, is selected by default. You can use this default or you can unselect it and create your own application classes. For application class configuration details, see [Chapter 6, “Using Application Classes.”](#)

**Note**

When your rule is configured, currently selected application classes appear at the top of the list.

- Step 5** **Attempt the following operations:**
- Select the operations **Read** and/or **Write** you are monitoring.

Step 6 On any of these files:

Click the **Insert File Set** link to enter a pre-configured file set here. When you click this link, a list of the File Sets you've configured appears here, allowing you to select one or more. Instead of file sets, you can list the literal files you want to protect, using the file paths (including wildcards).

For local system paths, you must specify the disk drive.

You can use a wildcard designation.

For information on entering file path literals here rather than using pre-configured File Sets, see the [“Using the Correct Syntax” section on page 2-18](#).



Note File monitor rules apply to files, not directories. You must make some file specification. A wildcard is acceptable to specify all files in a named directory.

Windows:

For example: `*:\Program Files\winnt*`

or `@system**` (this indicates all files below the system directory)

UNIX:

For example: `/etc/passwd`

For network machines (Windows only), enter

`\\<machine name>\<share>\<path>\<filename>`

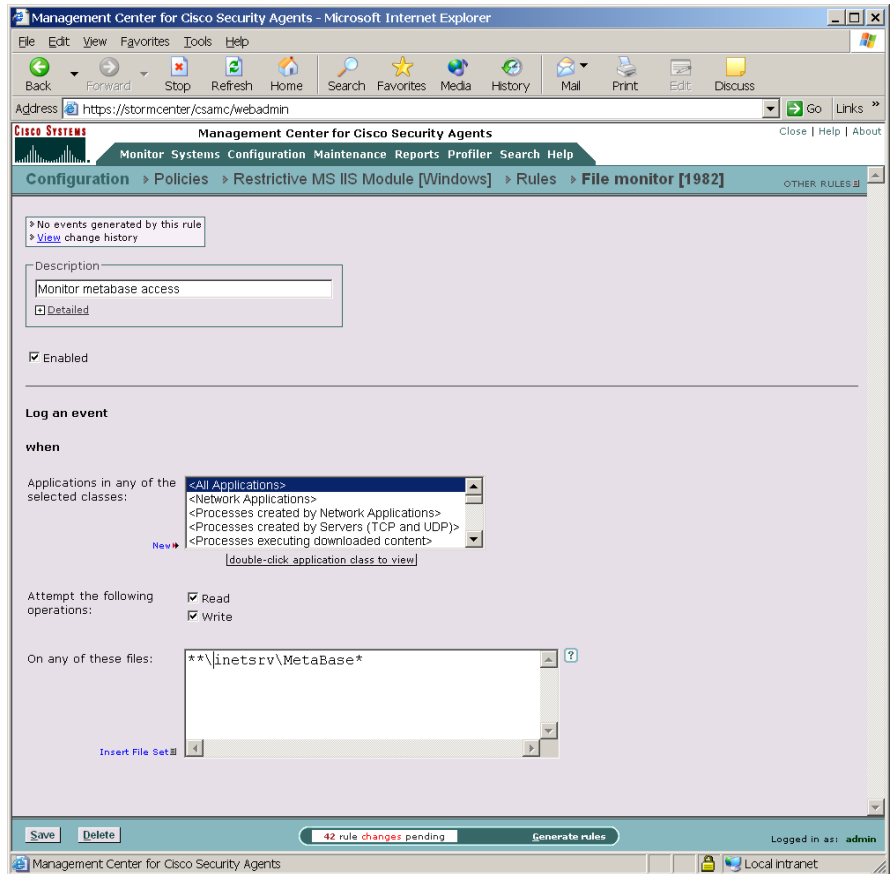
For example: `\\Backup_Server\finance\records\database.db`

You can enter more than one file path, but each entry must appear on its own line. For File Set configuration details, see the [“File Sets” section on page 7-6](#).

Step 7 When you are finished configuring your File monitor rule, click the **Save** button.

This rule is now part of your new policy. It takes effect when the policy is attached to a group and then downloaded by an agent on the network.

Figure 4-13 File Monitor Rule



Network Access Control

Use network access control rules to control access to specified network services and network addresses.

**Note**

The following instructions are a continuation of [Configuring Policies, page 4-13](#).

- Step 1** To add rules to your policy, click the **Add rule** link at the bottom of the rule list. A pop-up list of the available rule types appears.
- Step 2** Select the **Network Access Control** rule. This takes you to the configuration view for this rule type.
- Step 3** Enter the following information:
- **Description** Enter a description of this rule. This description appears in the list view for the policy. Optionally, expand the **+Detailed** field to enter a longer description.
 - **Enabled** Use this checkbox to enable this rule within the policy. (It is enabled by default.) By not selecting this checkbox, you can save this rule, but it will not be active in the policy and it will not be distributed to groups.

- Step 4 Take the following action**—Select an action type from the pulldown list. For further details on using allow versus deny rules, see [Writing Rules: Allow vs. Deny, page 4-6](#).
- **Add process to application class** Use this for defining dynamic application classes. See [Chapter 6, “Using Application Classes”](#) for details.
 - **High Priority Deny** Select this action type to create a deny rule that takes precedence over all other allow, deny, and query rules. For example, if you configure an allow rule that conflicts with this high priority deny, the high priority deny always takes precedence. Under the same circumstances, if the deny is not high priority and a conflicting allow rule exists in the same policy, the allow takes precedence.
 - **Allow** Select this action type to create a network access control rule that allows the applications you specify to perform the operation you select with the addresses you indicate. Because the default action of all policies is "allow", generally, you'll only want to configure allow rules as exceptions to existing deny rules within policies.
 - **Query User (Default Allow)** Select this action type to prompt the user when the action you indicate occurs. The user then has 5 minutes to answer Yes, No, Yes to all, or No to all. By default, it will be allowed unless the user decides otherwise. See [page 4-10](#) for details. (Query User options are not available for UNIX rules.)
 - **Query User (Default Deny)** Select this action type to prompt the user when the action you indicate occurs. The user then has 5 minutes to answer Yes, No, Yes to all, or No to all. By default, it will be denied unless the user decides otherwise. See [page 4-10](#) for details. (Query User options are not available for UNIX rules.)

Text used to query user If you are configuring a Query User rule, the text you type into this field is the same text that will appear in the Query User pop-up box to explain what is occurring on the system to the user.
 - **Deny** Select this action type to create a network access control rule that stops the application you indicate from performing the operation you select on the specified addresses.

Step 5 and

- **Log** Enable this checkbox to turn logging on for this rule. Generally, you will want to turn logging on for all deny rules. This means that the denied system action in question is logged and sent to the server at regular time intervals.
- **Take precedence over similar <action type> rules** Enable this checkbox to manipulate policy rule precedence so that this rule is evaluated before other similar rules. You should generally NOT require this checkbox. Do not use it without understanding how it works. See [Writing Rules: Manipulating Precedence, page 4-8](#) for details.

Step 6 When—Applications in any of the selected classes

Select one or more preconfigured application classes here to indicate the application(s) whose access to the listed services and addresses you want to exercise control over.

Note that the entry, <All Applications>, is selected by default. You can use this default or you can unselect it and create your own application classes. For application class configuration details, see [Chapter 6, “Using Application Classes.”](#)

Step 7 Attempt to act as a—Select server or client

From the pulldown menu, select **server** or **client** depending on the direction of the connection you are controlling.

If you are limiting a server’s contact with clients, select server here and enter the client(s) address in the host addresses field. If you are limiting a client’s contact with a server, select client here and enter the server(s) address in the host addresses field.

Step 8 for network services

Enter the literal protocol/port number combination for the service you want to control access to or click the **Insert Network Service** link to enter a pre-configured network service variable here. When you click this link, a list of the Network Service Variables you've configured appears here, allowing you to select one or more.

This field refers to either a server providing this service or a client accessing this service. For Network Service configuration details, see the [“Network Services” section on page 7-14.](#)

Step 9 Communicating with host addresses

Enter the literal network address(es) for the client/servers you want to control access to or click the **Insert Network Address Set** link to enter a pre-configured network address set variable here.

If you select server in the previous pulldown list, you enter client addresses here. If you select client in the previous pulldown list, you enter server addresses here. Note that you can use Network Address Set variables.

- You also can use the following "short hand" entry in Network Address Sets and in Network Access Control rules to indicate all local addresses on the agent system in question. The @symbol must appear at the start of the short hand name.

Use **@local** to indicate all local addresses on the agent system. You would want to use this if you are allowing different applications on a single system to talk to each other without opening these applications up to network access.

Step 10 Using these local addresses

Enter the literal network address(es) for the local system addresses you want to control (i.e. control clients making connections from or control servers making connections to). You can also click the **Insert Network Address Set** link to enter a pre-configured network address set variable here.

The addresses or address ranges you enter here are used to control the host initiating the network connection. For example, you could write a Network access control rule which would only allow laptop users to connect to an internal network database if their connection is coming through a VPN (i.e. machine using an allowed/disallowed address to make a connection, incoming or outgoing). If the connection attempt comes in through an ISP-assigned address that is not part of this rule, it would not be allowed.

You could also use this field to impose a restriction that only trusted addresses can read an internal server. If the connection is received from an internal system or via a VPN from a fixed, trusted address, it is allowed.

Use **@local** (the default) to indicate all local addresses on the agent system. You would want to use this if you are allowing different applications on a single system to talk to each other without opening these applications up to network access (intra-box connections).

Step 11 When you are finished configuring your Network access control rule, click the **Save** button.

This rule is now part of your policy. It takes effect when the policy is attached to a group and then downloaded by an agent on the network. You should note that new rules only apply to new connections. See the “[Preserving Application Process Classes](#)” section on page 6-6 for details.

**Caution**

In order to distribute rules to the correct hosts, you must associate policies with groups and then Generate rules. See [page 4-92](#) and [page 4-97](#) for instructions.

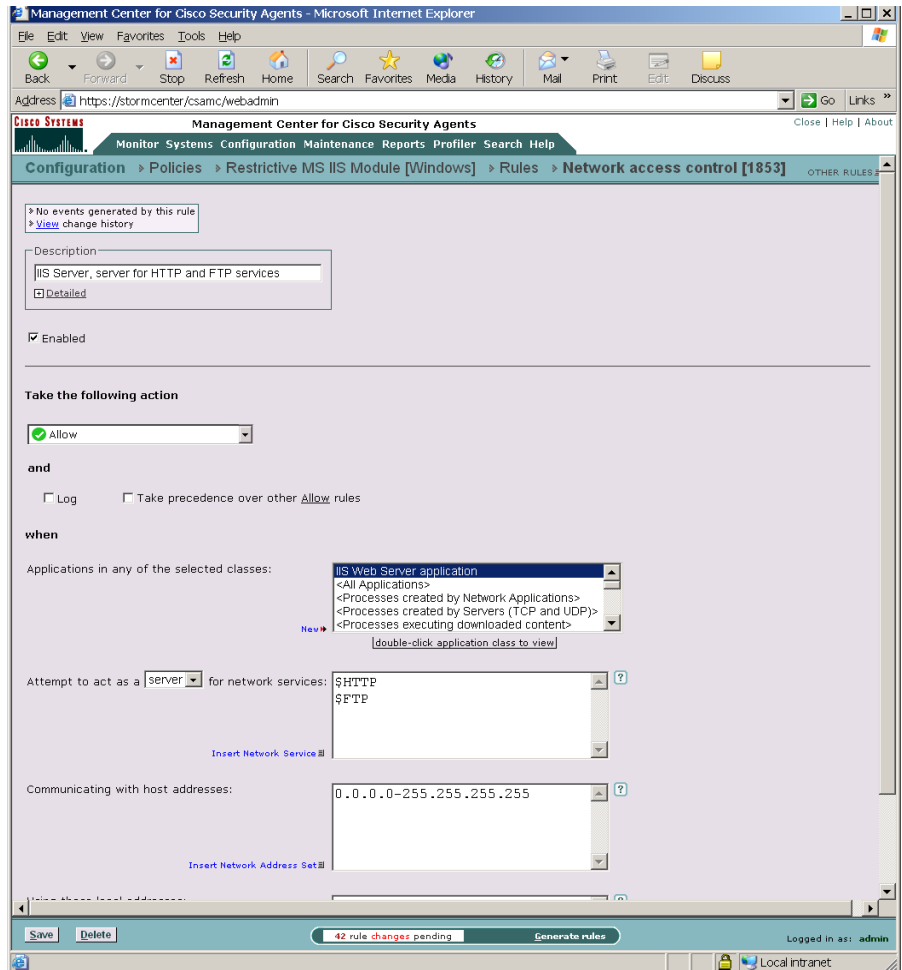
**Note**

You can use multiple access control rules in one policy.

**Note**

No network access control rule denial events are logged for any UDP port resulting from multicast packet signals. (If a collection of hosts have the same network access control rule and a broadcast such as UDP/138 were denied, then event messages would inundate CSA MC.)

Figure 4-14 Network Access Control Rule



Windows Only Rules

The following rules are only available for both Windows policies.

COM Component Access Control

Use COM component access control rules to allow or deny applications from accessing specified COM components. COM is the Microsoft Component Object Model, the technology that allows objects to interact across process and machine boundaries as easily as within a single process. Each of the Microsoft Office applications (Word, Excel, Powerpoint, etc.)

exposes an "Application" COM component which can be used to create macros or utility scripts. While this is useful functionality, it can be used maliciously by an inadvertently downloaded Visual Basic script.

An example would be the ILOVEYOU virus, which propagated by using the "Outlook.Application" COM component to send itself to each entry in the local address book. Using the COM component access control rule, you can protect specific COM components. For example, you could create a rule which limits access to Office components (Word.*, Outlook.*, Excel.*, etc.) only to the Office applications themselves. Non-Office applications (such as the Visual Basic scripting engine) would therefore be denied access to these components.

**Note**

CSA MC provides a COM component import utility which installs with each Cisco Security Agent. Running this utility extracts all COM component PROGID's and CLSID's for software running on the system in question. See the ["Using the COM Extract Utility" section on page 10-9](#) for information.

These instructions are a continuation of [Configuring Policies, page 4-13](#).

-
- Step 1** To add rules to your policy, click the **Add rule** link at the bottom of the rule list. A pop-up list of the available rule types appears.
- Step 2** Select the **COM component access control** rule. This takes you to the configuration view for this rule type (see [Figure 4-19](#)).

**Note**

This rule type is not available for UNIX policies.

- Step 3** Enter the following information
- **Description** Enter a description of this rule. This description appears in the list view for the policy. Optionally, expand the **+Detailed** field to enter a longer description.
 - **Enabled** Use this checkbox to enable this rule within the policy. (It is enabled by default.)
By not selecting this checkbox, you can save this rule, but it will not be active in the policy and it will not be distributed to groups.
- Step 4** **Take the following action**—Select an action type from the pulldown list. For further details on using allow versus deny rules, see [Writing Rules: Allow vs. Deny, page 4-6](#).
- **Add process to application class** Use this for defining dynamic application classes. See [Chapter 6, “Using Application Classes”](#) for details.
 - **High Priority Deny** Select this action type to create a deny rule that takes precedence over all other allow, deny, and query rules. For example, if you configure an allow rule that conflicts with this high priority deny, the high priority deny always takes precedence. Under the same circumstances, if the deny is not high priority and a conflicting allow rule exists in the same policy, the allow takes precedence.
 - **Allow** Select this action type to create a COM component access control rule that allows the applications you specify to perform the selected operation on the COM components you indicate. Because the default action of all policies is "allow", generally, you'll only want to configure allow rules as exceptions to existing deny rules within policies.
 - **Query User (Default Allow)** Select this action type to prompt the user when the action you indicate occurs. The user then has 5 minutes to answer Yes, No, Yes to all, or No to all. By default, it will be allowed unless the user decides otherwise. See [page 4-10](#) for details.

- **Query User (Default Deny)** Select this action type to prompt the user when the action you indicate occurs. The user then has 5 minutes to answer Yes, No, Yes to all, or No to all.

By default, it will be denied unless the user decides otherwise. See [page 4-10](#) for details.

Text used to query user If you are configuring a Query User rule, the text you type into this field is the same text that will appear in the Query User pop-up box to explain what is occurring on the system to the user.

- **Deny** Select this action type to create a COM component access control rule that stops the application you specify from performing the selected operation on the COM components you indicate.

Step 5 and

- **Log** Enable this checkbox to turn logging on for this rule. Generally, you will want to turn logging on for all deny rules. This means that the denied system action in question is logged and sent to the server at regular time intervals.
- **Take precedence over similar <action type> rules** Enable this checkbox to manipulate policy rule precedence so that this rule is evaluated before other similar rules. You should generally NOT require this checkbox. Do not use it without understanding how it works. See [Writing Rules: Manipulating Precedence, page 4-8](#) for details.

Step 6 When—Applications in any of the selected classes

Select one or more preconfigured application classes here to indicate the application(s) whose access to the selected COM components you want to exercise control over.

Note that the entry, <All Applications>, is selected by default. You can use this default or you can unselect it and create your own application classes. For application class configuration details, see [Chapter 6, “Using Application Classes.”](#)

Step 7 Attempt to access a COM component....matching any of the following component sets

Click the **Insert COM Component** link to select one or more pre-configured COM component sets for this rule. If you do not want to use a COM component set variable, using the correct syntax, enter a literal PROGID or CLSID (one per line) here. CSA MC provides a utility for extracting PROGID and CLSID information from systems running agent software. See the [“Using the COM Extract Utility”](#) section on page 10-9 for instructions.

PROGID's, use the following syntax:

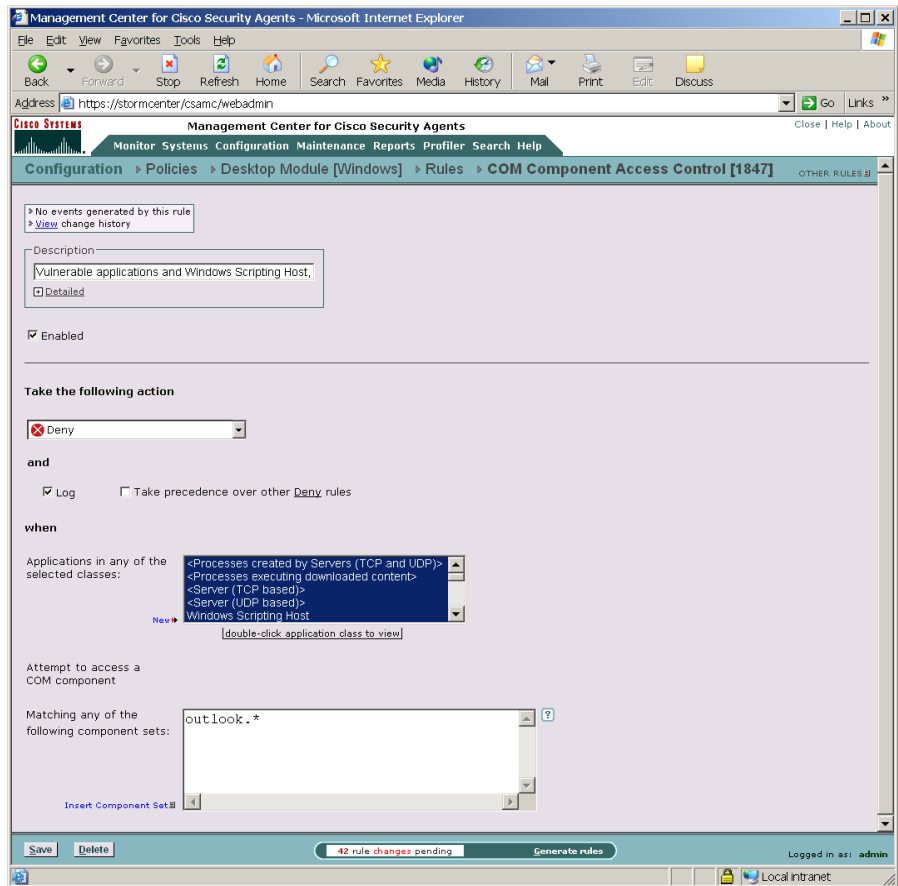
```
Outlook.Application
```

When entering CLSID's (uppercase hexadecimal) using the following syntax, you must include the brackets shown here:

```
{000209FF-0000-0000-C000-000000000046}
```

Step 8 When you are finished configuring your COM component access control rule, click the **Save** button.

Figure 4-15 COM Component Access Control Rule



File Version Control

Use the File version control rule to prevent users from running specified versions of applications on their systems. For example, if there is a known security hole in one or more versions of a particular application, this rule would prevent those specific versions from running, but would allow any versions not included in this rule to run unimpeded.

One particular example where this type of rule would be beneficial is in the case of Microsoft Security Bulletin (MS01-020). This bulletin states the following: "Because HTML e-mail messages are Web pages, Internet Explorer can render them and open binary attachments in a way that is appropriate to their MIME type. However, there is a flaw in the type of processing that is specified for certain unusual MIME types. If a malicious user creates an HTML e-mail message that contains an attachment that can be run and then modifies the MIME header information to specify that the attachment is one of the unusual MIME types that Internet Explorer handles incorrectly, Internet Explorer may run the attachment automatically when it renders the e-mail message."

Microsoft has a patch to correct this security problem, but the patch is only available for Internet Explorer 5.01 Service Pack 1 and IE 5.5. If users are running an earlier version of IE, they must upgrade to 5.01 or 5.5 and install the correct service packs and patches to correct the problem. Therefore, earlier versions of IE contain an unfixable security problem and you will want to prevent users from running these versions. The following configuration information uses the IE security bulletin as an example.

**Note**

Note that users can get around a File version control rule by copying the file in question to a different file name. Therefore you must assume that users are working in cooperation with you for these rule types to be successful. You could also create a File access control rule to prevent users from changing the application file name in question.

**Note**

These instructions are a continuation of [Configuring Policies, page 4-13](#).

Step 1

To add rules to your policy, click the **Add rule** link at the bottom of the rule list. A pop-up list of the available rule types appears.

Step 2

Select the **File version control** rule. This takes you to the configuration view for this rule type.

**Note**

This rule type is not available for UNIX policies.

- Step 3** In the File version control rule configuration view, enter the following information:
- **Description** Enter a description of this rule. This description appears in the list view for the policy. Optionally, expand the **+Detailed** field to enter a longer description.
 - **Enabled** Use this checkbox to enable this rule within the policy. (It is enabled by default.) By not selecting this checkbox, you can save this rule, but it will not be active in the policy and it will not be distributed to groups.
- Step 4** **Take the following action** (Select an action type from the pulldown list. For further details on using allow versus deny rules, see [Writing Rules: Allow vs. Deny, page 4-6.](#))
- **Add process to application class** Use this for defining dynamic application classes. See [Chapter 6, “Using Application Classes”](#) for details.
 - **High Priority Deny** Select this action type to create a deny rule that takes precedence over all other allow, deny, and query rules. For example, if you configure an allow rule that conflicts with this high priority deny, the high priority deny always takes precedence. Under the same circumstances, if the deny is not high priority and a conflicting allow rule exists in the same policy, the allow takes precedence.
 - **Allow** Select this action type to create a file version control rule that allows the applications you specify to perform the selected operation on the files you indicate. Because the default action of all policies is "allow", generally, you'll only want to configure allow rules as exceptions to existing deny rules within policies.
 - **Query User (Default Allow)** Select this action type to prompt the user when the action you indicate occurs. The user then has 5 minutes to answer Yes, No, Yes to all, or No to all. By default, it will be allowed unless the user decides otherwise. See [page 4-10](#) for details.

- **Query User (Default Deny)** Select this action type to prompt the user when the action you indicate occurs. The user then has 5 minutes to answer Yes, No, Yes to all, or No to all. By default, it will be denied unless the user decides otherwise. See [Querying the User, page 4-10](#) for more details.

Text used to query user If you are configuring a Query User rule, the text you type into this field is the same text that will appear in the Query User pop-up box to explain what is occurring on the system to the user.

- **Deny** Select this action type to create a file version control rule that stops the application version you specify from running on systems.

Step 5 and

- **Log** Enable this checkbox to turn logging on for this rule. Generally, you will want to turn logging on for all deny rules. This means that the denied system action in question is logged and sent to the server at regular time intervals.
- **Take precedence over similar <action type> rules** Enable this checkbox to manipulate policy rule precedence so that this rule is evaluated before other similar rules. You should generally NOT require this checkbox. Do not use it without understanding how it works. See [Writing Rules: Manipulating Precedence, page 4-8](#) for details.

Step 6 **When** An execution of the following

Enter the **File** you are prohibiting (You will enter the exact version in the next field.) This field accepts file entries for .exe, .dll, and .ocx files. Enter just the file name here. No path is required.

For example: `iexplore.exe`

You cannot use wildcard entries in this field.

Step 7 with version within these Version ranges:

Enter the version or version range (using a dash to indicate range) you are prohibiting of the file you entered in the previous field.

For example: 0-5.00.3314.2100
5.00.3314.2100-5.50.4522.1800

You can enter multiple, nonconsecutive ranges by entering versions on separate lines in this field.

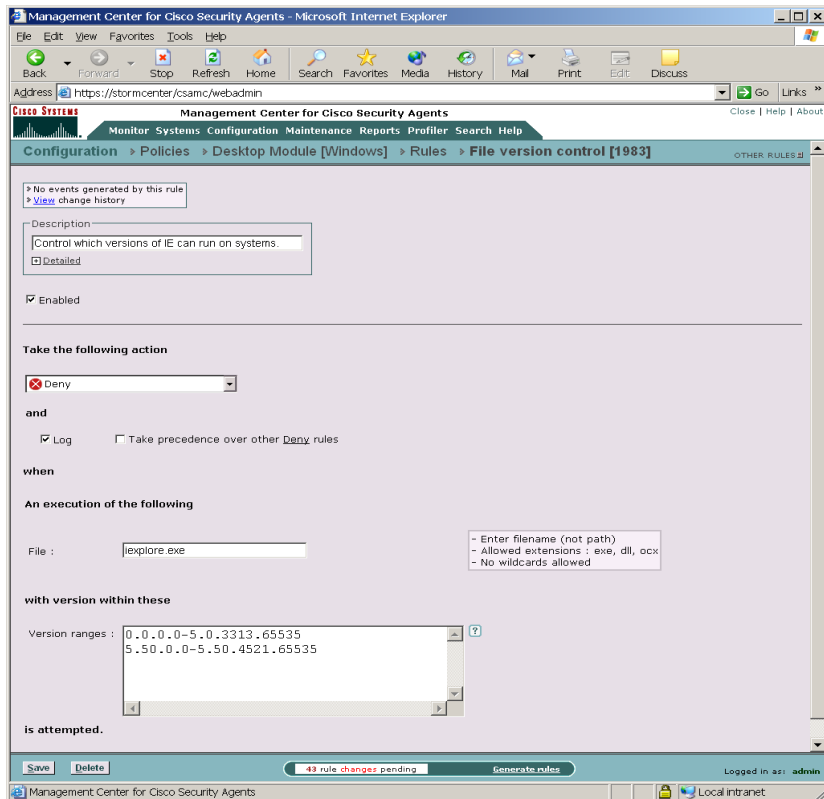
To locate the version of a file (*.exe, *.dll, or *.ocx), select the file and right click. Select **Properties**. Click the **Version** tab. The File version is normally 4 values separated by dots.

**Note**

When entering version numbers for Microsoft applications, refer to the Microsoft web site. Application version numbers accessible from the application itself sometimes correspond to slightly different version numbers in Microsoft version charts. For example, Microsoft Article number Q164539 was used to determine the version numbers for this File version control rule.

Step 8 Click the **Save** button when you are finished.

Figure 4-16 File Version Control Rule



Kernel Protection

Use the Kernel protection rule to prevent unauthorized access to the operating system. In effect, this rule prevents drivers from dynamically loading after system startup. You can specify exceptions to this rule for authorized drivers that you are allowing to load any time after the system is finished booting.

You can also use this rule to only detect unauthorized access to the operating system at any time.



Note These instructions are a continuation of [Configuring Policies, page 4-13](#).

- Step 1** To add rules to your policy, click the **Add rule** link at the bottom of the rule list. A pop-up list of the available rule types appears.
- Step 2** Select the **Kernel protection** rule. This takes you to the configuration view for this rule type.
- Step 3** Enter the following:
- **Description** Enter a description of this rule. This description appears in the list view for the policy. Optionally, expand the **+Detailed** field to enter a longer description.
 - **Enabled** Use this checkbox to enable this rule within the policy. (It is enabled by default.) By not selecting this checkbox, you can save this rule, but it will not be active in the policy and it will not be distributed to groups.

Step 4 Select one or more of the following checkboxes:

- **Control module loading after system startup**

This prevents drivers from dynamically loading after system startup. You can specify exceptions to this rule for authorized drivers that you are allowing to load any time after the system is finished booting.

- **User will be notified when any unauthorized module intercepts or modifies kernel functionality**



Caution

This part of the rule only detects unauthorized access to the operating system and does not prevent it. Upon this detection, you can impose stringent network restrictions by selecting the Restrict network connectivity... checkbox.

- **Restrict network connectivity upon detection of unauthorized module**

When unauthorized modules are detected, selecting this checkbox causes the system in question to lose network connectivity. It is essentially quarantined from the network until you investigate whether the module is harmful. If it is not harmful, you can push a new rule program out to the system allowing this module to load, at which time network connectivity will be restored. Alternatively, you can remove the module (clean the system) and push a new rule program down after the module is removed. The loading of a new rule program is the only way to restore network connectivity to the system.

Configure the edit fields in this rule using the wizard. (See the [“About the Event Management Wizard”](#) section on page 8-9 for details.) You should never type data into the fields of this rule. When an event is triggered, use the Wizard link from the event to configure an exception. Only create exceptions for actions you believe are safe. For example, virus-scanners and kernel debuggers might legitimately trigger this rule. The wizard enters module data in the following edit fields:

- **Module hashes to be excluded**

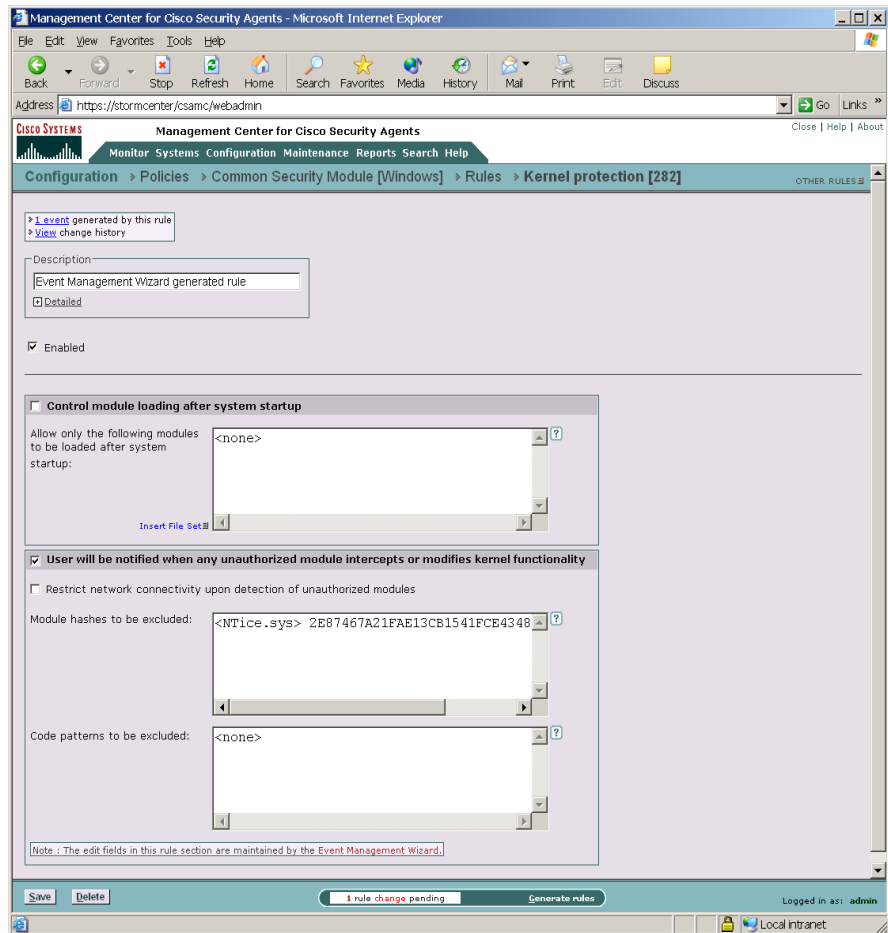
By default, this field contains <none> . The wizard enters fingerprints that identify kernel modules (e.g. drivers) into this field.

- **Code patterns to be excluded**

By default, this field contains <none> . The wizard enters code patterns (not inside any module) into this field.

Step 5 Click **Save** when finished.

Figure 4-17 Kernel protection rule



NT Event Log

Use the NT Event log rule to have specified NT Event Log items appear in the CSA MC Event Log for selected groups.



Note These instructions are a continuation of [Configuring Policies, page 4-13](#).

Step 1 To add rules to your policy, click the **Add rule** link at the bottom of the rule list. A pop-up list of the available rule types appears.

Step 2 Select the **NT Event log** rule. This takes you to the configuration view for this rule type (see [Figure 4-18](#)).



Note This rule type is not available for UNIX policies.

Step 3 Enter the following information:

- **Description** Enter a description of this rule. This description appears in the list view for the policy. Optionally, expand the **+Detailed** field to enter a longer description.
- **Enabled** Use this checkbox to enable this rule within the policy. (It is enabled by default.)
By not selecting this checkbox, you can save this rule, but it will not be active in the policy and it will not be distributed to groups.

Step 4 **Log events from the event log**

- **Include events matching the following** Select this radio button to specify the criteria for NT Event Log entries which you want to appear in the CSA MC Event Log.
- **Include all events except those matching the following** Select this radio button to specify the criteria for NT Event Log entries which you do not want to appear in the CSA MC Event Log. (All criteria not specified here will appear in the CSA MC Event Log.)



Note You can configure CSA MC to correlate NT event types logged across multiple systems. You can also correlate NT events received from virus scanners running on agent systems and quarantine contaminated files. See the [“Global Events” section on page 5-18](#).

Step 5 Criteria (You should select at least one for the rule to have any effect.)

- **Event Log Type** Select one or more checkboxes here to indicate which NT Event Log entries you want to appear (first radio button above) or which entries you want to not appear (second radio button above) in CSA MC Event Logs.

The choices are: **System, Application, Security**

- **Event Source** In the text field, enter (one per line) event source parameters you want to filter by.

The event source is the software that logged the event, which can be either an application name, such as `SQL Server`, or a component of the system or of a large application, such as a driver name. For example, `Elnkii` indicates the EtherLink II driver.

- **Event Severity (Type)** Select one more checkboxes to filter the viewing of events according to severity. If you select no checkboxes, all severity levels are included in the rule.

The choices are: **Information, Warning, Error, Audit Success, Audit Failure**

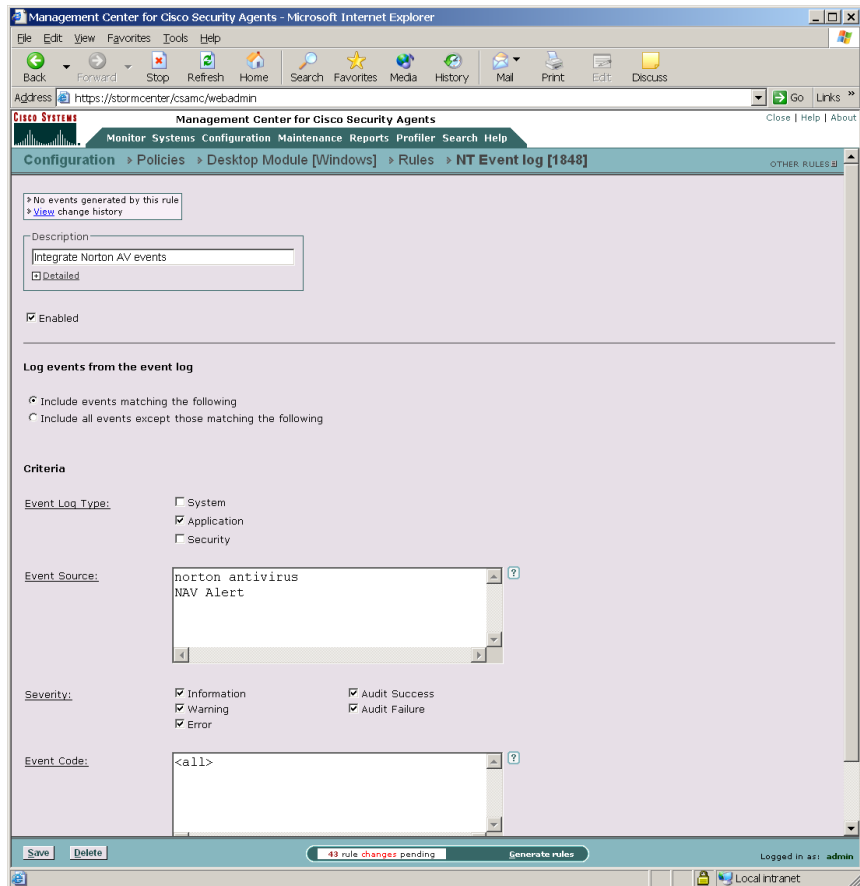
- **Event Code (Event ID)** In the text field, enter (one per line) event code parameters you want to filter by.

The event code is the number identifying the particular event type. For example, 6005 is the ID of the event that occurs when the Event log service is started. You can find the event IDs for Windows security events by searching for the following articles on the Microsoft web site: Q174074, Q299475, and Q301677.

Step 6 Click the **Save** button.**Note**

To receive messages logged by Norton AntiVirus, select the **Application** checkbox and enter `Norton AntiVirus` in the Event Source edit box. See the [“Global Events” section on page 5-18](#) for more information.

Figure 4-18 NT Event Log Rule



Registry Access Control

Use registry access control rules to allow or deny applications from writing to specified registry keys.



Note These instructions are a continuation of [Configuring Policies, page 4-13](#).

Step 1 To add rules to your policy, click the **Add rule** link at the bottom of the rule list. A pop-up list of the available rule types appears.

Step 2 Select the **Registry access control** rule. This takes you to the configuration view for this rule type.



Note This rule type is not available for UNIX policies.

Step 3 Enter the following information:

- **Description** Enter a description of this rule.
This description appears in the list view for the policy. Optionally, expand the **+Detailed** field to enter a longer description.
- **Enabled** Use this checkbox to enable this rule within the policy. (It is enabled by default.)
By not selecting this checkbox, you can save this rule, but it will not be active in the policy and it will not be distributed to groups.

Step 4 **Take the following action**—Select an action type from the pulldown list. For further details on using allow versus deny rules, see [Writing Rules: Allow vs. Deny](#), page 4-6.

- **Add process to application class** Use this for defining dynamic application classes. See [Chapter 6, “Using Application Classes”](#) for details.
- **High Priority Deny** Select this action type to create a deny rule that takes precedence over all other allow, deny, and query rules. For example, if you configure an allow rule that conflicts with this high priority deny, the high priority deny always takes precedence. Under the same circumstances, if the deny is not high priority and a conflicting allow rule exists in the same policy, the allow takes precedence.
- **Allow** Select this action type to create a registry access control rule that allows the applications you specify to perform the selected operation on the registry keys you indicate. Because the default action of all policies is "allow", generally, you'll only want to configure allow rules as exceptions to existing deny rules within policies.
- **Query User (Default Allow)** Select this action type to prompt the user when the action you indicate occurs. The user then has 5 minutes to answer Yes, No, Yes to all, or No to all. By default, it will be allowed unless the user decides otherwise. See [page 4-10](#) for details.
- **Query User (Default Deny)** Select this action type to prompt the user when the action you indicate occurs. The user then has 5 minutes to answer Yes, No, Yes to all, or No to all. By default, it will be denied unless the user decides otherwise. See [page 4-10](#) for details.

Text used to query user If you are configuring a Query User rule, the text you type into this field is the same text that will appear in the Query User pop-up box to explain what is occurring on the system to the user.
- **Deny** Select this action type to create a registry access control rule that stops the application you specify from performing the selected operation on the registry keys you indicate.

Step 5 and

- **Log** Enable this checkbox to turn logging on for this rule. Generally, you will want to turn logging on for all deny rules. This means that the denied system action in question is logged and sent to the server at regular time intervals.
- **Take precedence over similar <action type> rules** Enable this checkbox to manipulate policy rule precedence so that this rule is evaluated before other similar rules. You should generally NOT require this checkbox. Do not use it without understanding how it works. See [Writing Rules: Manipulating Precedence, page 4-8](#) for details.

Step 6 **When—Applications in any of the selected classes**

Select one or more preconfigured application classes here to indicate the application(s) whose access to the selected registry keys you want to exercise control over.

Note that the entry, <All Applications>, is selected by default. You can use this default or you can unselect it and create your own application classes. For application class configuration details, see [Chapter 6, “Using Application Classes.”](#)

Step 7 **Attempt to write to any of these registry entries**

Click the **Insert Registry Set** link to select one or more pre-configured registry sets for this rule. See the [“Included Registry Sets” section on page 7-20](#) for details on included operating system registry values.



Note You cannot enter registry literals here. You must create a registry set variable if you are not using pre-configured registry sets.

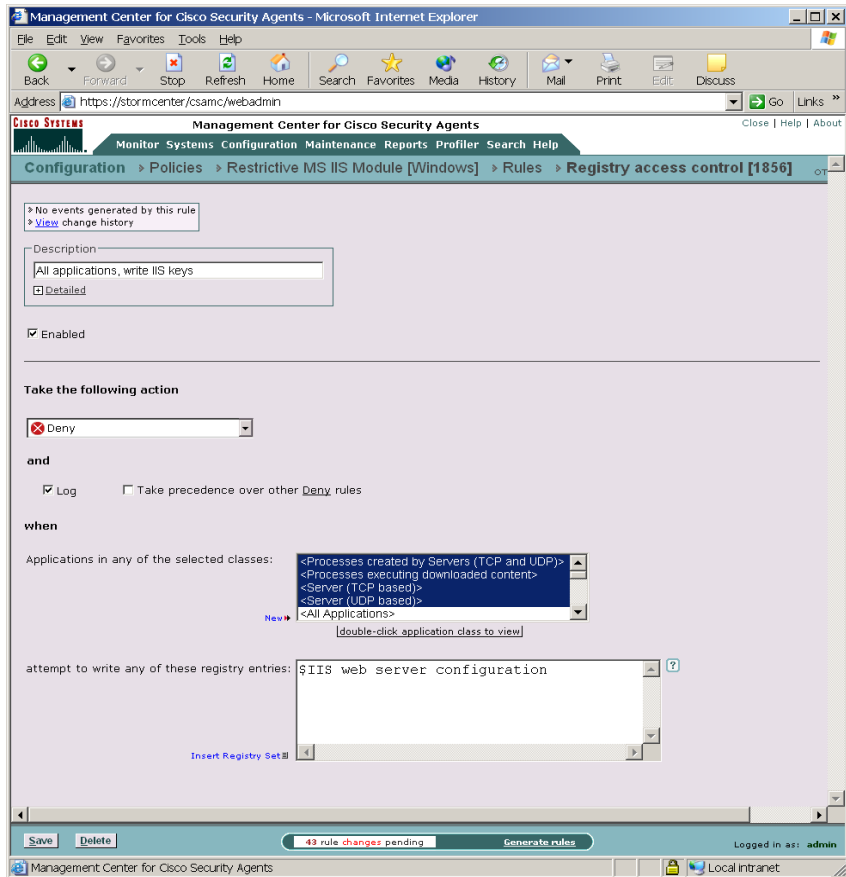
Step 8 When you are finished configuring your Registry access control rule, click the **Save** button.

This rule is now part of your new policy. It takes effect when the policy is attached to a group and then downloaded by an agent on the network.

**Caution**

In order to distribute rules to the correct hosts, you must associate policies with groups and then Generate rules. See [page 4-92](#) and [page 4-97](#) for instructions.

Figure 4-19 Registry Access Control Rule



Service Restart

Use the Service restart rule to have the agent restart Windows NT services that have gone down on a system or are simply not responding to service requests.



Note These instructions are a continuation of [Configuring Policies, page 4-13](#).

Step 1 To add rules to your policy, click the **Add rule** link at the bottom of the rule list. A pop-up list of the available rule types appears.

Step 2 Select the **Service restart** rule. This takes you to the configuration view for this rule type (see [Figure 4-20](#)).



Note This rule type is not available for UNIX policies.

Step 3 Enter the following information:

- **Description** Enter a description of this rule.
This description appears in the list view for the policy. Optionally, expand the **+Detailed** field to enter a longer description.
- **Enabled** Use this checkbox to enable this rule within the policy. (It is enabled by default.)
By not selecting this checkbox, you can save this rule, but it will not be active in the policy and it will not be distributed to groups.
- **Log** Enable this checkbox to turn logging on for this rule.

Step 4 **Restart the following service**

Enter a service here you want the agent to automatically restart should it go down for any reason. When entering services here, use the syntax found in the following locations:

- On Windows XP and Windows 2000: Start>Settings>Control Panel>Administrative Tools>Services "Name" field
- On Windows NT: Start>Settings>Control Panel>Services "Service" field

Step 5 When

Select one or both of the following checkboxes.

- Not responding to Service Control Manager The Windows Service Control Manager checks the status of system services and recognizes when a service is not responding. Selecting this checkbox causes the Cisco Security Agent to restart the specified service when it does not respond to the Windows Service Control Manager.
- Not responding to network requests for service: Select this checkbox and then choose a network service (such as HTTP) from the available pulldown list. The Cisco Security Agent will monitor whether the system is responding to network requests for the protocols in the network service. If not, it will restart the Windows NT service specified in this rule.

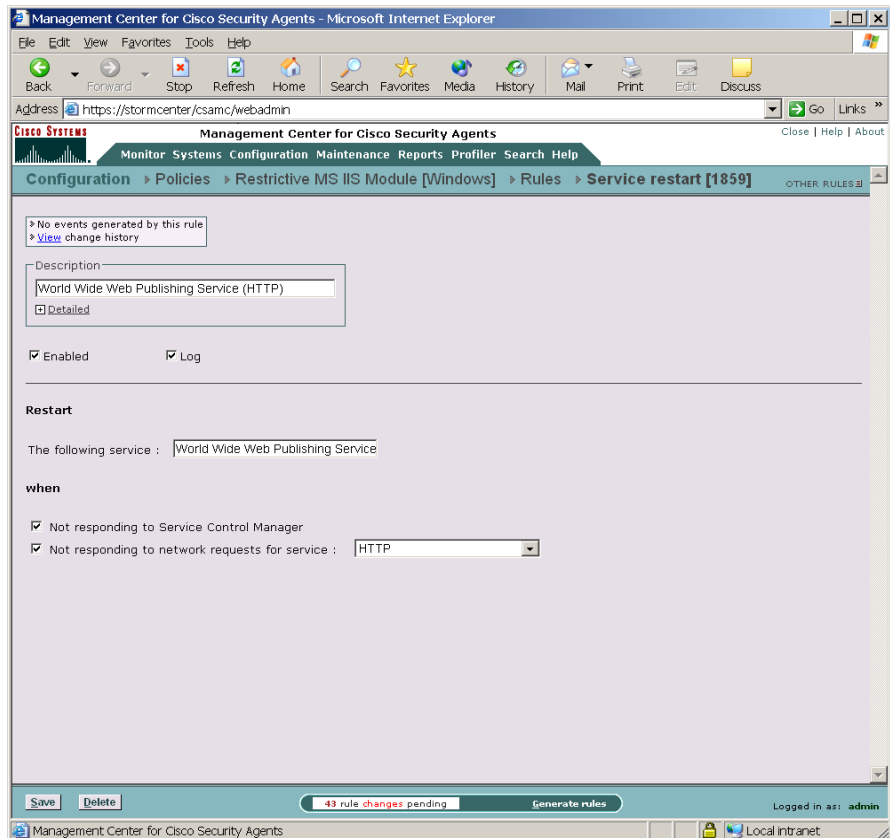
**Caution**

An agent must have the network shim installed in order for the "Not responding to network requests for service" feature to work.

Step 6 Click **Save** when finished.**Note**

The Service Restart rule is different from the Windows NT configurable restart service. Windows NT only restarts processes that have gone away. The agent restarts a process that experiences a failure of any kind.

Figure 4-20 Service Restart Rule



Sniffer and Protocol Detection

Use the Sniffer and protocol detection rule to cause an event to be logged when non-IP protocols and packet sniffer programs are detected running on systems.

Non-IP protocols, such as IPX, AppleTalk, and NetBEUI, are used to provide distributed computing workgroup functions between server and clients and/or sharing between peer clients.

A packet sniffer (also controlled by this rule type) is a program that monitors and analyzes network traffic. Using this information, a network manager can troubleshoot network problems. A sniffer can also be used illegitimately to capture data being transmitted on a network. Sensitive information such as login names and passwords can be extracted from this data and used to break into systems.

The Sniffer and protocol detection rule is a monitoring tool. By adding this rule to a policy, you are causing an event to be logged when any non-IP protocols and packet sniffer programs are detected running on systems which receive this rule.

**Note**

You can use the Sniffer and protocol detection rule page to configure exceptions to this monitoring rule. If you select any non-IP protocols or enter any packet sniffer programs here, you are allowing them to run on systems without generating events. Only non-IP protocols and packet sniffer programs which you explicitly exclude as part of the rule will not cause events to be logged. Otherwise, all are monitored when you add this rule to a policy.

These instructions are a continuation of [Configuring Policies, page 4-13](#).

-
- Step 1** To add rules to your policy, click the **Add rule** link at the bottom of the rule list. A pop-up list of the available rule types appears.
- Step 2** Select the **Sniffer and protocol detection** rule. This takes you to the configuration view for this rule type (see [Figure 4-19](#)).

**Note**

This rule type is not available for UNIX policies.

- Step 3** Enter the following information:
- **Description** Enter a description of this rule. This description appears in the list view for the policy. Optionally, expand the **+Detailed** field to enter a longer description.
 - **Enabled** Use this checkbox to enable this rule within the policy. (It is enabled by default.)
By not selecting this checkbox, you can save this rule, but it will not be active in the policy and it will not be distributed to groups.

Step 4 Select one or more preconfigured **Standard protocols** here to be excluded as part of this rule. The protocols you select here are the only non-IP protocols that will not generate events when they are detected.

If the non-IP protocol(s) you want to exclude are not included in the Standard Protocols list, enter your own in the **Non-standard protocols and packet sniffers** text field. By default, TCP/IP Protocol is already excluded.

This is also where you should enter any packet sniffer programs you want to exclude from this rule. (Find the names for these programs in Cisco Security Agent log files or in system registries.) For example, enter:

`PacketDriver`

In this example, Windump is the application. The libcap packet capture driver registers using the name PacketDriver.

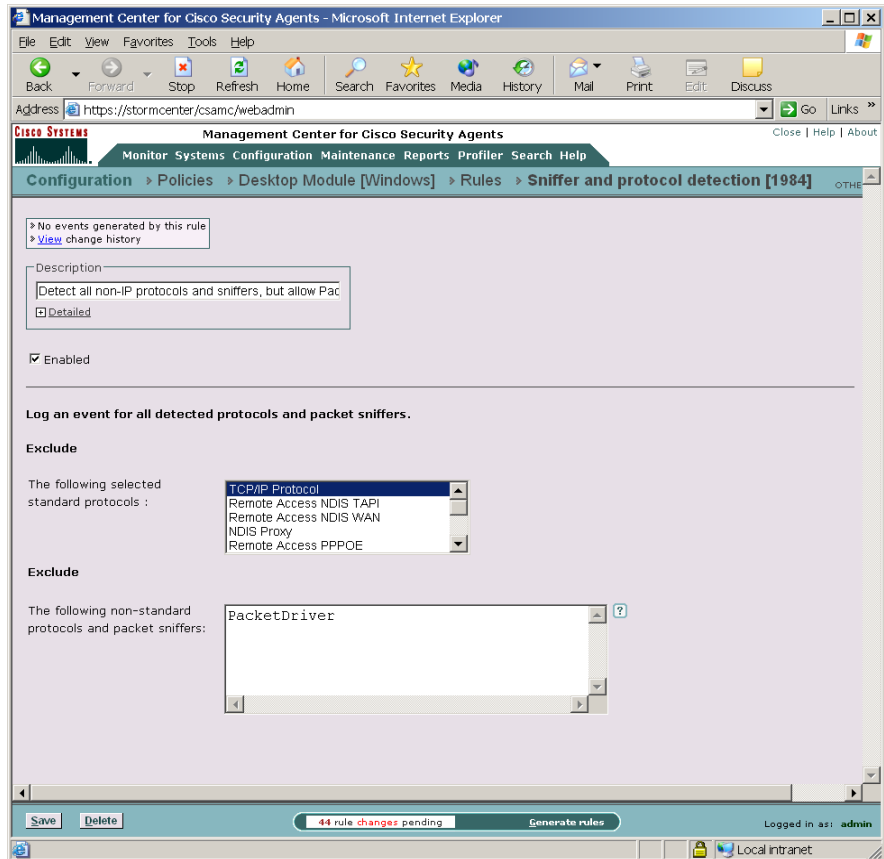
Step 5 Click the **Save** button.



Note

If you have multiple sniffer and protocol detection rules, the exceptions are combined.

Figure 4-21 Sniffer and Protocol Detection Rule



UNIX Only Rules

The following rules are only available for UNIX policies.

Network Interface Control

Use the Network interface control rule to specify whether applications can open a device and act as a sniffer (promiscuous mode). A packet sniffer is a program that monitors and analyzes network traffic. Using this information, a network manager can troubleshoot network problems. A sniffer can also be used illegitimately to capture data being transmitted on a network. Sensitive information such as login names and passwords can be extracted from this data and used to break into systems.

These instructions are a continuation of [Configuring Policies, page 4-13](#).

-
- Step 1** To add rules to your policy, click the **Add rule** link at the bottom of the rule list. A pop-up list of the available rule types appears.
- Step 2** Select the **Network interface control** rule. This takes you to the configuration view for this rule type (see [Figure 4-22](#)).
- Step 3** Enter the following information:
- **Description** Enter a description of this rule.
This description appears in the list view for the policy. Optionally, expand the **+Detailed** field to enter a longer description.
 - **Enabled** Use this checkbox to enable this rule within the policy. (It is enabled by default.)
By not selecting this checkbox, you can save this rule, but it will not be active in the policy and it will not be distributed to groups.

Step 4 **Take the following action**—Select an action type from the pulldown list. For further details on using allow versus deny rules, see [Writing Rules: Allow vs. Deny, page 4-6](#).

- **Add process to application class** Use this for defining dynamic application classes. See [Chapter 6, “Using Application Classes”](#) for details.
- **High Priority Deny** Select this action type to create a deny rule that takes precedence over all other allow, deny, and query rules. For example, if you configure an allow rule that conflicts with this high priority deny, the high priority deny always takes precedence. Under the same circumstances, if the deny is not high priority and a conflicting allow rule exists in the same policy, the allow takes precedence.
- **Allow** Select this action type to create a registry access control rule that allows the applications you specify to perform the selected operation on the registry keys you indicate. Because the default action of all policies is "allow", generally, you'll only want to configure allow rules as exceptions to existing deny rules within policies.
- **Deny** Select this action type to create a registry access control rule that stops the application you specify from performing the selected operation on the registry keys you indicate.

Step 5 and

- **Log** Enable this checkbox to turn logging on for this rule. Generally, you will want to turn logging on for all deny rules. This means that the denied system action in question is logged and sent to the server at regular time intervals.
- **Take precedence over similar <action type> rules** Enable this checkbox to manipulate policy rule precedence so that this rule is evaluated before other similar rules. You should generally NOT require this checkbox. Do not use it without understanding how it works. See [Writing Rules: Manipulating Precedence, page 4-8](#) for details.

Step 6 When—Applications in any of the selected classes

Select one or more preconfigured application classes here to indicate the application(s) whose access to the selected resource(s) want to exercise control over.

Note that the entry, <All Applications>, is selected by default. You can use this default or you can unselect it and create your own application classes. For application class configuration details, see [Chapter 6, “Using Application Classes.”](#)

Step 7 Attempt the following operations

Select one or more of the following checkboxes:

- Open a stream connection to the NIC driver
- Put the NIC into promiscuous mode



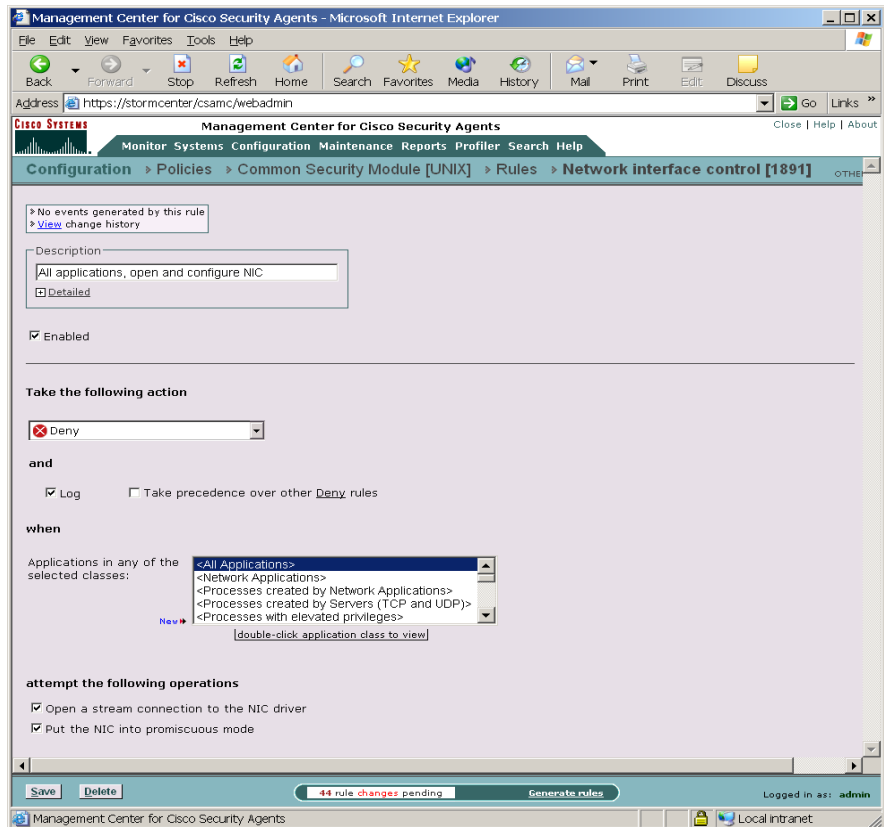
Note If you have selected the Allow radio button, when you select to "Put the NIC into promiscuous mode", the "Open a stream connection to the NIC driver" checkbox is also automatically selected. It must be enabled for promiscuous mode to work.

Conversely, if you have selected a Deny radio button, when you select the "Open a stream connection to the NIC driver" checkbox, the "Put the NIC into promiscuous mode" checkbox is also automatically selected. If you deny one, the other is automatically denied as well.

Step 8 When you are finished configuring your rule, click the **Save** button.

Note If you are using remote management tools and you are configuring a Network interface control rule to deny "all applications" from opening a stream connection to the NIC and operating in promiscuous mode, you may want to make an exception for the remote management application (if you want to run snoop).

Figure 4-22 Network Interface Control Rule



Resource Access Control

Use the Resource access control rule to protect systems from symbolic link attacks. In this type of attack, an attacker attempts to determine the name of a temporary file prior to its creation by a known application. If the name is determined correctly, the attacker could then create a symbolic link to the target file for which the user of the application has write permissions. The application process would then overwrite the contents of the target file with its own output when it tries to write the named temporary file.

For example, a directory such as /tmp is writable by everyone. An attacker could create a symbolic link in this directory to a protected file such as etc/shadow. This would then grant the attacker access to this sensitive information via a symbolic link from the /tmp directory.

By enabling the resource access control rule, you can prevent "suspicious" symbolic links from being followed. A suspicious symbolic link is one that meets the following criteria:

- The parent directory is a temporary directory such as /tmp and usr/tmp
- The symbolic link's owner is different from the parent directory's owner
- The symbolic link's owner is different from the effective UID of the process

These instructions are a continuation of [Configuring Policies](#), page 4-13.

Step 1 To add rules to your policy, click the **Add rule** link at the bottom of the rule list. A pop-up list of the available rule types appears.

Step 2 Select the **Resource access control** rule. This takes you to the configuration view for this rule type (see [Figure 4-24](#)).



Note This rule type is **ONLY** available for UNIX policies.

Step 3 Enter the following information:

- **Description** Enter a description of this rule. This description appears in the list view for the policy. Optionally, expand the **+Detailed** field to enter a longer description.

- **Enabled** Use this checkbox to enable this rule within the policy. (It is enabled by default.)
By not selecting this checkbox, you can save this rule, but it will not be active in the policy and it will not be distributed to groups.

Step 4 Select the **Symbolic Link Protection** checkbox to turn on that functionality.

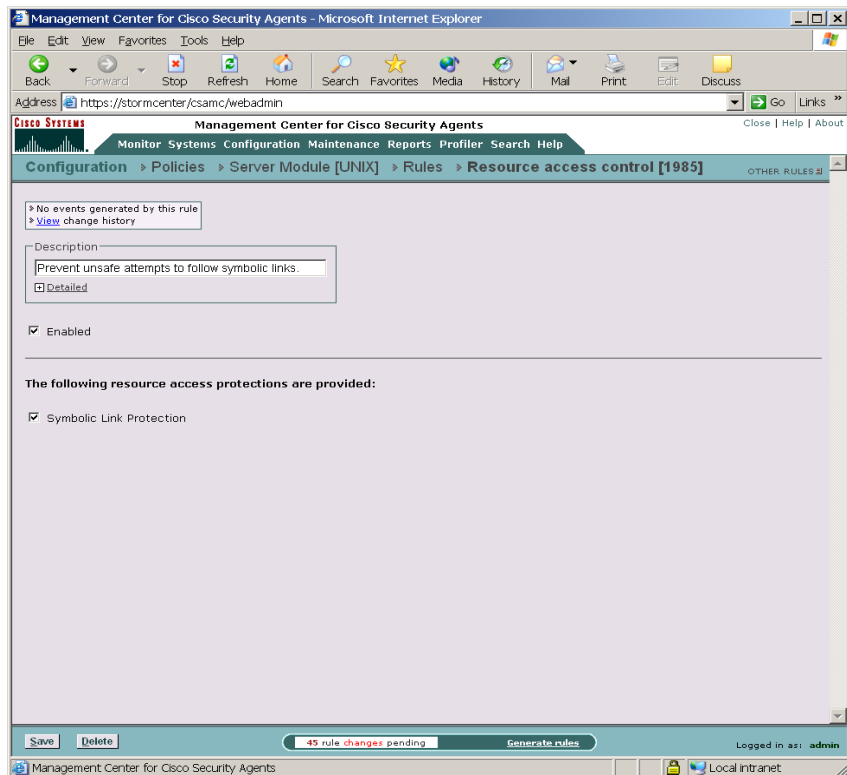
Step 5 Click the **Save** button.



Caution

Symbolic Links: If you create a File access control rule to protect a symbolic link, **ONLY** that symbolic link is protected. The underlying resource, unless also specified, is **NOT** protected. For example, a File access control rule written for `/etc/hosts` does not protect `/etc/inet/hosts`. Similarly, a File access control rule written for `/etc/inet/hosts` does not protect `/etc/hosts`. If you want to protect a symbolic link and its underlying resource, both must be specified in the rule.

Figure 4-23 Resource Access Control Rule



Rootkit / kernel Protection

Use the Rootkit / kernel protection rule to prevent unauthorized access to the operating system. In effect, this rule prevents drivers from dynamically loading after boot time. You can specify exceptions to this rule for authorized drivers that you are allowing to load any time after the system is finished booting.

**Note**

These instructions are a continuation of [Configuring Policies, page 4-13](#).

- Step 1** To add rules to your policy, click the **Add rule** link at the bottom of the rule list. A pop-up list of the available rule types appears.
- Step 2** Select the **Rootkit / kernel protection** rule. This takes you to the configuration view for this rule type (see [Figure 4-24](#)).

**Note**

This rule type is **ONLY** available for UNIX policies.

- Step 3** Enter the following information:
- **Description** Enter a description of this rule.
This description appears in the list view for the policy. Optionally, expand the **+Detailed** field to enter a longer description.
 - **Enabled** Use this checkbox to enable this rule within the policy. (It is enabled by default.)
By not selecting this checkbox, you can save this rule, but it will not be active in the policy and it will not be distributed to groups.

- Step 4** **Allow only the following modules to be loaded after boot**

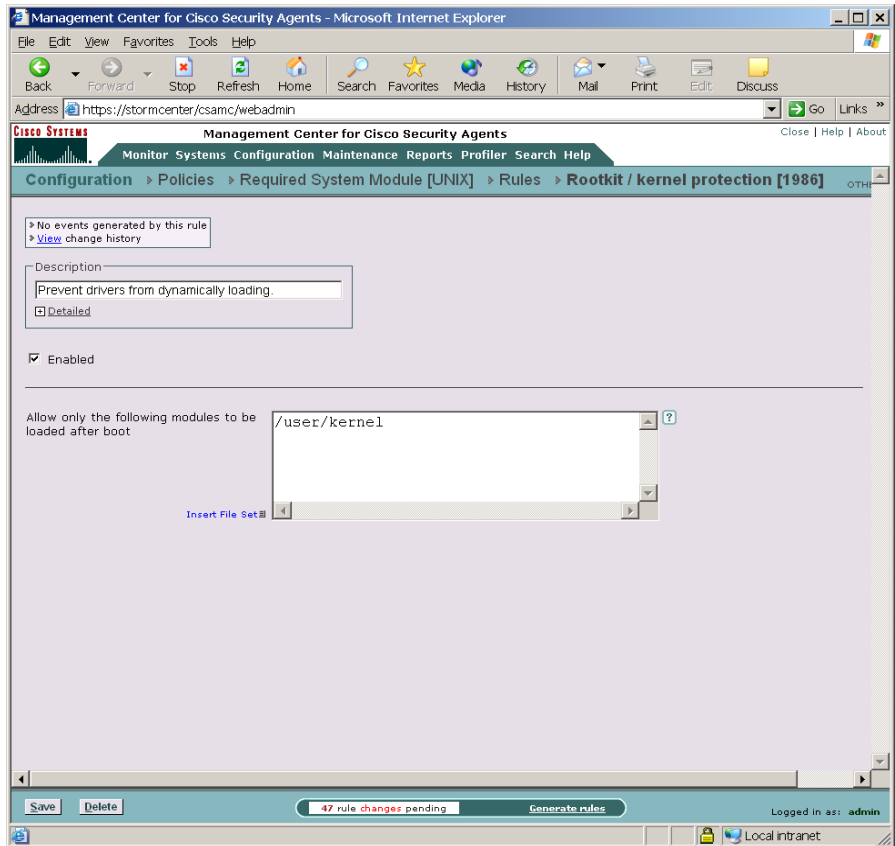
By default, this field contains <none> which in turn prevents all drivers from loading after boot time. You can make exceptions to this rule by entering the names of drivers you want to exempt from the rule and therefore allow to load at any time.

**Caution**

If you enter files sets which use "content-matching" constraints, via the Insert File Set link, the content-matching constraints are ignored.

- Step 5** Click the **Save** button.

Figure 4-24 Rootkit / kernel Protection Rule



Syslog Control

Use the Syslog control rule to have specified Solaris Syslog items appear in the CSA MC Event Log for selected groups.



Note These instructions are a continuation of [Configuring Policies, page 4-13](#).

Step 1 To add rules to your policy, click the **Add rule** link at the bottom of the rule list. A pop-up list of the available rule types appears.

Step 2 Select the **Syslog control** rule. This takes you to the configuration view for this rule type (see [Figure 4-25](#)).



Note This rule type is **ONLY** available for UNIX policies.

Step 3 Enter the following information:

- **Description** Enter a description of this rule.
This description appears in the list view for the policy. Optionally, expand the **+Detailed** field to enter a longer description.
- **Enabled** Use this checkbox to enable this rule within the policy. (It is enabled by default.)
By not selecting this checkbox, you can save this rule, but it will not be active in the policy and it will not be distributed to groups.

Step 4 **Log events from syslog**

- **Include events matching the following** Select this radio button to specify the criteria for Syslog entries which you want to appear in the CSA MC Event Log.
- **Include all events except those matching the following** Select this radio button to specify the criteria for Syslog entries which you do not want to appear in the CSA MC Event Log. (All criteria not specified here will appear in the CSA MC Event Log.)



Note You can configure CSA MC to correlate Syslog events logged across multiple systems. See the [“Global Events” section on page 5-18](#).

Step 5 **Criteria** (You should select at least one for the rule to have any effect.)

- **Event Source** In the text field, enter (one per line) event source parameters you want to filter by.
The event source is the software that logged the event, which can be an application name such as `/sbin/dhcpagent`, a kernel level driver module such as `scsi`, or the `unix` kernel itself.
- **Facility** Select one or more items from the list box you want to appear (first radio button above) or which entries you want to not appear (second radio button above) in CSA MC Event Logs.
- **Priority** Select one more checkboxes by which to filter the viewing of events according to priority. If you select no checkboxes, all priorities are included in the rule.
- **Message Pattern** In the text field, enter (one per line) message patterns you want to match and filter by. To match, the string you enter must literally appear somewhere within the message.

Step 6 Click the **Save** button.

Syslog rule configuration examples

For Example:

Configure a syslog rule to log warning messages such as the one listed below:

```
Apr 29 13:46:35 myhost /sbin/dhcpagent[39]: [ID 929444
daemon.warning] configure_if: no IP broadcast specified for
eri0
```

To get every message of category "warning" from the `/sbin/dhcpagent` daemon, you would configure your syslog rule in the following manner (See [Figure 4-25](#)):

Select the "Include events matching the following" radio button and enter:

- Facility: daemon
- Event Source: `/sbin/dhcpagent`
- Priority: Warning checkbox

- Message Pattern: <all>

For Example:

Configure a syslog rule to log failed su root attempts such as the one listed below:

```
Apr 29 13:49:23 myhost su: [ID 810491 auth.crit] 'su root'  
failed for haxor on /dev/pts/4
```

To get messages for failed su root attempts, you would configure your syslog rule in the following manner:

Select the "Include events matching the following" radio button and enter:

- Facility: auth
- Event Source: su
- Priority: Alert and Above checkbox
- Message Pattern: root

For Example:

Configure a syslog rule to include all events but exclude all lockstat-related messages such as the one listed below:

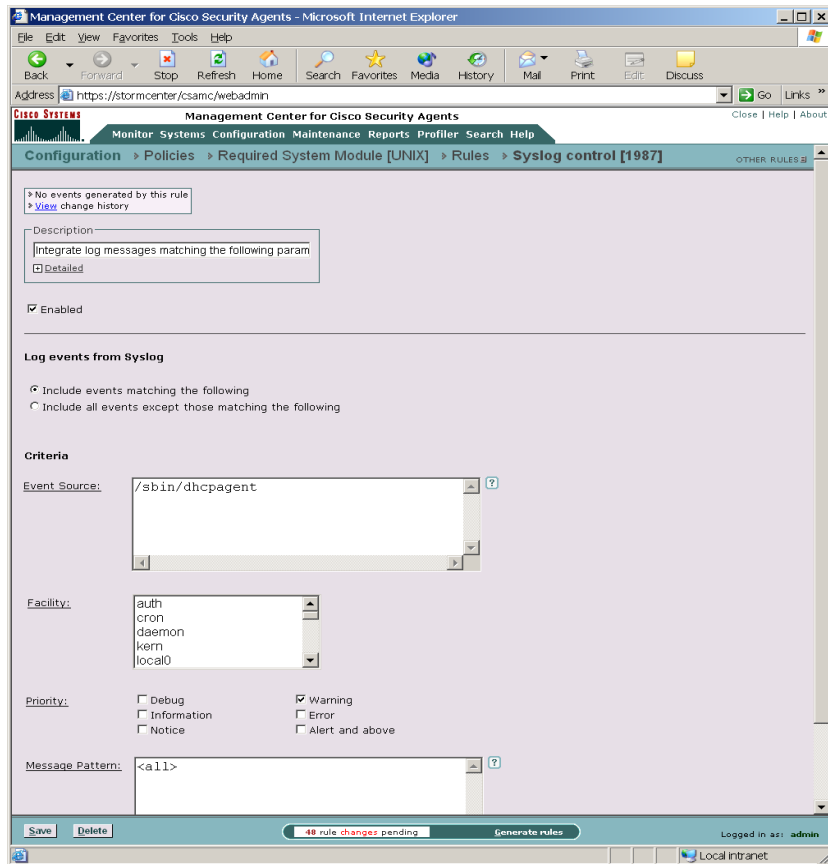
```
Apr 29 13:46:43 myhost genunix: [ID 936769 kern.info] lockstat0  
is /pseudo/lockstat@0
```

To log all events except for lockstat-related messages, configure your rule in the following manner:

Select the "Include events except those matching the following" radio button and enter:

- Facility: kern
- Event Source: <all>
- Priority: all checkboxes
- Message Pattern: lockstat

Figure 4-25 Syslog Control Rule



Attaching Policies to Groups

When you configure a policy, you are combining access control rules and/or system correlation rules under a common name. That policy name is then attached to a group of hosts and it uses the rules that comprise the policy to control the actions that are allowed and denied on those hosts. See [Configuring Policies, page 4-13](#).

CSA MC gives you the option of attaching a policy to a group using the **Modify policy associations** link in the Group configuration page or attaching a group to a policy using the **Modify group associations** link in the Policy list view page. (You can use the Modify policy associations link to attach multiple policies to a group and use the Modify group association link to attach one policy to multiple groups.)

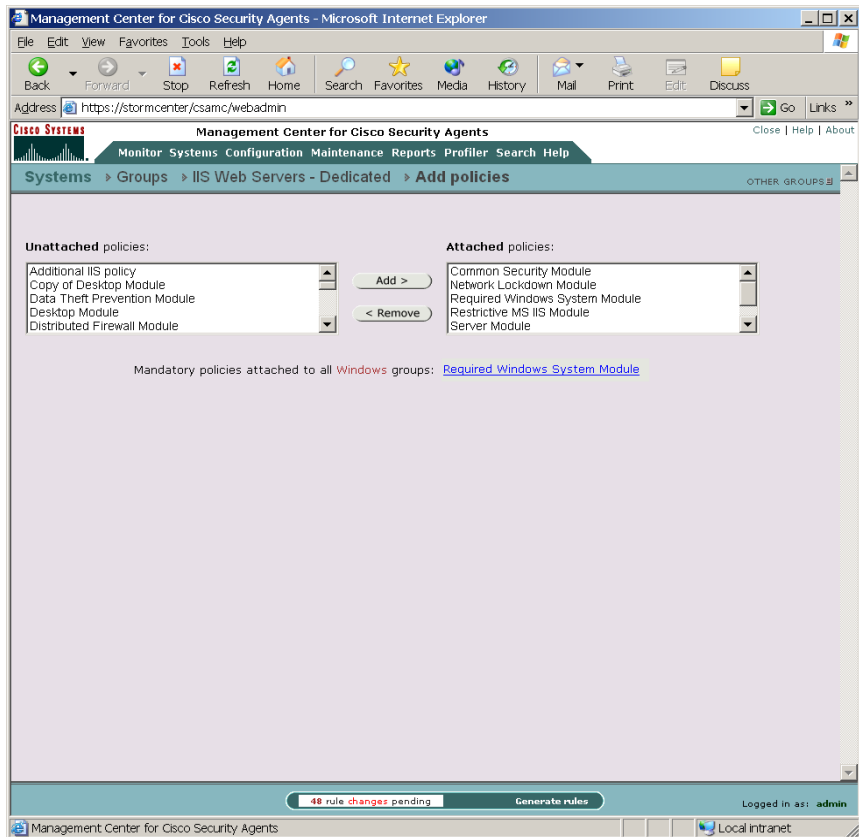
To attach a policy or policies to an existing group using the **Modify policy associations** link in the Group configuration page, do the following.

-
- Step 1** Attach a policy to a particular group by accessing that group's edit view. From **Systems** in the menu bar, click on **Groups** to access the group's list view.
 - Step 2** From the group list view, click the link for the group you want to attach a policy to. This brings you to that group's edit view.
 - Step 3** From the edit view, click the **Modify policy associations** link. This takes you to a page containing swap boxes. The left box contains the policies not attached to this group. The right box contains policies that are attached to this group.
 - Step 4** To add an existing policy to this group, select the policy in the left box and click the **Add** button. The selected policy moves to the right box and is now attached to the group.
 - Step 5** Group configuration pages display all rules, in order of precedence, which are applied to the group in question. See [Figure 4-27](#).



Note To remove a policy from a group, select the policy in the right box and click the **Remove** button. It moves back to the left box. (The policy is not deleted from the database, it is just no longer applied to the group.) Although the selected policy is no longer attached to the group, this is not apparent in the GUI until you click the **Generate rules** link in the bottom frame and then the **Generate** button.

Figure 4-26 Attaching Policies

**Note**

You can try out policies on host systems by selecting Test Mode for a group. Selecting Test Mode and enabling logging on rules attached to "test mode" groups causes the agent to log designated denied events triggered by policies but not take any actions on those events.

Using Test Mode

Test Mode is useful when you are installing a new host or are modifying a host configuration and want to understand the ramifications without actually impacting host operation. In Test Mode, the agent will not deny any action or operation even if an associated policy says it should be denied. Instead, the agent will allow the action but log an event if a deny or query rule is triggered (if logging is enabled for the rule) and log an event when an allow rule with logging enabled is triggered. This helps you to understand the impact of deploying a policy on a host before enforcing it. If examining the logs shows you that the policy is working as intended on a group, you can then remove the Test Mode designation.

When using Test Mode, you'll likely also want to enable Verbose Logging mode. This way, the agent will not suppress any log messages as it normally does when several of the same log message are received.

When an agent running in Test Mode sends events to CSA MC, event log messages are preceded with the words "Test mode". There are some exceptions to this. For example, event log messages related to detected events such as port scans and malformed packets are not preceded by the words "Test mode." Event detection (not prevention) messages appear the same in the event log regardless if Test mode is on or off.



Note

If a host belongs to a group with Test Mode selected, all policies associated with that host are in test mode (even if the host is part of another group that does not have Test Mode selected), not just the policies applied to the test group. Therefore, test mode applies to the host as a whole, not to policies.



Caution

You should be aware that putting a deployed "live" policy into Test Mode turns off all security that the policy in question had been providing. Keep this in mind when using Test Mode to analyze how policies are working.

When all policies on a system are running in Test Mode, the system in question is acting as an Intrusion Detection System (IDS). When Test Mode is disabled and the policies go live, the system becomes an Intrusion Prevention System.

Figure 4-27 Test Mode Enabled

The screenshot shows the Management Center for Cisco Security Agents web interface in Microsoft Internet Explorer. The browser address bar shows the URL `https://stormcenter/cssamc/webadmin`. The page title is "Management Center for Cisco Security Agents". The navigation menu includes "Monitor Systems", "Configuration", "Maintenance", "Reports", "Profiler", and "Search Help". The breadcrumb trail is "Systems > Groups > Test Mode Systems".

The main content area displays the configuration for "Test Mode Systems". A "Quick links" box contains the following links:

- [Modify host membership](#)
- [Modify policy associations](#)
- [View related events](#)
- [Explain rules](#)

The configuration fields are as follows:

- Name: Test Mode Systems
- Description: Systems operating in test mode
- Detailed
- Target operating system: Windows
- Polling interval (sec): 600
- Profiler: Disabled
- Test mode
- Verbose logging mode
- No user interaction [?](#)

The "Attached Policies" section contains the following table:

Name	Description
Email Quarantine	Prevent Email applications from accessing known viruses
Required Windows System Module	Policy module to allow critical Windows functions

The "Combined Policy Rules" section is currently empty. At the bottom of the page, there are "Save" and "Delete" buttons, a status bar indicating "48 rule changes pending", and a "Generate rules" button. The user is logged in as "admin".

Generating Rule Programs

**Caution**

When you make changes to existing CSA MC configurations, they are saved in the database, but they are not yet distributed to the agents across your network. You *must* click the **Generate rules** link in the bottom frame of CSA MC to first view all new and edited configurations and then distribute them to the agents. (When you have pending changes, the line beneath Generate rules link flashes.)

The Generate rule programs view displays the status of all non-distributed database items with the name of the administrator who made the configuration changes. A **Details** link appears beside each edited configuration item. Click this link to view what modifications were made to the configuration in question.

Once you've checked these modifications, you can either go back and change or delete configurations or you can click the **Generate** button (in the bottom frame) to distribute all updates.

**Note**

Before you generate rule programs and distribute them to agents, you can view all database changes, including the time the changes were made and the administrator who made them by accessing the **Audit Trail** view from the Maintenance drop-down list. See the [“Using Audit Trail” section on page 2-6](#) for information.

**Caution**

If you have set the Group polling interval too low for too many hosts, CSA MC warns you of this fact when you attempt to generate rules. If the average polling frequency (number of agents polling per second) is greater than 100, rule generation is not allowed. If that average is between 15 and 100 you are advised to increase the polling interval.

Figure 4-28 Generate Configuration

Management Center for Cisco Security Agents - Microsoft Internet Explorer

Address: https://stormcenter/csamc/webadmin

Management Center for Cisco Security Agents

Monitor Systems Configuration Maintenance Reports Profiler Search Help

Generate Rule Programs

Warning :
The following policies are not attached to any hosts or groups:

- [File System Lockdown Module](#) [U]
- [Insecure Management Module](#) [U]
- [Restrictive SendMail Module](#) [U]
- [CiscoWorks Restrictive VMS Module](#) [W]
- [Data Theft Prevention Module](#) [W]
- [Distributed Firewall Module](#) [W]
- [File Integrity Module](#) [W]
- [Windows Security Events Module](#) [W]
- [Windows XP Help Center Module](#) [W]

8 changes since the last rule program generation:

Action	Time
Modify policy 'Email Quarantine' [Details]	5/30/2003 3:14:39 PM
Add File access control rule to policy 'Email Quarantine'	5/30/2003 3:14:50 PM
Modify File access control rule in policy 'Email Quarantine' [Details]	5/30/2003 3:15:12 PM
Create FS variable 'Untitled_1'	5/30/2003 3:15:21 PM
Modified file set variable 'Known virus files' [Details]	5/30/2003 3:18:23 PM
Modify File access control rule in policy 'Email Quarantine' [Details]	5/30/2003 3:19:07 PM
Modify File access control rule in policy 'Email Quarantine' [Details]	5/30/2003 3:38:49 PM
Add policy 'Email Quarantine' to group 'Default Desktops'	5/30/2003 4:47:50 PM

Press the Generate button to create and distribute rule programs based on the current configuration:

3 rule changes pending

Logged in as: admin

Management Center for Cisco Security Agents Local intranet



Using System Correlation Rules

Overview

Management Center for Cisco Security Agents provides preconfigured rules you can add to your policies that allow CSA MC to correlate events across multiple systems. When these rules are triggered by one or more errant system actions across a network, the server registers this occurrence and automatically sends out new rules to all Cisco Security Agents to prevent this action from executing on any additional systems. The Cisco Security Agent also uses heuristics to detect and terminate suspicious activities on systems, such as network worm infections.

Add the preconfigured rules described in this chapter to your policies just as you would other rules. For details on adding rules to policies and attaching policies to groups, refer back to [Chapter 4, “Building Policies.”](#)

This section contains the following topics.

- [Event Correlation and Heuristics, page 5-2](#)
- [Global Events, page 5-18](#)

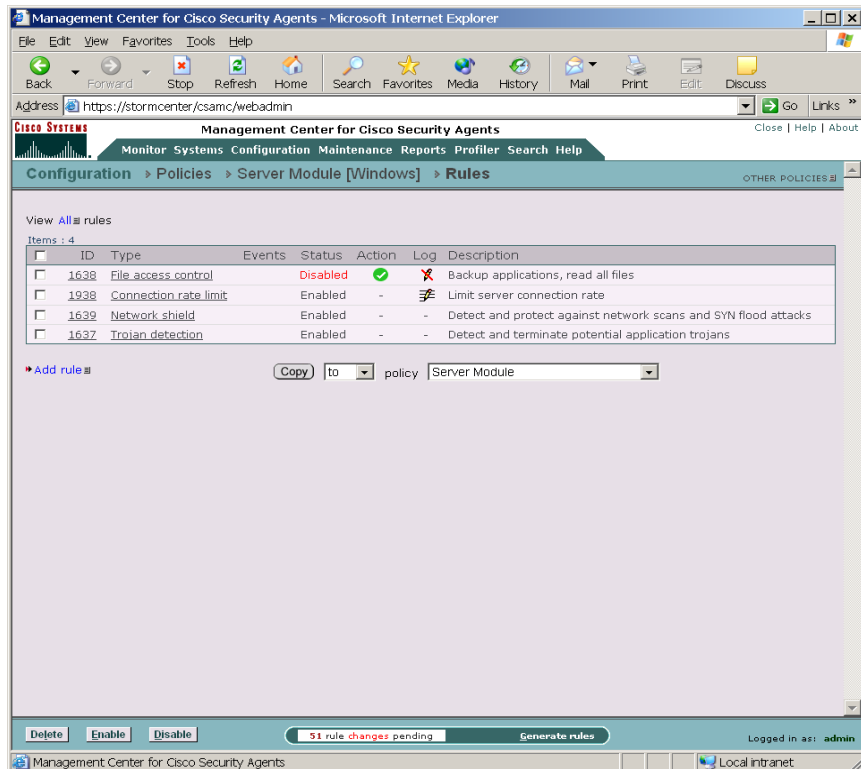
Event Correlation and Heuristics

The Network shield rule which controls SYN flood protection and port scan detection, the Network worm protection rule, and the Trojan detection rule are some examples of preconfigured rules you can add to your policies in the same way you add other rules. These are basic system hardening, event correlation, and heuristic features that should be applied in most cases. Some are used in the Windows Server Module policy shown in [Figure 5-1](#).

**Note**

While you can add several access control rules to one policy, you can only add once instance of the Network shield rule and the Network worm protection rule to a single policy.

Figure 5-1 Server Module



Network Worm Protection

Network and email worms are some of the most commonly spread and costly attacks affecting corporate networks today. Some well-known worms include ILOVEYOU, Anna Kournikova, and variations thereof. These worms easily infected systems, passing undetected through most security software until virus scanner vendors provided updates to detect these virus signatures. Even with this detection capability, if the worm is modified in any way, it is again undetectable by virus scanners.

**Note**

This rule type is not available for UNIX policies.

When a worm of this type is received through email and executed by unsuspecting users, it generally attempts to send copies of itself to all entries in the email address book of the user. In doing this, the worm modifies registry keys, writes its own script files, and modifies existing files. This makes file recovery difficult and it can cause users to invoke the virus again when they attempt to open these infected files.

When the Network worm protection rule is triggered on a system, a Query User pop-up box appears. This box provides text explaining the suspicious system action that is occurring and what the application in question is attempting to do. When prompted, the user must choose one of the following buttons:

- Yes: Allows the application access to the resource in question.
- Terminate: This stops the system action in question and terminates the application that triggered the worm protection rule.

Network Worm Event Correlation

If you select the email worm correlation checkbox in the Global Event Correlation page (see [Global Events, page 5-18](#)), when a worm is detected, other agents will be notified to prevent the spread of this virus. Under these circumstances, the agent(s) report the file name the worm was written into. If at least two agents report worms writing to the same file name within an hour, the file is added to a dynamic list (@dynamic) of quarantined files. If you have a rule configured to stop dynamically quarantined files in a deployed policy, no further agents can open the contaminated file during the quarantine time frame. See also, [page 4-44](#) for information on using @dynamic in File access control rules.

Trojan Detection

A Trojan is a form of malicious programming code that is installed on a system by an unsuspecting user either thinking that he or she is running some other type of program, or as a result of some other activity such as reading an attachment to an email message. Once installed, Trojans may allow others to access and

virtually take over a system across the network. Other Trojans may be set up to automatically send mail messages or other types of network traffic (including system passwords) while the system owner is unaware of what is occurring.

**Note**

This rule type is not available for UNIX policies. Refer to the Buffer overflow rule information on [page 5-16](#) for similar UNIX functionality.

When the Trojan detection rule is triggered on a system, a Query User pop-up box appears. This box provides text explaining the suspicious system action that is occurring and what the application in question is attempting to do. When prompted, the user must choose one of the following buttons (The button options available for each Trojan rule vary. They may display any of the following):

- **Yes:** Allows the application access to the resource in question. (This option is not available if a password stealing Trojan is detected).
- **Yes to all:** Allows the application access to all related query user protected resources, with no further queries appearing.
- **No:** Denies the application access to the resource in question.
- **Terminate:** This stops the system action in question and terminates the application that triggered the Trojan rule, effectively killing the Trojan program.

**Caution**

In some cases, a Trojan might hide itself in another application, such as Internet Explorer. Then the application seen as being the Trojan program (for example, Word or Internet Explorer) is a legitimate application on your system. Some Trojans have the ability to present themselves in this manner. Pressing the Terminate button in such cases will kill the legitimate program. However, this may be necessary.

It could be useful, especially in the case of server systems, to use a service restart rule in conjunction with a trojan rule. This way, if you are forced to press the Terminate button when queried and you subsequently terminate the application in question, a service restart rule will cause the application to automatically restart.

Use the Trojan detection rule in a policy to detect and prevent Trojans from performing malicious acts on individual systems and networks. The included Trojan detection rule lets you enable several different types of Trojan detection.

- Trapping of keystrokes by network applications
(Detect applications that attempt to capture system keystrokes.)
If the system in question is unattended, the default response of "No" is automatically taken.
- Injecting code into other applications
(Detect applications that have been marked as downloaded content attempting to write code to space owned by other applications. e.g. injecting a malicious .dll into a privileged process)
If the system in question is unattended, the default response of "No" is automatically taken.
- Accessing memory owned by other applications
(Detect applications that attempt to interfere with the memory space of other applications or detect Trojans attempting to hide in another executable to escape detection and gain permissions to access other resources.)
If the system in question is unattended, the default response of "No" is automatically taken.
- Stealing local passwords
(Detect applications that attempt to steal local system passwords.)
If the system in question is unattended, the default response of "No" is automatically taken.
- Downloading and invoking executable files
(Detect applications that download executables and immediately attempt to execute them. This could be downloaded code as a result of a buffer-overflow attach or an executable downloaded by an application such as an email client or web browser.)
If the system in question is unattended, the default response of "Terminate" is automatically taken.

- Downloading and invoking ActiveX controls
(Detect applications that download ActiveX controls and immediately attempt to execute them.)
If the system in question is unattended, the default response of "No" is automatically taken.
This functionality limits applications from downloading ActiveX controls (signed and unsigned). This type of behavior is generally typical of a web browser and sites that require the downloading of ActiveX can trigger this rule. But the rule also covers a Trojan scenario in which a malicious application attempts to act like a web browser. Note that this rule may be unnecessary if system web browser settings are configured with a "High" security level that would restrict the downloading of ActiveX controls.
- Accessing system functions from code executing in data or stack space
(This behavior may be symptomatic of a buffer overflow attack and the agent prompts the user if this behavior is detected on the system.)
If the system in question is unattended, the default response of "Terminate" is automatically taken.
 - Patterns to be excluded: Use the Wizard from the Event log message in question to exclude a particular pattern when you are seeing buffer overflow events you believe are harmless.

**Note**

If an application is currently not enforcing any ActiveX download or Accessing system functions rules, that application must be restarted for any newly applied ActiveX download or Accessing system functions rules to take effect.

You also have the ability to select specific **application classes to exclude** from the various Trojan detection types you designate. For example, in some cases, debuggers may perform actions that can be misconstrued as Trojan behavior. Therefore, you would want to create an application class, and select it as an exclusion to one or more Trojan detection features.

**Note**

If you have multiple similar Trojan detection rules, the application class exceptions are combined.

**Note**

Additionally, if you distribute software updates over your network, you would want to exclude that application in the Downloading and invoking executable files—Trojan detection rule.

Figure 5-2 Trojan Detection Rule

The screenshot shows the Management Center for Cisco Security Agents web interface in Microsoft Internet Explorer. The browser address bar shows `https://stormcenter/cssamc/webadmin`. The page title is "Management Center for Cisco Security Agents". The navigation menu includes "Monitor", "Systems", "Configuration", "Maintenance", "Reports", "Profiler", and "Search". The breadcrumb trail is "Configuration > Policies > Server Module [Windows] > Rules > Trojan detection [1637]".

The main content area displays the configuration for the "Trojan detection [1637]" rule. It includes the following sections:

- Description:** "Detect and terminate potential application trojans". There is a checkbox for "Detailed" which is currently unchecked.
- Enabled:** A checked checkbox.
- User will be prompted when an application exhibits the following behaviors:**
 - Trapping of keystrokes by applications:** Select any application classes to be excluded: <none>, Apache Log Utilities, Apache Log Utilities_V3.2.0.27011.
 - Injecting code into other applications:** Select any application classes to be excluded: <none>, Apache Log Utilities, Apache Log Utilities_V3.2.0.27011.
 - Writing memory owned by other applications:** Select any application classes to be excluded: <none>, Apache Log Utilities, Apache Log Utilities_V3.2.0.27011.
 - Stealing local passwords:** Select any application classes to be excluded: <none>, Apache Log Utilities, Apache Log Utilities_V3.2.0.27011.
 - Downloading and invoking executable files:** Select any application classes to be excluded: <none>, Apache Log Utilities, Apache Log Utilities_V3.2.0.27011.
 - Downloading and invoking ActiveX controls:** Select any application classes to be excluded: <none>, Apache Log Utilities, Apache Log Utilities_V3.2.0.27011.
 - Accessing system functions from code executing in data or stack space:** Select any application classes to be excluded: <none>, Apache Log Utilities, Apache Log Utilities_V3.2.0.27011.

At the bottom of the configuration area, there are "Save" and "Delete" buttons, a status bar indicating "51 rule changes pending", and a "Generate rules" button. The user is logged in as "admin".

Replicate Feature

When you make rule changes and click the Save button for rule types that contain multiple checkboxes, such as Trojan detection (Network shield and Buffer overflow rules also provide this feature) a "replicate" link appears beside the "Saved changes" message at the top of the rule page. Click on **replicate** to access a pop-up box. From this box, you select other policies that contain Trojan detection rules and choose to propagate the same change(s) you made on the current page to Trojan detection pages in other policies. If the change you make to one Trojan rule page is a change you need to make to all Trojan rules in all your policies, this is a quick way to propagate those changes on a wide or even global scale.

Network Shield

The Network shield rule provides network protocol stack hardening capabilities. The features available here require that the network shim be enabled on an agent system. If the network shim is not enabled, these rules have no effect when applied. See the [“Network Shim Optional” section on page A-4](#).



Note

The information provided in this manual, in this section especially, assumes a basic knowledge of TCP/IP. A good source for further reading on the topic is the book *Internetworking with TCP/IP*, Douglas R. Comer and David L. Stevens, Prentice Hall, Inc.

The Network shield rule lets you enable several different types of protocol stack hardening features. In most cases, you can also turn logging on or off for each individual feature. They are as follows.

IP Security checks

- Discard invalid IP headers

Enabling this feature causes the Cisco Security Agent to perform an integrity check on the IP packet header. This includes performing a consistency check on the IP header, on the length of the IP header, and on the number of bytes in the packet. Should any of these checks fail, the packet is dropped. Additionally, if an IP checksum fails, the packet is dropped. IP options and IP fragments are validated as well and dropped if they are found to be invalid. (This defeats attacks such as Teardrop, Boink, and Ping of death.)
- Discard invalid IP addresses

IP addresses are determined to be invalid for several reasons: if the source address is a multicast address, if the TCP connection is to a broadcast address. Enable this feature to protect against these types of attacks.
- Prevent IP source routing

Enable this feature to prevent IP options which control explicit routing instructions for packets. With IP source routing (an IP header option) the originator of a packet can try to partially or completely control the path through the network to the destination.
- Prevent trace route

Enable this feature to prevent the mapping of network topology via trace route.

Transport Security checks

- Discard invalid transport (TCP/UDP/ICMP) headers

Enable this feature to ensure that transport headers are the proper length and that they are consistent (have enough data in the packet for them to fit). This enforces that certain fields have valid values and that certain combinations of TCP flags are legal. This defeats attacks such as a Christmas Tree scan.

- Detect TCP/UDP port scans

Port scanning is a common method for finding weaknesses at a site by determining what network services are being run. An attacker attempts to connect to port after port on a target system, mapping ports to identify network services and machine type vulnerabilities. Use port scan detection to log an event (one per minute) when an attempt is made to scan the system for an open port. Information is also gathered on the number of different source IP addresses perpetrating the scan and it reveals the source address of the latest scan attempt.



Note

If you select the network scans correlation checkbox in the Global Event Correlation page (see [Global Events, page 5-18](#)), when scans are detected across several machines, CSA MC correlates these events and generates an additional event to warn of this correlation.

In most cases, you should apply port scan detection to servers and end-user systems in your enterprise.

You should be aware that although port scan events are always sent to the central server (when port scan detection is enabled for a host), there is a threshold for suppressing port scan event log messages in order to avoid false positives. Once the threshold is reached, the port scan message is logged. Also note that even suppressed messages are used for event correlation purposes.



Note

This feature only DETECTS port scans, logging detected attempts but not protecting against them. For port scan prevention capabilities, you must enable the Cloak system checkbox.

- Cloak system (prevent unauthorized port scans)

This cloaking capability causes a system to not respond to connectivity tests (the system will not reply to a ping request) and to not respond to service requests with connectivity error messages.

When cloaked, the system can hide itself from view on the network. A system generally sends out error messages when a remote machine sends a request for a service which is not running on the system. Often, this is how remote machines locate other systems and obtain network information about the system in an attempt to target it for an attack. By not responding, this prevents both UDP and TCP-based port scans of the system and basically hides it on the network.

**Note**

If you are running an "allowed" server on a system and you have cloaking enabled, connection requests to this service are honored and your machine is viewable for the service you're offering.

- Prevent TCP blind session spoofing (randomize TCP sequence number)

Enabling this feature causes agents to make TCP sequence numbers unpredictable.

A server accepting connections using predictable TCP sequence numbers may be tricked into accepting a connection from a malicious source that is spoofing a trusted host. This prevents that vulnerability.

(This rule type is not available for UNIX policies.)

- Prevent TCP SYN floods

SYN flooding is a type of denial of service attack. It occurs when a TCP/IP connection request is received from a return address that is not in use (i.e. a non-existent host for a spoofed address) resulting in a half open connection. An abundance of half open states on a server can prevent legitimate connections from being established. Using SYN flood protection prevents this attack from succeeding.

(This rule type is not available for UNIX policies.)

You should apply SYN flood protection to servers that are external to your network and not protected by a firewall. Firewalls generally provide this protection.

In addition to providing protection against SYN flood attacks, information is also collected on the SYN flood by logging (once per minute) any connection request received without a response.

- Block ICMP covert channels

Enabling this feature causes agents to drop unsolicited echo responses.

The Cisco Security Agent validates that the echo response date matches the echo request data. This way, ping cannot be used as a transport for communicating.

- Block dangerous ICMP messages

Some ICMP messages may be used to gather information about a machine in an attempt to attack it (i.e. a time stamp request or a redirect). This data, when obtained, can be used to gather system information which can be used to exploit the system. Enable this feature to protect against this type of attack.

- Block malicious packets

Enable this feature to block packets which are technically legal, but are known exploits against protocol stacks (e.g. UDP packet storm or RF poison).

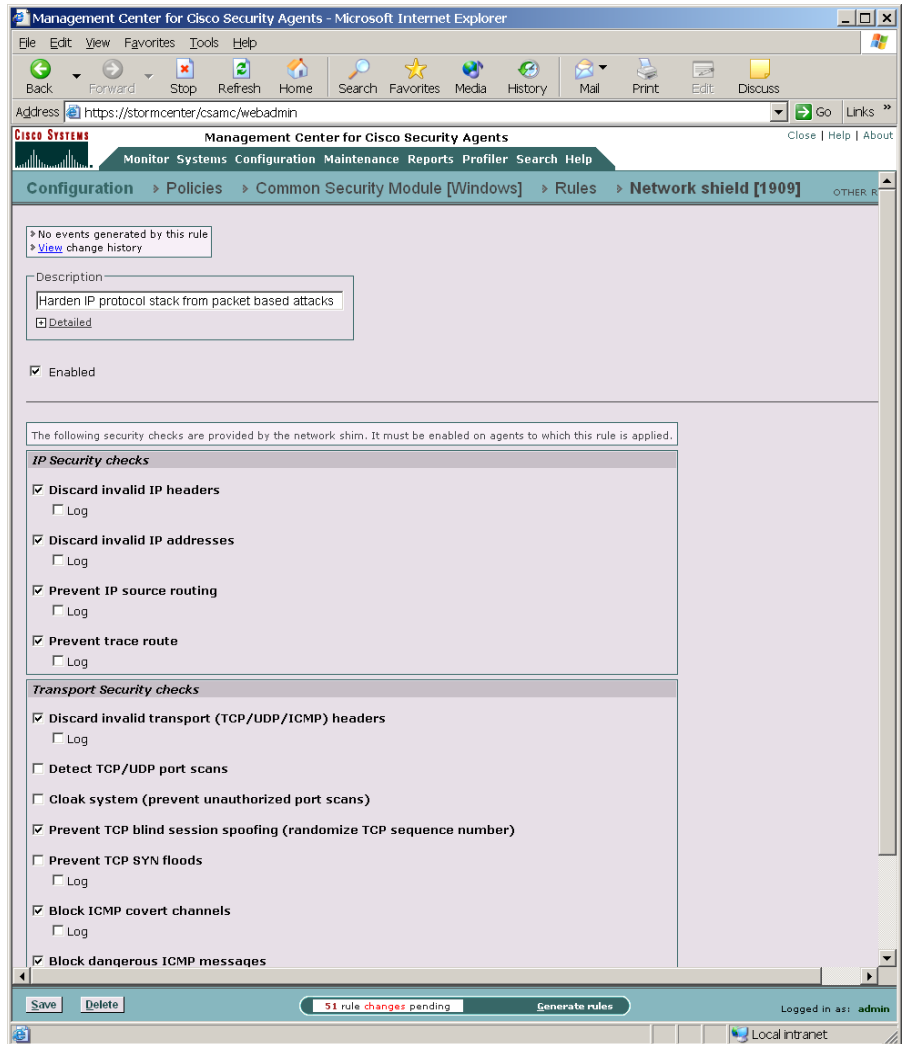
System Startup Security checks

- Restrict network connectivity

Prevent non-essential network connections during system startup. This check is automatically disabled when the agent service starts and policies (including those which govern allowed network connections) are enforced. This protects the system from network-based attacks at boot-time before the agent service has started.

(This rule type is not available for UNIX policies.)

Figure 5-3 Network Shield Rule

**Note**

Refer to [Replicate Feature, page 5-9](#) for details on easily propagating the changes you make to one Network shield rule to other Network shield rules in other policies.

Buffer Overflow

A buffer overflow is a description of what happens when two conditions are met: Firstly, an application is coded in a manner such that it trusts that all users of that application will provide the application with reasonable and expected data. Secondly, the application is provided much larger quantities of data than it is capable of correctly handling. When these events come together, an application can behave in unexpected and unintentional ways.

For applications with special privileges, this can result in external users gaining access to machine resources and privileges which they normally would not be able to acquire. In other words, a hostile, network-based attack on a privileged, trusted application via buffer overflows can result in a undesirable parties gaining access to your system.

In the case of UNIX operating systems, there are three distinct types of buffer overruns which can occur, based upon the type of memory space involved: stack, data, and heap.

- Stack space is used to store data and information which is local to the piece of code currently being executed in an application, and contains stored away control flow information for the application.
- Data space is used to store data with fixed sizes which needs to be shared among different parts of an application. Often, content in data space has been given initial values.
- Heap space is dynamically given out to applications, with the intent that it is relatively short-lived, of varying size based upon the input datasets, and is frequently visible to numerous sub-components of an application.



Note

This rule is UNIX specific. Some corresponding Windows functionality is available from the Trojan detection rule page.

Configure the Buffer overflow rule to protect one or all of the spaces described.

- Step 1** To add rules to your policy, click the **Add rule** link at the bottom of the rule list. A pop-up list of the available rule types appears.
- Step 2** Select the **Buffer overflow** rule. This takes you to the configuration view for this rule type ([Figure 5-4](#)).

Step 3 Select one or more of the following checkboxes to prevent the associated buffer overflow attack from occurring.

- Executing code in stack space

This checkbox enables the "noexec_user_stack" system variable (if not already turned on) for all processes. This prevents the execution of instructions from stack memory. You can select an application to exclude from this rule in the list box.

- Executing key system calls in unsafe contexts

Use this checkbox to prevent certain system calls (e.g. those which grant extra privileges or start new processes) from occurring if they are invoked in an unsafe manner, or if they appear to have come from a corrupted or invalid context.

- Buffer overflow by Solaris executables

Enable this checkbox to detect buffer overflow conditions which occur in Solaris executables. This feature provides protection from stack buffer overflows to a number of commonly used libc routines. As a large number of attacks on Solaris systems are based upon buffer overflow attacks, it is recommended that you enable this feature.

- Use of %n argument in printf calls

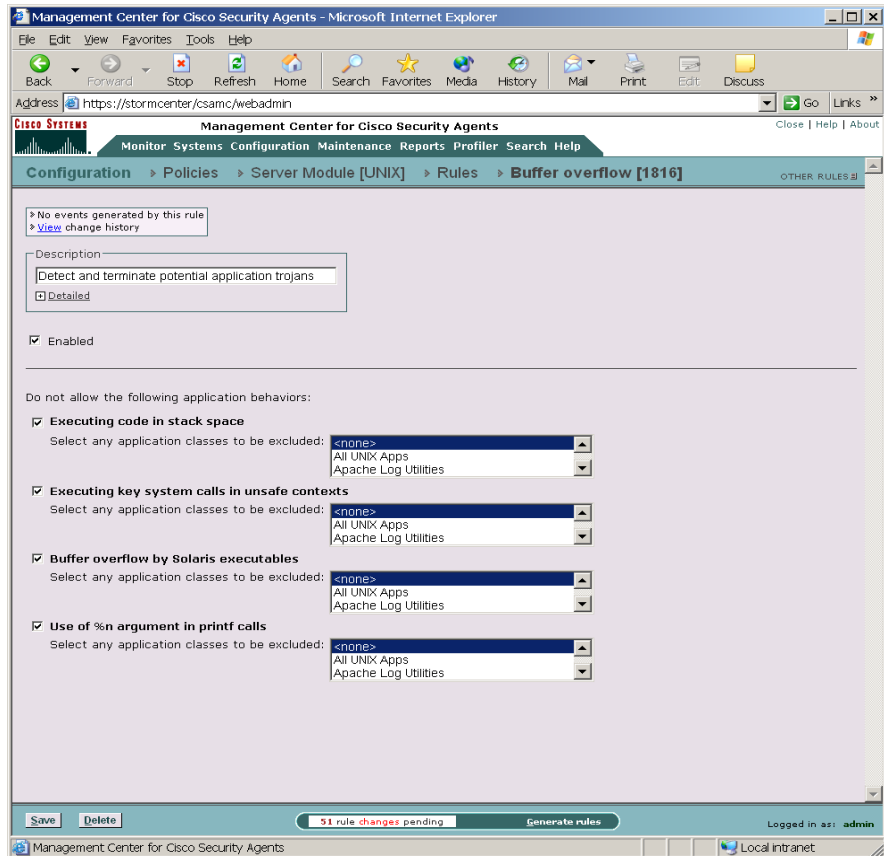
Use this checkbox to prevent the usage of the '%n' *printf() format qualifier. Numerous attacks utilize the '%n' format on *printf() routines to gain access to program control flow information.

You also have the ability to select specific **application classes to exclude** from the various Buffer overflow types you designate. If you select an application in the available list beside a checkbox rule, that rule does not apply to the selected application class. If you have multiple, similar Buffer overflow rules, the application class exceptions are combined.

**Note**

Refer to [Replicate Feature, page 5-9](#) for details on easily propagating the changes you make to one Buffer overflow rule to other Buffer overflow rules in other policies.

Figure 5-4 Buffer Overflow Rule



Global Events

The Management Center for Cisco Security Agents lets you enable correlation functions for particular types of events. In each case, you must have a corresponding rule enabled in a policy for the global event correlation to take place. If you do not enable global event correlation, individual events are logged by system agents but similar events across multiple agents are not correlated by the central CSA MC.

Correlation

The **Event Correlation** page, accessible from the menu bar as follows **Configuration>Global Event Correlation** (see [Figure 5-5](#)), provides the following capabilities:

- Correlate network scans

With this checkbox enabled, correlated port scans and ping scans across multiple agent systems are logged separately as a correlated event in addition to the individual port scan and ping scan events that continue to be logged.

Note that you must have a Network shield rule with Portscan detection enabled in a policy deployed to the agent(s) in question for these event types to be detected and logged.

The threshold and time frame for correlating network scans are values you can configure. By default, this feature is set to log a correlated event when netscans across 5 systems occur within 60 minutes.

- Correlate email worm events and add contaminated files to list of dynamically quarantined files

With this checkbox enabled, email worm events logged across multiple agent systems are correlated and the contaminated file that triggered the event is added to a dynamic list of quarantined files that CSA MC maintains for up to one hour. If you have a rule configured to stop dynamically quarantined files in a deployed policy, no further agents can access the contaminated file. See [page 4-44](#) for information on using @dynamic in File access control rules.

If you do not enable this checkbox, email worm event correlation does not take place, but individual email worm events are logged.

Note that you must have a Network worm detection rule in a policy deployed to the agent(s) in question for these event types to be detected and logged.

The threshold and time frame for correlating email worm events are values you can configure. By default, this feature is set to log a correlated event when email worms are detected across 2 systems within 60 minutes.

- Correlate events received from operating system event logs and generate a summary event
 - Log individual events in addition to summary event

With this checkbox enabled, events from multiple systems are correlated based on the NT event code, NT event severity, NT event source, and NT event log type. If 2 systems log the same NT event type within 30 minutes, a correlated summary event is logged.

Note that you must have an NT event log rule in a policy deployed to the agent(s) in question for these events to be uploaded to the CSA MC log.

If you do not enable this checkbox, NT event correlation does not take place, but individual NT events are logged in accordance with the NT event log rule you have configured.

**Note**

In this case, there is an additional checkbox (Log individual events in addition to summary events) to control whether the individual events are logged in addition to the summary event. If you do not enable this checkbox, but you do enable the Correlate events checkbox, only correlated summary events will log, NOT individual events. This can be useful if NT event log messages are filling up your CSA MC logfile.

- Correlate events received from virus scanners and add contaminated files to list of dynamically quarantined files

With this checkbox enabled, NT events logged by virus scanners running on agent systems are received and correlated by CSA MC. Contaminated files detected by virus scanners are added to a dynamic list of quarantined files which CSA MC maintains for up to one hour. If you have a rule configured to stop dynamically quarantined files in a deployed policy, no further agents can receive the contaminated file. See [page 4-44](#) for information on using @dynamic in File access control rules.

**Note**

This feature only works with Norton AntiVirus. To receive these virus events, you must have an NT event log rule in a policy deployed to the agent(s) in question for these events to be uploaded to the CSA MC logfile. In the NT event log rule, you must enter Norton AntiVirus in the Event Source field. See the [“NT Event Log” section on page 4-67](#) for details.

The threshold and time frame for correlating events received from virus scanners are values you can configure. By default, this feature is set to log a correlated event when virus scan events across 2 systems occur within 60 minutes.

**Note**

To view the files that are added to the dynamically quarantined files list, click the numbered link beside **Quarantined File Events** (only available when there are events to view). This link is located beside the last checkbox on the Global Event Correlation page. It takes you to the pertinent event log messages. Read the messages there to locate the names of quarantined files.

Figure 5-5 Event Correlation Page

The screenshot shows the 'Global Event Correlation' configuration page in a web browser. The page title is 'Management Center for Cisco Security Agents - Microsoft Internet Explorer'. The browser address bar shows 'https://stormcenter/csamc/webadmin'. The page content includes the following configuration options:

- Correlate network scans** [Events : none]
 - Log a message if systems report this event within minutes
- Correlate email worm events and add contaminated files to list of dynamically quarantined files** [Events : none]
 - Log a message if systems report this event within minutes
- Correlate events received from operating system event logs and generate a summary event** [Events : none]
 - Log individual events in addition to summary event
 - Log a message if systems report this event within minutes
- Correlate events received from virus scanners and add contaminated files to list of dynamically quarantined files** [Events : none]
 - Log a message if systems report this event within minutes

At the bottom of the page, there is a 'Save' button, a status bar indicating '51 rule changes pending', and a 'Generate rules' button. The user is logged in as 'admin'.



Using Application Classes

Overview

Access control rules are application-centric. The application classes, those shipped with CSA MC and the ones you configure yourself, are the key to the rules you build as part of your security policies.

This chapter explains the application classes shipped with CSA MC and provides instructions for creating new static and dynamically defined application classes.

This section contains the following topics.

- [About Application Classes, page 6-2](#)
- [Dynamic Application Classes, page 6-11](#)
- [Create New Application Classes from Rule Pages, page 6-20](#)
- [Application Class Management, page 6-22](#)

About Application Classes

When you create rules, you must decide which applications are performing the operations you are allowing or denying as part of the rule. Once you know this, you configure the application as an "application class" in CSA MC and select it as part of your rule.

Application classes are groupings of application executable files that you combine under one name, generally as part of a File Set Variable, see the [“File Sets” section on page 7-6](#). For example, you can enter `netscape.exe` and `iexplore.exe` under the heading of Web Browsers. Then you can select Web Browsers in the application field for your rule and apply restrictions to the actions that both Netscape and Internet Explorer can perform on specified resources.

Processes Created by Application Classes

When applications are invoked, they often spawn other processes as part of the action they are performing. Therefore, when you create an application class, CSA MC gives you the option of including or excluding child processes created by the original applications you define as part of the application class (see [page 6-6](#) for details).

Removing Processes from Application Classes

Processes are part of a configured application class when they are running on the system. When the process stops running, the CSA MC application classification for that process also ends. Should the process begin again, it may or may not fall into the same application class depending on the process's behavior and on the definition of the application class. Therefore, all application classifications are ephemeral and are constantly being re-evaluated and classified on the system.

The application class configuration page lets you control how long a process maintains a certain application classification. In general, you do not have to specify a time frame. You should only put a time limit on an application classification if you are configuring rules that require it for a particular reason. For example, you may want to create special process start rules for an application. The classification of the process could be configured to time out once the system is finished booting.

Shell Scripts and Application Classes

On UNIX systems, the agent allows control over shell scripts which satisfy both of the following conditions:

- the script begins with an interpreter string (e.g., `#!/bin/bash`)
- the script is executed directly on a command line, e.g., `"$foo.sh"`.

Therefore, if you have an application class "foo.sh", a process satisfying the above conditions becomes a member of that application class.

Note that a shell may be launched by various methods which do not meet those conditions, e.g., `"$. foo.sh"`, or `"$ cat foo.sh | /bin/sh"`. Note also that if you happen to have an application class for a script's interpreter -- say, `/bin/bash` -- when you invoke the script, the process becomes a member of the `/bin/bash` application class.

If a user has write access to the disk, and can execute commands, then using the name of a shell script in a rule to DENY actions may not make sense. For example, denying access by `foo.sh` to modify `/etc/hosts` does not improve the protection of `/etc/hosts` as the user could just run `'vi /etc/hosts'`. It would make more sense to deny everything access to a file, and then permit known good scripts access to that file.



Caution

If the user can copy a script (or re-implement it) to a file of their choice, then any Deny rules would be avoided.



Note

On Windows, when writing rules for script application classes, you can create the rule for either the script itself or for the interpreter. (Scripts are handled by script interpreters.) If you write the rule for the interpreter, it will include the script handled by that interpreter.

Included Application Classes

CSA MC ships with several pre-configured application classes. Those application classes appear inside brackets (see [Figure 6-1](#)) in the rule application class selection list boxes. You can view them in the Application Class list page, but you cannot edit them in the application class configuration view.

Figure 6-1 Included Application Classes



Non-configurable application classes include:

- **Network Applications:** A network application would include any process that connects as a client or accepts a connection as a server and has in some manner accessed the network. The process would fall into this network application class after it has accessed the network. (This does not include applications that communicate only with other applications on the same system.)
- **Processes created by Network Applications:** This includes any process that is launched by a network application. For example, one network process may create another process that attempts to download code. This is one way viruses are propagated.
- **Processes executing downloaded content:** This includes any downloaded executable or any process that is interpreting downloaded content. This ships as a pre-configured "built-in" application class. But unlike other built-ins, you can change the definition of downloaded content. See [Configure Downloaded Content, page 6-9](#) for details.
- **Processes created by Servers (TCP and UDP):** This includes any TCP or UDP process invoked by a server (falling into the categories detailed in the two following bullet points).

- Server (TCP based): This application class includes all processes that have accepted a connection on TCP port numbers under 1024. It also includes all processes that have accepted a connection on TCP port numbers greater than 1024 if the server makes a subsequent outgoing connection.
- Server (UDP based): This application class includes all processes that have received a datagram on UDP port numbers under 1024.
- Processes with elevated privileges: This application class is only available for UNIX file and application control rules. It includes processes that have elevated user privileges for users other than root, such as ping. Using such processes is a common way to attempt a system break-in. Note that this elevated privilege designation does not apply to processes when the user is logged in as root.
- Remote clients (Remote application exception: available only in File access control rules and COM component access control rules): When a remote machine accesses resources over the network that are protected locally by an agent, the agent sees the remote access attempt as coming from a "remote application." The actual application that is used to open the resource in question cannot be determined on the local system. All remote access attempts are seen by the local system as being invoked by a remote application.

Therefore, if you are writing rules for file shares or COM components on a machine that other machines can access over the network, you must include <All Applications> or <Remote clients> as your application class. Otherwise, the rule will not work as expected in regard to remote access to those resources.

- System Process (available only in Network Access Control rules): Using this application class, you can control network access for the operating system itself (as opposed to applications running on the operating system).

**Caution**

Any application class that you define does not include the system process. If you want to include the system process in a rule, you must select the included, built-in <All applications> or <System process> classes.

Preserving Application Process Classes

You should be aware that all application process classes are preserved when your policies are changed if those processes (application classes) are used in an existing policy. For example, processes that have been classified by CSA MC as descendants or as network applications are preserved if the application classes that included them are changed in any way.

On policy changes, process name-based application classes are re-evaluated. Old application class memberships are not lost, only new memberships are gained.

Configuring Static Application Classes

Access control rules are application-centric. Meaning that when you write your rules, you should understand that the application(s) you select are really the heart of each rule. In your file, network, registry, and COM rules, you are controlling what applications can do to the files, addresses, registry keys, and COM components you specify. So, when you begin creating rules, think in terms of the applications your enterprise as a whole uses and the manner in which you want to limit an application's ability to perform undesired actions.

See also [Included Application Classes, page 6-4](#).

To create an application class, do the following:

-
- Step 1** Move the mouse over **Configuration** in the menu bar and select **Application Classes** from the drop-down list that appears. The list of existing Application classes is displayed. CSA MC ships with several pre-configured applications. Some Application classes appear within brackets. These are built-in CSA MC application classes and you cannot edit them.
 - Step 2** Click the **New** button to create a new application class. This takes you to the application class configuration view (see [Figure 6-2](#)).



Note If you have "All" designated as the operating system type for your administrator session, you are prompted to select whether this is a Windows or a UNIX application. See the [“Administration by Operating System” section on page 2-5](#) for details.

- Step 3** Enter a **Name** for the application class you are creating. It is important to use a descriptive name that you can easily recognize in the application selection list that appears in the rule views.
- Step 4** Enter a **Description** for your application class. This description becomes visible in the application class list view.
- Step 5** **Enable this application class for the following selected products:**
You can choose to have an application class not appear in application class fields in rules or in the Profiler utility. For example, if you have a long list of application classes and you only want to view specific classes in rule pages or only view them in Profiler pages, you can choose to have application classes appear or not appear in products you select. By default, all application classes appear in all application class fields in all installed products.



Note You can also change the products for which application classes are enabled from the **Application Class Management** window available from **Configuration** in the menu bar. See [Application Class Management, page 6-22](#).

- Step 6** Under **Add process to application class**, for a static application class, do the following:
- Leave the default **when created from one of the following executables** radio button selected. Then enter the executable file names (one per line) for the applications you are grouping together in this application class.
- See [Configuring Dynamic Application Classes, page 6-13](#) for details on that feature.



Note You can enter preconfigured File Set variables in the executables edit field by clicking the **Insert File Set** link. To learn more about File Sets, see the “[File Sets](#)” section on page 7-6.

Step 7 Remove process from application class: Select the checkbox beside **After** and enter a time frame in **seconds** to configure an application classification which expires after a period of time. Only use this feature if you have a rule that requires it. In general, you do not have to specify a time-out for application classifications. See [Removing Processes from Application Classes, page 6-2](#) for details.

For UNIX application classes, you have the additional option of selecting the **When session association is voided** checkbox. Selecting this checkbox causes the application classification to be removed when a process disassociates itself from the current TTY session. For example, when an application class exists for applications descended from "superuser", you might not want the process to continue having the application class of the superuser shell.

Step 8 When applications are invoked, they often spawn other processes as part of the action they are performing. When you create an application class, select one of the following radio buttons to determine when processes spawned by the applications in the application class are also included.

- Only this process
- This process and all its descendents
- Only descendents of this process

(Creating an application class for "Only descendents of this process" is useful when making exceptions to a rule that is written for the main process itself. For example, you can write a rule allowing IIS to talk on the network, but create another rule denying descendents of the IIS process from talking on the network.)

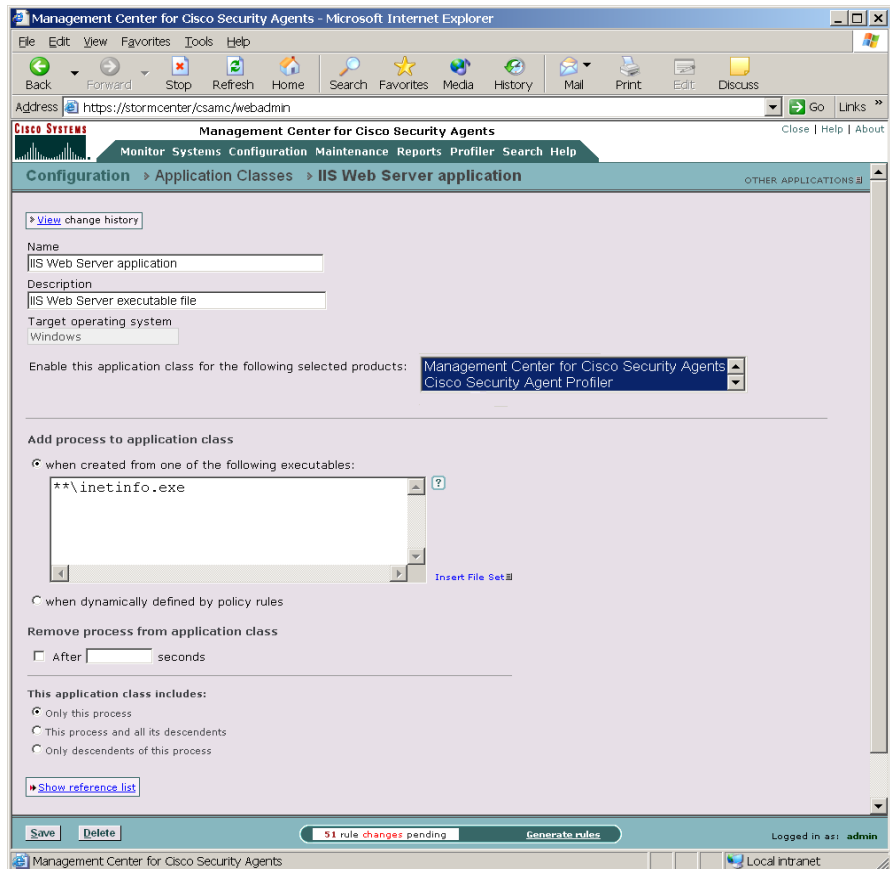
Step 9 When you are finished, click the **Save** button. This application class name, IIS Web Server application, now appears in the application list view and in the application selection fields for rule configurations. When you select it in a rule, you are indicating all the executables that comprise it.



Note You can use the Compare button in the Application Class list view to compare and merge similar application classes. See the [“Comparing Configurations”](#) section on [page 4-19](#) for details on using the Compare tool.

See [Dynamic Application Classes, page 6-11](#) for information on that application class type.

Figure 6-2 Static Application Class



Configure Downloaded Content

CSA MC ships with a pre-configured application class called <Processes executing downloaded content>. This class includes any downloaded executable or any process that is interpreting downloaded content. CSA MC uses the downloaded application class in several pre-configured policies as a way to track executables or scripts that have been modified by "network aware" applications.

It is via the network that most viruses are delivered to systems. Therefore, policies are generally configured to track data of this type to ensure that it is not malicious content intent on doing harm to your networks or systems.

When you access the page for downloaded content, you can see the following configuration. A process is considered executing downloaded content when:

- It belongs to any of the following selected application classes: Download Directory Executables (pre-configured selection)

This lets you specify certain applications as inherently untrusted. The preconfigured entry, Download Directory Executables, includes files located in the system cache and temp directories. These directories are where email attachments and browser files are stored.

- It is interpreting scripts and belongs to any of the following selected application classes: Network Applications, Web browser applications (pre-configured selection)

This is intended to capture processes that download scripts and execute them internally. For example, this could be a web browser downloading and executing java script from a remote site.

- Its executable was written to disk by any of the following selected applications classes: Network Applications, Email applications (pre-configured selection)

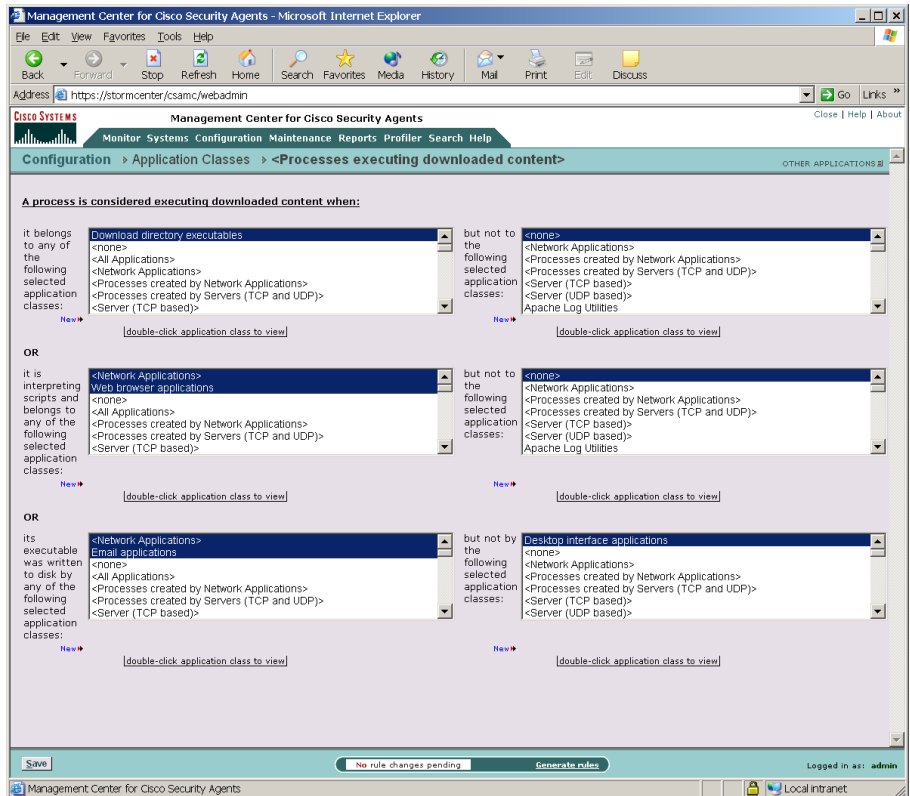
This is intended to capture executables written by network-aware applications and possibly downloaded from the network.

If necessary, you can edit any of these downloaded content fields and make any necessary exclusions to change the global definition of downloaded content. You may want to do this if you are experiencing false positives due to an application being seen as downloaded content and therefore possibly dangerous when this is not the case. On the other hand, you may want to use this page to add file share names or removable media so that data accessed from the share or media type is tracked as downloaded content. You can edit this page to make your policies more or less restrictive in regards to tracking downloaded content and dictating what it can and cannot do.

**Caution**

You should not edit the downloaded content application class unless you fully understand how it works with your policy rules. In most cases, the pre-configured downloaded content application class should be sufficient.

Figure 6-3 Downloaded Content Configuration



Dynamic Application Classes

The configurable application classes described in the previous pages are considered static application classes. Basically, in a static application class, a process is added to the class based on the name of its executable file (or the process name). Alternatively, you can build an application class based on an application's behavior rather than by a specific application executable name. This would be a dynamic application class defined by process behavior on a system.

There are already built-in dynamically defined application classes in CSA MC. For example, the <Processes executing downloaded content> application class is a "built-in" dynamically defined class.

One example of an instance in which you might need a dynamic application class would be if you are writing rules for email clients but you do not know all the different email applications that are being used throughout your corporate network. In this case, you could use a dynamic application class. Any process appearing to act as a client for SMTP (you can use whatever criteria you decide to define what an email application is) would fall into a dynamic email application class that could be used in rules quarantining dangerous email messages.

**Note**

A dynamically defined application class can be used in any rule where a static application class can be used.

Define a dynamic application class by doing the following:

- Create a new application class and select the **Processes dynamically defined by policy rules** radio button. (Do not enter any process names in the Application class page edit field.)
- Configure an application-builder rule to define your dynamic application class.

**Note**

Configuring the dynamic application class is only the first step. It does not become populated by processes until it is selected in a rule that will be used to define it.

For example, create a new File access control rule and select the **Add to application class** radio button as the rule action. Then choose the name of the dynamic application class (created in the first bullet point) from the pulldown list. Configure the remaining rule parameters. This rule type takes precedence over all others in the policy, but it does not override other rules in the policy the way allow, deny, and query rules do when triggered.

- Configure another rule to control the actions of this dynamic application class. As processes are added to this dynamic application class, those same processes will be used in all other rules in which the dynamic class is selected.

The following section provides an example of defining and using a dynamic application class in a policy

Configuring Dynamic Application Classes

Continuing to use the email client example, we will create an application class that will be dynamically populated by email client applications. You might want to do this if you are writing rules to protect email applications, but you do not know what email applications are being used across your network. Using this dynamic class, rules will restrict email clients based on detected behavior, such as using SMTP to access an email server, rather than by explicitly defining email application executables.

To create a dynamic application class, do the following:

-
- Step 1** Move the mouse over **Configuration** in the menu bar and select **Application Classes** from the drop-down list that appears. The list of existing Application classes is displayed.
 - Step 2** Click the **New** button to create a new application class. This takes you to the application class configuration view (see [Figure 6-4](#)).



Note If you have "All" designated as the operating system type for your administrator session, you are prompted to select whether this is a Windows or a UNIX application.

- Step 3** Enter a **Name** for the dynamic application class you are creating. It is important to use a descriptive name that you can easily recognize in the application selection lists that appear in the rule views.
For this example, we will create a new dynamic class called *Email clients_dynamic*. We will use this class to determine what email client applications are running on systems. Then we will add this dynamic class to an existing email quarantine rule.
- Step 4** Enter a **Description** for your application class.

- Step 5** Under **Add process to application class**, for a dynamic application class, do the following:
- Select the **when dynamically defined by policy rules** radio button. (Do not enter any process names in the edit field.)

- Step 6** **Remove process from application class**: Select the checkbox beside **After** and enter a time frame in **seconds** to configure an application classification which expires after a period of time. Only use this feature if you have a rule that requires it. In general, you do not have to specify a time-out for application classifications. See [Removing Processes from Application Classes, page 6-2](#) for details.

For UNIX application classes, you have the additional option of selecting the **When session association is voided** checkbox. Selecting this checkbox causes the application classification to be removed when a process disassociates itself from the current TTY session. For example, when an application class exists for applications descended from "superuser", you might not want the process to continue having the application class of the superuser shell.

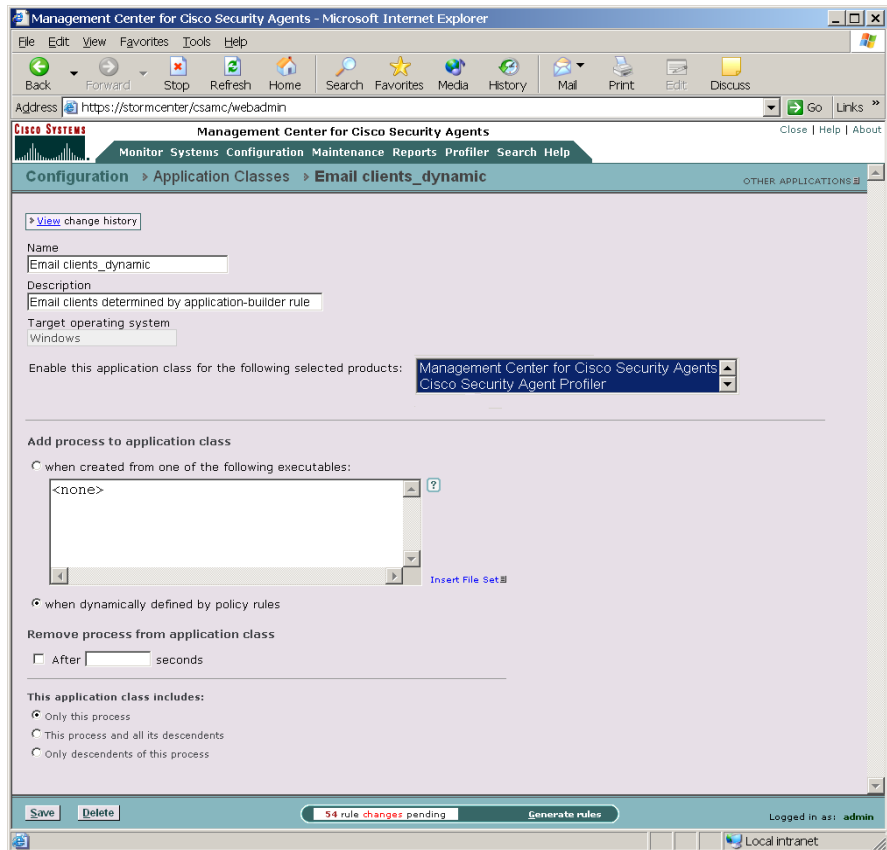
- Step 7** When applications are invoked, they often spawn other processes as part of the action they are performing. When you create a dynamic application class, you can select one of the following radio buttons (just as you can when you create a static application class) to determine when processes spawned by the applications in the dynamic application class are also included.
- For this example, we will leave the default, Only this process, selected.

- Only this process
- This process and all its descendents
- Only descendents of this process

- Step 8** When you are finished, click the **Save** button. This dynamic application class name now appears in the pulldown list beside the **Add to application class** radio button in access control rules and in all application selection fields.

Next we will use this dynamic class in an application-builder rule that will define the class.

Figure 6-4 Dynamic Application Class



Configure an Application-Builder Rule

In this example, we are going to use a Network access control rule to define our dynamic application class. You can use any access control rule type as your application-builder rule. We are adding this rule to the Desktop Module that ships with CSA MC. (Remember, your dynamic application class is not populated with applications until an application-builder rule is triggered by the process's behavior and added to the class.)

**Note**

Creating dynamic application classes from the Application control rule is a bit different than creating them from other rule types. Because this rule has two application class fields, you can choose to add the current application to the dynamic class or choose to add the new application that is invoked by the first application to the dynamic class.

**Caution**

Dynamic application class process membership is temporary and is based on a running process meeting the criteria in the application-builder rule. When the process is no longer running on a system, it is no longer included in the dynamic class.

To prevent errors or unexpected behavior, you should avoid selecting the dynamic application class for a rule within a policy that does not also include the corresponding application-builder rule. Both the application-builder rule and the subsequent rule(s) that use the dynamic application class should co-exist within the same policy—although this is not required.

Step 1

To configure the application-builder rule which will dynamically create a new Email client class, access the W-Desktop Module policy and click the **Modify rules** link.

**Note**

This is only an example. This is not intended to recommend that you add this rule to the Desktop Module. This simply shows you what you can do if you do not know all the Email clients being used on systems across your network.

Step 2

Click **Add rule** and select **Network access control**.

- Step 3** In the Network access control rule, configure the following (see [Figure 6-5](#)):
- Enter a description
 - Select the **Add process to application class** radio button. Select the dynamic application class, *Email clients_dynamic*, from the corresponding pull-down list.



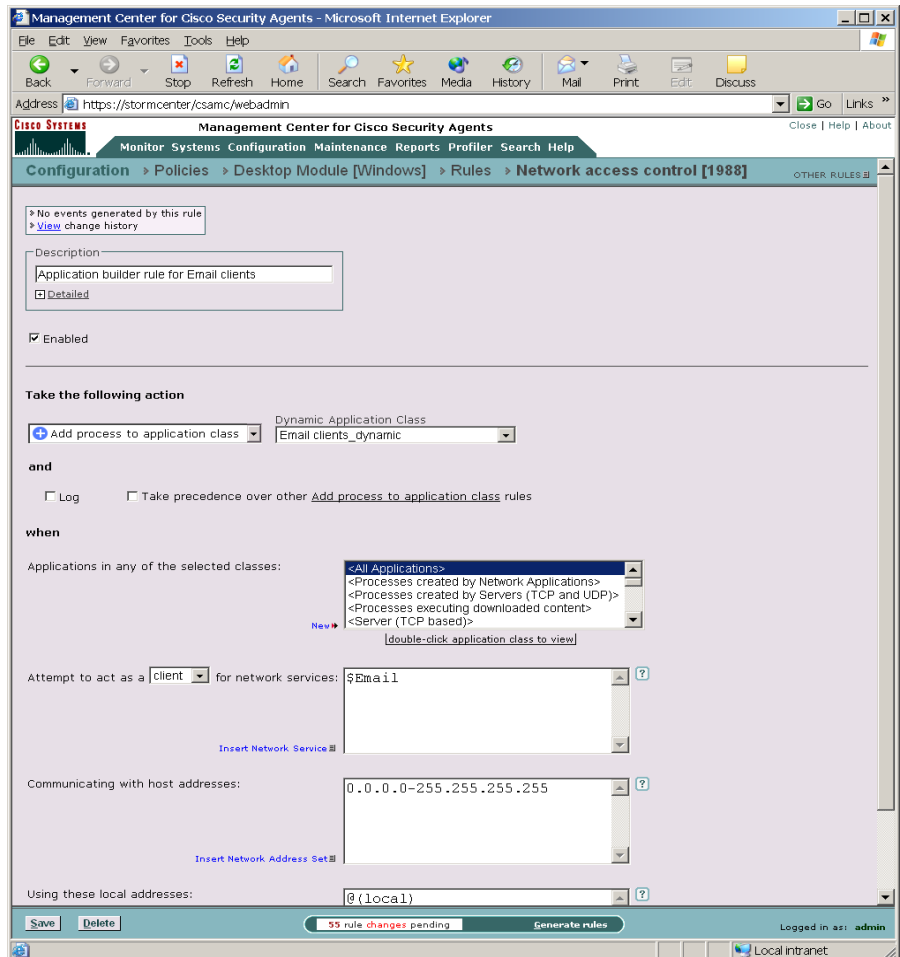
Note This rule type takes precedence over all other types but it does not override them. The only action of this rule is to build the application class for any subsequent rules within the policy that make use of it.

- Leave the default, **<All Applications>**, selected in the Application class field. This way, all applications that trigger the rule have the potential of being added to the dynamic class. You could select another application class here if you only want specific applications to fall into the dynamic class.
 - Select **client** from the pull-down list and select the pre-configured variable, \$Email, from the list of configured Network services.
 - Leave the default of 0.0.0.0-255.255.255.255 entered in the host addresses field.
- Step 4** Click **Save**.

Now, based on the application-builder rule we've just configured, any application which uses the network services, SMTP, POP3, IMAP3 or IMAP2 as a client to access any system on the network, will fall into the *Email clients_dynamic* application class.

Next we will select this dynamic application class in a rule within this same policy.

Figure 6-5 Application-Builder Rule



Configure a Rule Using a Dynamic Application Class

In this example, we are going to use a File access control rule to control the actions of a dynamic application class.

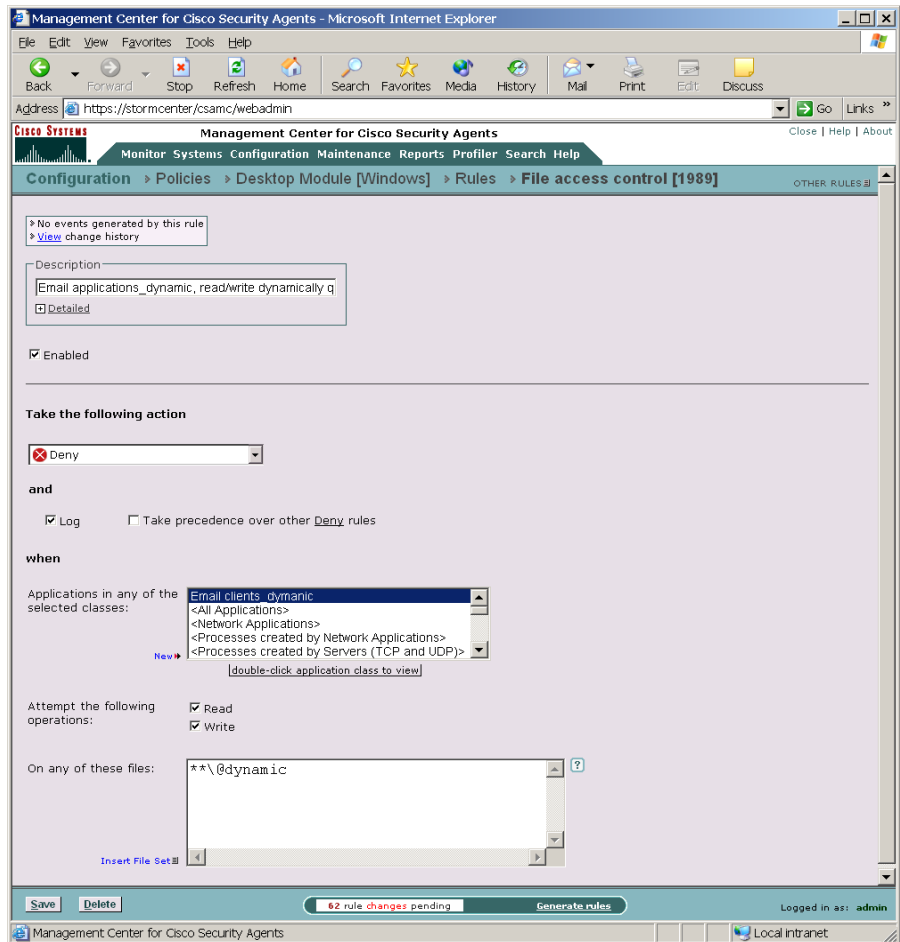
-
- Step 1** Configure this rule in the same manner in which you configure any other rule. For this example, once again access the W-Desktop Module policy (This is the policy already containing our application-builder rule.) and click the **Modify rules** link.
- Step 2** Click **Add rule** and select **File access control**.
- Step 3** In the File access control rule, configure the following (see [Figure 6-6](#)):
- Enter a description
 - Select the **Deny** radio button.
 - Select the dynamic application class, **Email clients_dynamic**, in the Application class list box.
 - Select the **read** and **write** checkboxes.
 - Enter **@dynamic** in the files field.
- Step 4** Click **Save**.

This rule will prevent any email application that falls into the selected dynamic email client class from reading or writing any dangerous, quarantined files.

**Note**

Because of the manner in which rule precedence works, the application-builder rule takes precedence over all others and it builds the application class for the rules that follow it such as the File access control rule in this example. (The application-builder rule does NOT override other rules as allow, deny, and query rules do when triggered.)

Figure 6-6 Rule with Dynamically Defined Application

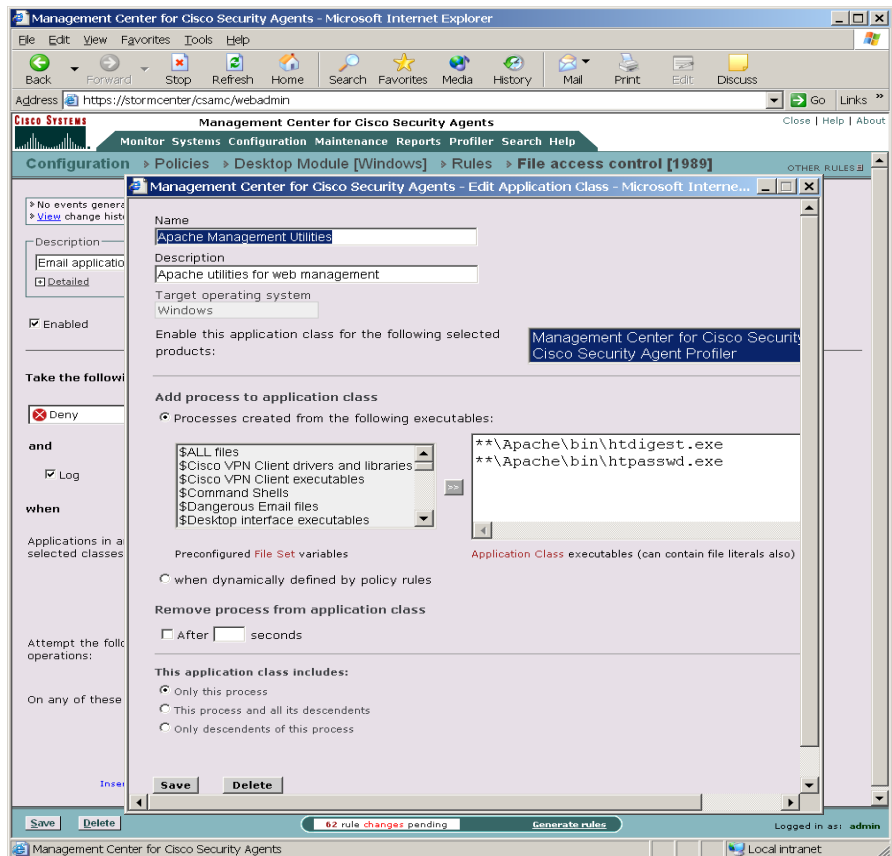


Create New Application Classes from Rule Pages

You can create a new application class from a rule page and have that application class be available to the rule you're currently configuring and to all other rules as well.

From the rule page, click the **New** link beside the Application class selection field to access configuration window (see Figure 6-7). Configure your new application class and click **Save**. It is now available for selection in the rule page.

Figure 6-7 New Application Class—Rule page



Also available for Application classes from the rule page, is the ability to view the configuration parameters for a selected application class. Double-click an application class in the rule page to view its configuration page.

Application Class Management

The Application Class Management page (available from the **Configuration** option in the menu bar) allows you to pare down the application class selection fields in the rule pages and in the Profiler utility. If you have a long list of application classes and you only want to view specific classes in rule configuration pages or only view them in the rule pages or only view them for Profiler, you can choose to have application classes appear or not appear in products you select.

Note that selecting certain application classes to not appear in certain products does not delete those application classes. They will still appear in the main Application Class list page. They simply will not appear in the application class selection fields in the product in question.

By default, all application classes appear in all application class fields in all installed products. You can decide what products an application class will appear in at the time you configuration the application class or you can do it from this main page. The result is the same.

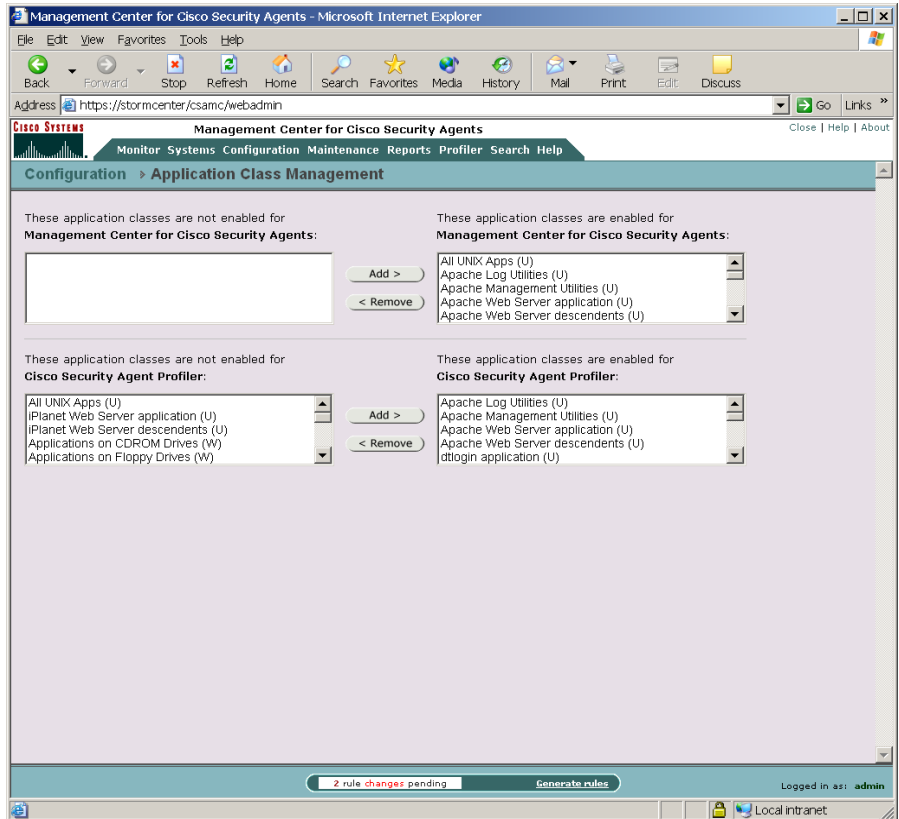
To enable or disable an application for the main product or profiler, do the following:

-
- Step 1** Move the mouse over **Configuration** in the menu bar and select **Application Class Management** from the drop-down list that appears. In the Application Class Management page (see [Figure 6-8](#)) there are swap box fields for CSA MC and Profiler.

The application classes appearing in the swap box(es) on the right side of the window are enabled for the task in question. Those appearing in the left swap box(es) are disabled, or will not appear, in the task in question.

- Step 2** Select an application class and click the **Add** or **Remove** buttons to move the selected class to the other swap box. This action enables or disables the application for the product. (It does not delete the application class.)

Figure 6-8 Application Class Management Window





Configuring Variables

Overview

Configuration variables are named configuration data items that you create for repeated use in other configuration items such as file access control rules, network access control rules, and alerts. You can group files together, as well as network addresses, and network services. Once configured, you enter these global variables in corresponding fields for other CSA MC items.

You use configuration variables to help build the rules that form your policies. Using variables makes it easy for you to maintain policies by letting you make any necessary modifications in one place and having those changes instantiated across all rules and policies.

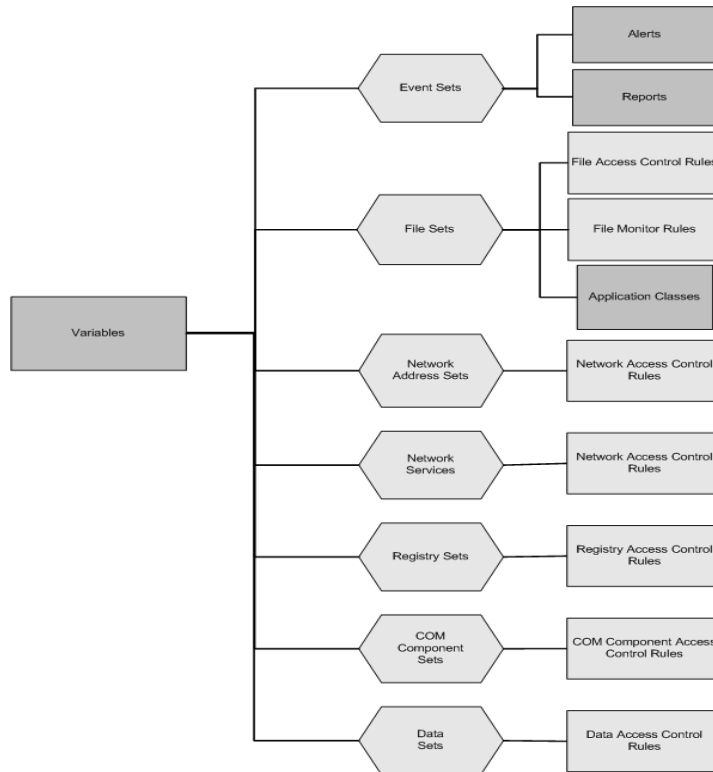
This section contains the following topics.

- [Where Variables are Used, page 7-2](#)
- [Data Sets, page 7-3](#)
- [Network Address Sets, page 7-11](#)
- [Network Services, page 7-14](#)
- [Registry Sets, page 7-17](#)
- [COM Component Sets, page 7-21](#)

Where Variables are Used

Figure 7-1 displays how variables relate to access control rules. In the diagram, variables (Event Sets, File Sets, Network Address Sets, Network Services, Registry Sets, COM Component Sets, Data Sets) are shown on the left and the rule types they can be applied to are shown on the right.

Figure 7-1 Variable Use in Rules



**Note**

Using variables is optional (note that Application Classes are included in this diagram, but they are not optional). Nearly all the information used in variable configurations can also be entered directly into corresponding rule configuration fields. Variables are simply a tool meant to simplify the creation of rules, especially if the same configurations are used in multiple rules.

**Note**

You can use the Compare button in Variable list views to compare and merge similar variables. See the [“Comparing Configurations” section on page 4-19](#) for details on using the Compare tool.

See [Chapter 8, “Event Logging Alerts”](#) for details on configuring Event Sets.

Data Sets

Configure data sets for use in data access control rules. Data sets are groupings of data strings under one common name. These strings represent a set of patterns that will be matched against the URI portion of HTTP requests. The name of the data set is then used in rules that control data access permissions and restrictions. All the data parameters that exist under that name are then applied to the rule where the name is used.

CSA MC ships with several pre-configured Data Sets you can use. The pre-configured data sets group patterns to match based upon the following:

- functional associations of meta-characters (e.g. "(" and ")")
- examples of known classes of attacks
- Web server specific exploits

The following is an example of an HTTP request attempting to execute an attack by invoking a command shell to obtain a directory listing. A data set of this syntax, *cmd.exe*, would stop not only this exploit but any other exploit trying to make use of a command shell.

```
GET /scripts/..%25c%25c../winnt/system32/cmd.exe?/c+dir
```




Note Not all pre-configured data sets are used in pre-configured policies. For example, some attack fingerprints or command arguments might be acceptable on one deployment of a web server, but not be acceptable for a different deployment. Therefore, pre-configured data sets used in shipped policies may require modification if legitimate, but blocked meta-characters are being used by a web server.



Note Additionally, modifying the preconfigured data sets allows you to block a pattern which specifically matches a new/old exploit or attack.

To configure a data set, do the following.

Step 1 From the menu bar **Configuration** drop-down list, move the mouse over **Variables**. A cascading menu with further selections appears.



Step 2 Select **Data Sets** from the cascading menu. Any existing data set configurations are shown.

Step 3 Click the **New** button to create a new data set. This takes you to the data set configuration view.

Step 4 In the available edit fields, enter the following information. (Note that you can click the Quick Help question mark beside each field for syntax information.):

- **Name** This is a unique name for this data set. Generally, it's a good idea to adopt a naming convention that lets you quickly enter data set names in a corresponding rule configuration field.
- **Description** This is a line of text that is displayed in the list view helping you to identify this particular data set configuration.

Step 5 **Patterns matching**— Enter the data strings here (one per line) to which you want to impose restrictions. By default, this field has an <all> entry indicating all strings. When you click inside this field, the <all> disappears so that you can enter your own data. This pattern is used by HTTP Web servers to match against the request URI (Uniform Resource Identifier) to enforce allow/deny Data access control rules.



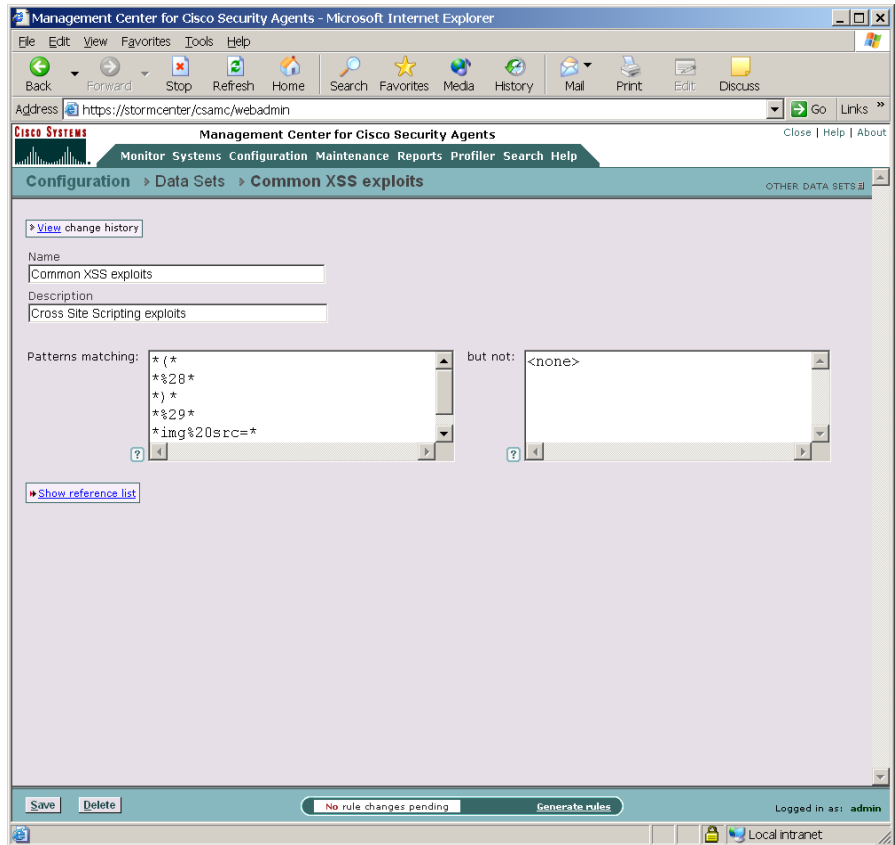
Note When entering data patterns, the “*” character is a generic wildcard specification.

Step 6 **but not**—Make exceptions to the data strings you’ve entered in the directories matching field. By default, this field has a <none> entry indicating no exceptions. When you click inside this field, the <none> disappears so that you can enter your own exceptions.

Step 7 When all required information is entered, click the **Save** button to save your data set in the CSA MC database.

You can now enter this data set name by clicking the **Insert Data Set** link in the data access control rule files field.

Figure 7-2 Data Set Configuration View



File Sets

Configure file sets for use in file access control rules, file monitor rules, and application classes. File sets are groupings of individual files and directories under one common name. This name is then used in rules that control file permissions and restrictions. All the file parameters that exist under that name are then applied to the rule where the name is used.

CSA MC ships with several pre-configured File Sets you can use.

To configure a file set, do the following.

-
- Step 1** From the menu bar **Configuration** drop-down list, move the mouse over **Variables**. A cascading menu with further selections appears.
 - Step 2** Select **File Sets** from the cascading menu. Any existing file set configurations are shown.
 - Step 3** Click the **New** button to create a new file set. This takes you to the file set configuration view (see [Figure 7-3](#)).



Note If you have "All" designated as the operating system type for your administrator session, you are prompted to select whether this is a Windows or a UNIX file set.

- Step 4** In the available edit fields, enter the following information (See the [“Using the Correct Syntax”](#) section on page 2-18. You can also click the Quick Help question mark beside each field for syntax information.):
 - **Name** This is a unique name for this file set. Generally, it’s a good idea to adopt a naming convention that lets you quickly enter file set names in a corresponding rule configuration field. When using configuration variables in file access rules, network access rules and application classes, you must enter the variable name preceded by a dollar sign:

For example, if you have a file set variable named `cgi_files`, you must enter `$cgi_files` into any edit field where you are using this variable. The dollar sign tells CSA MC that this is a variable value.
 - **Description** This is a line of text that is displayed in the list view helping you to identify this particular file set configuration.

Step 5 Directories matching—Enter the directories and files here (one per line) to which you want to impose restrictions.

By default, this field has an <all> entry indicating all directories. When you click inside this field, the <all> disappears so that you can enter your own directory restrictions. When entering directory restrictions, use the following syntax (See the [“Using the Correct Syntax”](#) section on page 2-18):

Windows example:

```
c:\Winnt\system32\*
*:\Program Files\Netscape\**
```

UNIX example:

```
/user/bin/*
/user/adm/sg/**
```



Note You can protect directory paths as well as files on UNIX systems (You cannot do this on Windows.) You should use a file literal rather than a file set variable if you are protecting directories. See the [“Using the Correct Syntax”](#) section on page 2-18 for details.

Step 6 but not—Make exceptions to the files and directories you’ve entered in the directories matching field. For example:

Windows example:

```
c:\Winnt\system32\temp
```

This means that any temp files in the system32 folder and its subfolders are ignored by the restrictions you apply using this file set.

UNIX example:

```
/etc/passwd
```

By default, this field has a <none> entry indicating no exceptions. When you click inside this field, the <none> disappears so that you can enter your own exceptions.

Step 7 Files Matching—Enter the names of the files you are controlling access to.

You can use wildcards here to indicate all of a specific file type. For example, *.exe to specify all executables.

By default, this field has an <all> entry indicating all files. When you click inside this field, the <all> disappears so that you can enter your own file restrictions.

Step 8 but not—Make exceptions to the file names you enter in the Files Matching field. For example, all executables, but not regedit.exe.

By default, this field has a <none> entry indicating no exceptions. When you click inside this field, the <none> disappears so that you can enter your own exceptions.



Note Use **@dynamic** (Windows only) in the File set text field to indicate all files that have been quarantined by CSA MC as a result of correlated email worm events or correlated virus scanner log messages. Files are quarantined by CSA MC for up to one hour. This list updates automatically (dynamically) as logged quarantined files are received.

To view the files that are added to the dynamically quarantined files list, click the numbered link beside **Quarantined File Events**. This link is located beside the last checkbox on the Global Event Correlation page. It takes you to the pertinent event log messages. Read the messages there to locate the names of quarantined files.

See the [“Global Events” section on page 5-18](#).

Step 9 (UNIX only instruction) File Sets created for UNIX have an additional configuration field. In the **Content Matching** edit fields, click the **Insert content** link and optionally select one or more file types to match against. Available file types are as follows:

- executable object: All binary executable files
- interpreter file: All files that begin with #!
- java class file: All java class file types

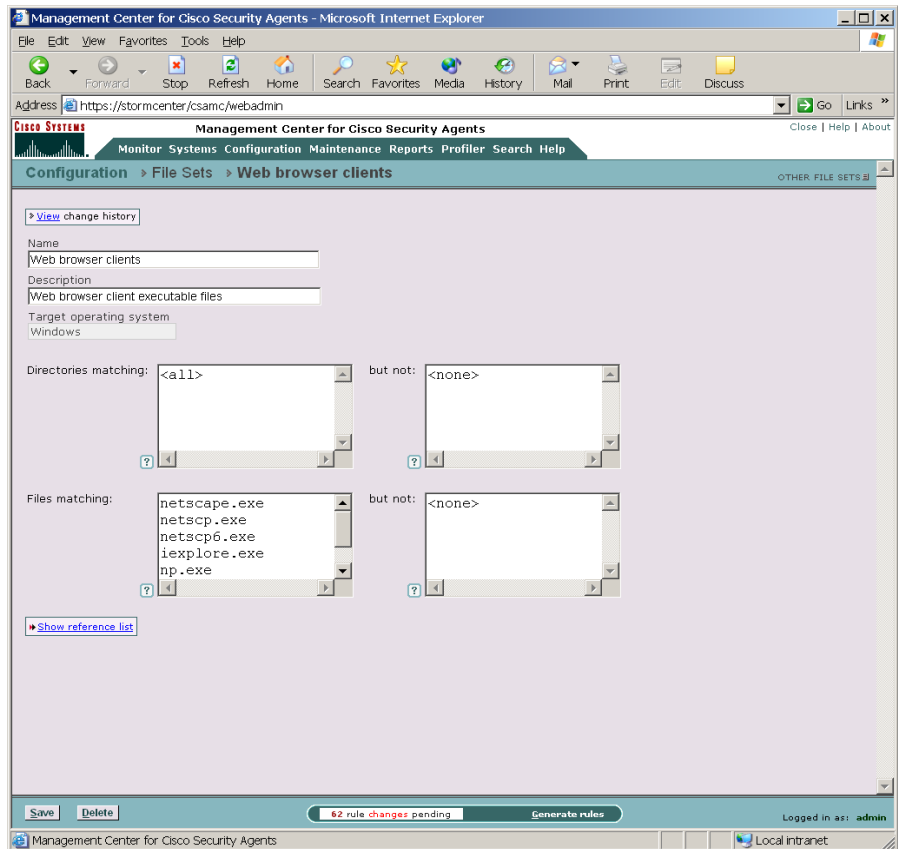
Step 10 When all required information is entered, click the **Save** button to save your file set in the CSA MC database.

You can now enter this file set name by clicking the **Insert File Set** link in the application class files field and in the file access control rule files field.



Note At the top of each variable page, there is a **View change history** link. Click this link to go to a page which lists all the changes that have been made to the item in question. This View change history link is also available for application classes, policies, and rules.

Figure 7-3 File Set Configuration View



Network Address Sets

Configure network address sets for use in network access control rules to impose restrictions on specified IP addresses or a range of addresses. Once configured, you can simply enter the name of the address set in any network access control rules you create.

To configure network address sets, do the following.

-
- Step 1** From the menu bar **Configuration** drop-down list, move the mouse over **Variables**. A cascading menu with further selections appears.
 - Step 2** Select **Network Address Sets** from the cascading menu. Any existing address set configurations are shown.
 - Step 3** Click the **New** button to create a new network address set. This takes you to the configuration view (see [Figure 7-4](#)).
 - Step 4** In the available edit fields, enter the following information (See the [“Using the Correct Syntax” section on page 2-18](#). You can also click the Quick Help question mark beside each field for syntax information.):
 - **Name** This is a unique name for this address set. When using configuration variables in file access rules, network access rules and application classes, you must enter the variable name preceded by a dollar sign:

For example, if you have a network address set variable named Finance systems, you must enter `$Finance systems` into any edit field where you are using this variable. The dollar sign tells CSA MC that this is a variable value.
 - **Description** This is a useful line of text that is displayed in the list view and helps you to identify this particular set of addresses.

Step 5 Enter address ranges—In the available edit field, enter a single address, range of addresses.

By default, this field has a <none> entry indicating no addresses. When you click inside this field, the <none> disappears so that you can enter your own addresses. When entering directory restrictions, use the following syntax:

Put each entry on its own line. For address ranges, use a hyphen to indicate the range. Address ranges are inclusive.

For example: 128.66.24.130
 128.67.2.10-20

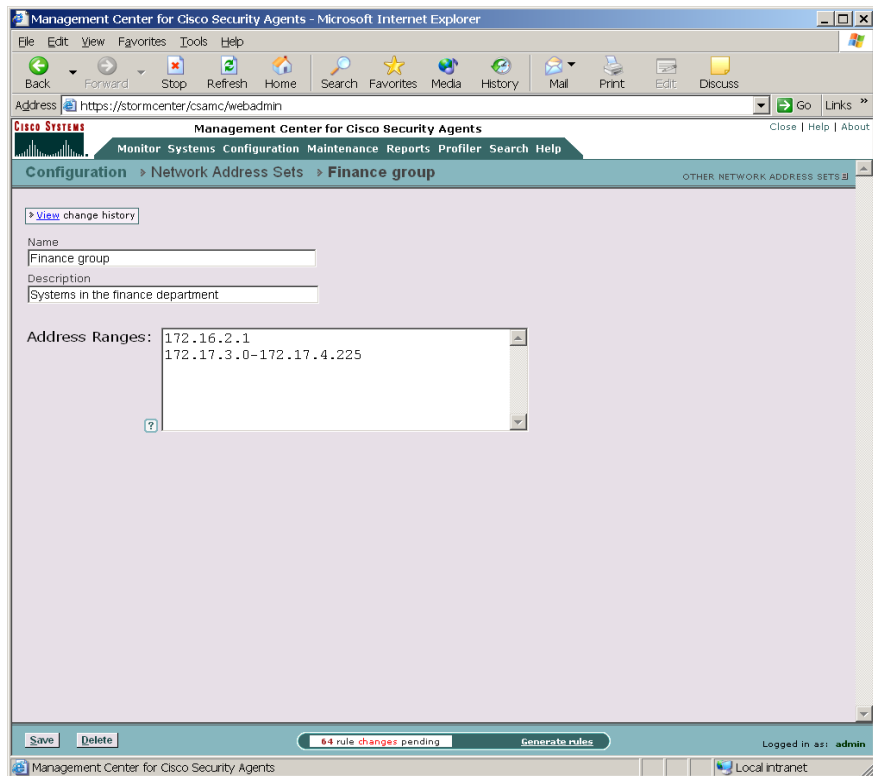
Use **@local** to indicate all local addresses on the agent system. You would want to use this if you are allowing different applications on a single system to talk to each other without opening these applications up to network access.



Caution

On UNIX platforms, IPV6 addresses are not officially supported; however, an IPV6 connection will work as the applied rules dictate if the address in question is covered by the "all" addresses range (0.0.0.0-255.255.255.255 includes IPV6 addresses) or by @local. Local addresses on the agent system (indicated by @local) also include IPV6 addresses.

Figure 7-4 Network Address Set Configuration View



Step 6 When all required information is entered, click the **Save** button to save your address set in the CSA MC database.



Note You can now enter this network address set name by clicking the **Insert Network Address Set** in the network access control rule host addresses field.

Network Services

Configure network services for use in network access control rules to add preconfigured protocol and port number restrictions. You can restrict by initial connection ports, and when applicable, by subsequent client/server connection.

CSA MC ships with several pre-configured network services you can use.

To configure network services, do the following.

-
- Step 1** From the menu bar **Configuration** drop-down list, move the mouse over **Variables**. A cascading menu with further selections appears.
 - Step 2** Select **Network Services** from the cascading menu. Any existing configurations are shown.
 - Step 3** Click the **New** button to create a new network service variable. This takes you to the configuration view ([Figure 7-5](#)).
 - Step 4** In the available edit fields, enter the following information (See the [“Using the Correct Syntax” section on page 2-18](#). You can also click the Quick Help question mark beside each field for syntax information.):

- **Name** This is a unique name for this network service configuration. This name is case insensitive. Generally, it’s a good idea to adopt a naming convention that lets you quickly enter network service variables in network access control rule configuration fields. When using configuration Variables in file access rules, network access rules and application classes, you must enter the variable name preceded by a dollar sign.

For example, if you have a network service variable named FTP Service, you must enter `$FTP Service` into any edit field where you are using this variable. The dollar sign tells CSA MC that this is a variable value.

- **Description** This is a useful line of text that is displayed in the list view and helps you to identify this particular configuration.

- Step 5** **Ports used for initial connection**—Enter a tcp or udp protocol and corresponding port or port range to indicate a restriction according to the system that is initiating the connection

By default, all port fields have a <none> entry indicating no ports. When you click inside each field, the <none> disappears so that you can enter your own port restrictions. For example:

Use the following syntax:

```
TCP/21
UDP/1025-65535
```

- Step 6** **Ports used for subsequent connections initiated by the client**—Enter a tcp or udp protocol and corresponding port to indicate a restriction according to subsequent connections initiated by any client as part of the same session. For example:

```
TCP/1024-65535
```

- Step 7** **Ports used for subsequent connections initiated by the server**— Enter a tcp or udp protocol and corresponding port to indicate a restriction according to subsequent connections initiated by any server back to the client as part of the same session. For example:

```
TCP/1024-65535
```

Some protocols, such as ftp, create additional connections as part of the same session started by the initial connection. The port numbers used for these additional connections are defined in the subsequent port fields. When a network service is used in an allow rule, once an initial connection is established, the subsequent connections will also be allowed, but only to the process that participated in the initial connection.

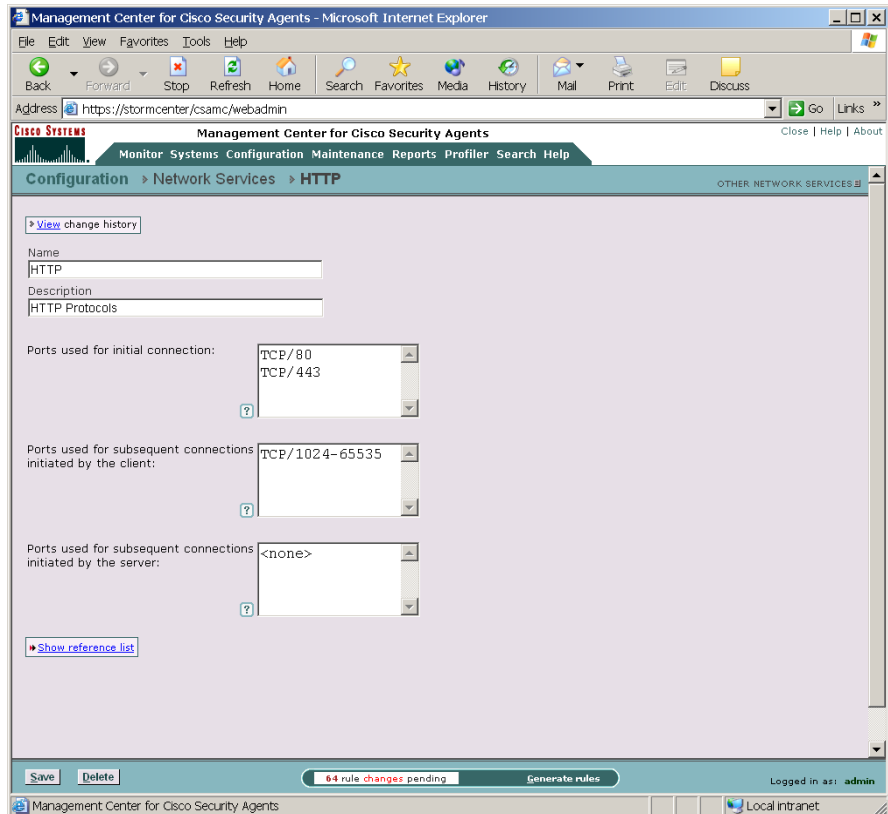
In some cases, an application may want to offer a temporary service port for callback data connections. An ephemeral port is a temporary system-assigned port for this purpose. You can specify an ephemeral port range for a Network service as follows (See the [“Using the Correct Syntax”](#) section on page 2-18 for more details):

```
TCP/ephemeral
UDP/ephemeral
```

Step 8 When all required information is entered, click the **Save** button to save your event set in the CSA MC database.

You can now enter this network service name by clicking the **Insert Network Service** link in the network access control rule network services field.

Figure 7-5 Network Services Configuration View



Registry Sets

A variety of viruses invoke themselves using registry settings. Use the preconfigured registry sets in registry access control rules to prevent viruses from writing to registry values popular with viruses.

This variable is not available for UNIX configurations.



Caution

If you attempt to create your own registry sets to include in a rule, you should note that the ability to restrict registry access is an extremely powerful tool. Critical applications may not function as a result of a misconfigured registry restriction. Therefore, registry values should be as specific as possible. All rules restricting registry access should first be run in **Test Mode** to ensure that no unintended restrictions have been configured.

Registry sets are groupings of registry keys and settings under one common name. This name is then used in rules that allow or deny registry write operations. All the registry restriction parameters that exist under that name are then applied to the rule where the name is used.

To view preconfigured registry sets, do the following.

-
- Step 1** From the menu bar **Configuration** drop-down list, move the mouse over **Variables**. A cascading menu with further selections appears.
 - Step 2** Select **Registry Sets** from the cascading menu. Any existing registry set configurations are shown.
 - Step 3** Click the **New** button to create a new registry set variable. This takes you to the configuration view (see [Figure 7-6](#)).
 - Step 4** Enter the following.
 - **Name** This is a unique name for this registry set.
 - **Description** This is a line of text that is displayed in the list view helping you to identify this particular registry set configuration.

Step 5 Registry keys matching—You *must* enter a value in this field if you are creating a registry set.

The registry key fields (matching and exclusions) must begin with a wildcard or specification of a registry hive. There must be at least one non-wildcarded component in a registry key other than the hive itself.

Hives are one of the following strings:

HKLM—refers to the HKEY_LOCAL_MACHINE

HKCR—refers to HKEY_CLASSES_ROOT

HKCC—refers to HKEY_CURRENT_CONFIG

HKU—refers to HKEY_USERS (HKU* refers to all users)

Table 7-1 Example valid and invalid registry key entries

\MSSQLSERVER	This is a valid entry.
\M*	This is invalid because there is no non-wildcard component. (M* is wildcarded.)
HKLM\SOFTWARE\CSCOpX**	This is a valid entry.
FOO\SOFTWARE\CSCOpX**	This is invalid (FOO is not a hive).

Step 6 but not—Make exceptions to registry keys.

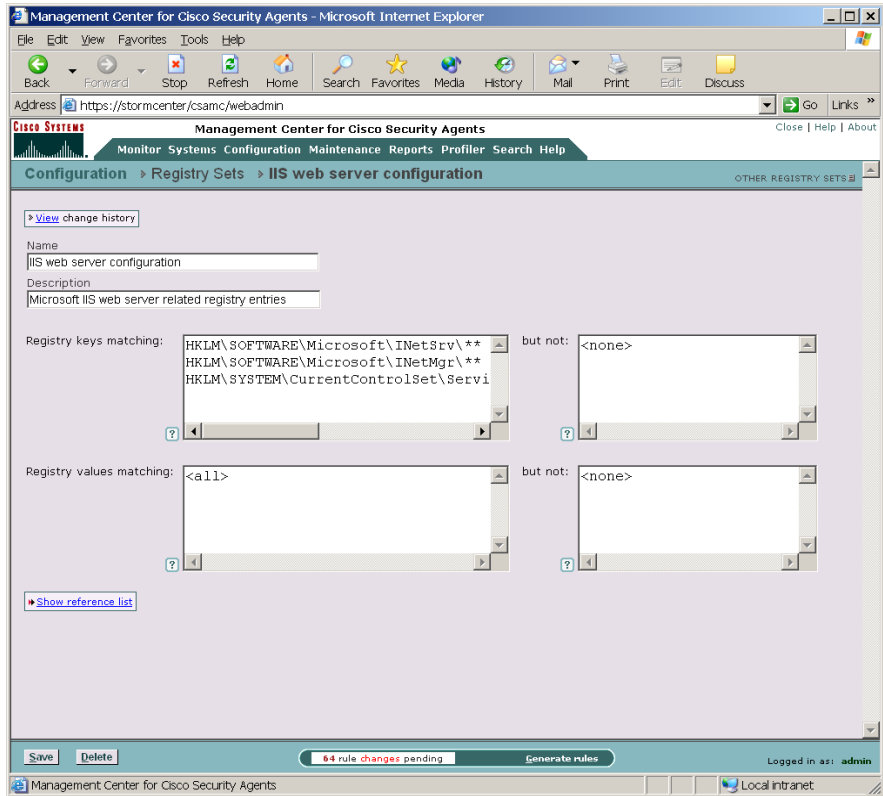
Step 7 Registry values matching— Enter the registry values you are controlling access to.

Step 8 but not—Make exceptions to registry values.

Step 9 When all required information is entered, click the **Save** button to save your registry set in the CSA MC database.

You can enter this registry set name by clicking the **Insert Registry Set** link in the registry access control rule registry entries field.

Figure 7-6 Registry Set Configuration View



Included Registry Sets

CSA MC ships with several pre-configured registry sets you can use in your registry access rules. Some are application specific, others are operating system specific. This section describes a sample of the included operating system specific registry keys.

- Run Keys are used to register programs so that the system will invoke them as a service. Viruses can make use of this key to become persistent.

Protecting this registry value by creating a rule to prevent writing to run keys can prevent the type of virus described above from invoking and propagating itself.



Note

It is important to note that if users have administrator privileges on their systems and are installing software, this type of rule may trigger and prevent that installation. In such cases, using a Query User rule would be most effective. This way, if users are installing software, they themselves can prevent the agent from stopping the installation by answering "Yes" to the query to allow the install. However, if users are not installing software, this type of Query User rule triggering on a system could be treated as a serious issue and the user should answer "No to all" to disallow the action.

- Shell commands are used to tell your system how to open a file based on the file format. This is how the system knows which application to use when opening a particular file.

Viruses can exploit this by having the registry setting invoke the virus along with the application being opened. In this case, the application would open correctly and the virus could silently begin doing harm.

BootExecute tells the system which executables should be run at system startup time.

- Reboot operations tell the system which operations should begin at system startup time. If programs have been uninstalled, the reboot operation also tells the system which files and services should be deleted on the next reboot and startup.

Viruses can exploit this registry setting by marking particular files for copying, overwriting, or deleting on startup. For example, a virus may attempt to delete a system service that could possibly detect the virus itself. By deleting this service at startup, the virus can go undetected.

**Note**

It is important to note that if users have administrator privileges on their systems and are uninstalling software, this type of rule may trigger and prevent the uninstall. In such cases, using a Query User rule would be most effective. This way, if users are uninstalling software, they themselves can prevent the agent from stopping the uninstall by answering "Yes" to the query to allow the action. However, if users are not uninstalling software, this type of Query User rule triggering on a system could be treated as a serious issue and the user should answer "No to all" to disallow the action.

COM Component Sets

Configure COM component sets for use in COM component access control rules. COM objects are groupings of COM Program IDs (PROGID's) and/or COM Class IDs (CLSID's) under one common name. This name is then used in COM component access control rules to allow or deny access to the COM component set name. All COM components that match the entries of a given component set are relevant to the rule in which the set is used.

You can also use pattern matching when creating COM component sets. For example, entering "Word.*" would match "Word.Application" and "Word.Document".

CSA MC ships with several pre-configured COM component sets you can use as well.

**Note**

This is not available for UNIX configurations.

**Note**

CSA MC provides a COM component import utility which installs with each Cisco Security Agent. Running this utility extracts all COM component PROGID's and CLSID's for software running on the system in question. See [page 7-23](#) for instructions.

To configure a COM component set, do the following.

-
- Step 1** From the menu bar **Configuration** drop-down list, move the mouse over **Variables**. A cascading menu with further selections appears.
- Step 2** Select **COM Component Sets** from the cascading menu. Any existing COM component set configurations are shown.
- Step 3** Click the **New** button to create a new COM component set. This takes you to the configuration view (see [Figure 7-7](#)).
- Step 4** In the available edit fields, enter the following information (See the [“Using the Correct Syntax” section on page 2-18](#). You can also click the Quick Help question mark beside each field for syntax information.):
- **Name** This is a unique name for this COM component set. Generally, it’s a good idea to adopt a naming convention that lets you quickly enter COM component set names in a corresponding rule configuration field.
 - **Description** This is a line of text that is displayed in the list view helping you to identify this particular COM component set configuration.
- Step 5** **PROGID’s/CLSID’s matching**—Enter the COM component PROGID’s or CLSID’s here (one per line) to which you want to impose restrictions.

By default, this field has an <all> entry indicating all PROGID’s and CLSID’s. When you click inside this field, the <all> disappears so that you can enter your own restrictions.

When entering PROGID’s, use syntax as shown in the following example:

```
Outlook.Application
```

When entering CLSID’s (uppercase hexadecimals), using the following syntax (You must include the brackets shown here.):

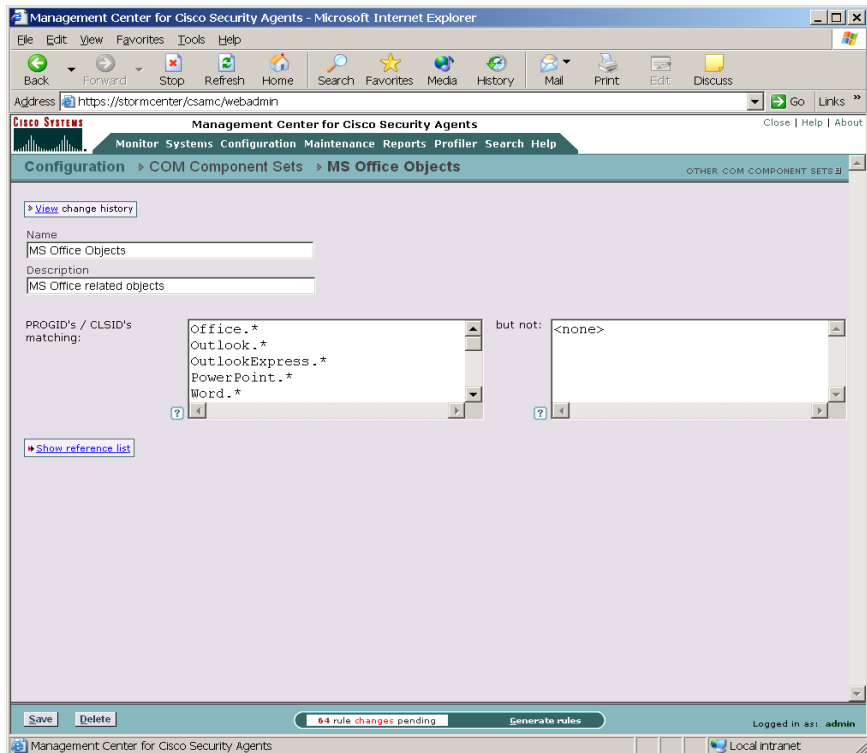
```
{000209FF-0000-0000-C000-000000000046}
```

- Step 6** **but not**—Make exceptions to PROGID’s or CLSID’s you’ve entered in the PROGID’s/CLSID’s matching field.

By default, this field has a <none> entry indicating no exceptions. When you click inside this field, the <none> disappears so that you can enter your own exceptions.

When all required information is entered, click the **Save** button to save your COM component set in the CSA MC database.

Figure 7-7 COM Component Set Configuration View



COM Component Extract Utility

CSA MC provides a COM component extraction utility, called `extract_com`, which installs in the `Cisco\CSAgent\bin` directory with each Cisco Security Agent. Running this utility extracts all COM component PROGID's and CLSID's for software installed on the system in question and places this data into a text file. You can cut and paste these ID's from the text file into your COM component sets and access rules.

See the "Using the COM Extract Utility" section on page 10-9.



Event Logging Alerts

Overview

Events and messages logged by Cisco Security Agents can be viewed from CSA MC. You can also control the type of alert sent out based on the severity level of the logged event, the specific event, and the host that generated the alert. You can configure CSA MC to send email, issue SNMP traps, beep pagers, log to a text file, and execute custom programs.



Note

You can configure the CSA MC Event Log to display events from the Cisco Security Agent system's NT Event Log. See the [“NT Event Log”](#) section on [page 4-67](#).

This section contains the following topics.

- [The Event Log, page 8-2](#)
- [Event Monitor, page 8-5](#)
- [Event Log Management, page 8-6](#)
- [How Logging Works, page 8-8](#)
- [About the Event Management Wizard, page 8-9](#)
- [Event Sets, page 8-20](#)
- [Configuring Alerts, page 8-27](#)

The Event Log

The Event Log view, available from the Monitor category in the menu bar, lets you view system events provided by registered agents according to designated time frames, event severity levels, and the system that generated the event.

The information displayed at the top of the Event Log page (controlled by the settings in the Change Filter window, see next section) tells you the following:

- Filter by eventset: This displays the name of the Event Set, if any, used to filter the event log view.

or Define a filter with the following parameters:

- Event log generation time: This is the date and time of the last event log generation (the last time this page was refreshed).
- Time range: This is the current time range set for the event log filter.
- Severity: This is the current minimum and maximum severity range set for the event log filter.
- Host: This displays which hosts have generated the events viewable in the event log (set as part of the filter).
- Rule ID: Enter the ID number for a rule to search for events generated by that rule.
- Events per page: This is the current value set for the number of events displayed on each page of the event log (set as part of the filter).
- Filter text: Enter a text string here to either include or exclude in your event message search.

Start date and End date

To search events, click the **Change Filter** link to access a pop-up window from which you can enter search criteria such as Start and End Date time frames. You can refer to the following points for entering time frame information, but note that most reasonable time frames are recognized by CSA MC:

- You can select a preconfigured Event Set by which to filter the event log or
- You can specify a relative time using any of the following terms: tomorrow, yesterday, today, now, next, ago, year, month, week, day, hour, minute, and second.

- Enter a specific time using any of the follow time formats: hh:mm:ss. If no meridian (AM or PM) is specified, hh is interpreted on a 24 hour clock (0-23). Note that entering minutes and/or seconds is optional.
- Enter a specific month and day with optional year in the formats: mm/dd?/yy?, monthname dd ?, yy? (Note that the question marks here indicate that the information inside the question marks is optional.) The default year is the current year.

Minimum and Maximum Severity Settings

From the Minimum and Maximum Severity pulldown list, select a severity level and click the View button to see all events within the designated severity levels that have been logged within the time frame you've specified. Select from the following.

- Informational
- Notice
- Warning
- Error
- Alert
- Critical
- Emergency

Host

You can filter the Event Log by host systems. All is the default here. All events generated by systems registered with the server are displayed. You select a specific host from the pulldown list and view only events generated by the selected host.

Events / page

Enter the number of events per page you want to display up to a *maximum of 500 events* per page. The event log displays the most recent number of events based on the value you enter. You can page forward through links to view additional pages matching the query.



Note

You can configure the CSA MC Event Log to display events from the agent system's NT Event Log. See the “NT Event Log” section on page 4-67.

Figure 8-1 Event Log View

Management Center for Cisco Security Agents - Microsoft Internet Explorer

Address: https://stormcenter/csamc/webadmin

Management Center for Cisco Security Agents

Monitor > Event Log

Viewing 1175 - 1126 of 1225 events [change filter](#)

Event log generation time : 1/3/2003 10:44:28 AM
 Severity : Information - Emergency
 Host : All
 Events per page : 50

Filter by eventset: Management events - [edit](#)

Define filter:

Start date: [] e.g.: 24 hours ago
 End date: []
 Minimum Severity: Information
 Maximum Severity: Emergency
 Host: All
 Rule ID: 172
 Events per page: 50
 Filter text: [] Include Exclude

#	Date	Host	Severity	Event
1175	3/25/2003 9:54:56 AM	cheshire	Alert	The p on po Detail
1174	3/25/2003 9:49:07 AM	sneezy	Alert	The p port 1 Detail
1173	3/25/2003 9:44:46 AM	cheshire	Alert	The p port 1 Detail
1172	3/25/2003 9:43:49 AM	kolsch	Alert	The p port 1 Detail
1171	3/25/2003 9:29:03 AM	-	Information	Administrator 'admin' logged in from 10.20.10.115 (S242).
1170	3/25/2003 9:28:57 AM	-	Information	Session for administrator 'admin' marked as expired (S241).
1169	3/25/2003 9:28:06 AM	kolsch	Warning	The process 'E:\Program Files\TextPad 4\TextPad.exe' (as user SNEEZY\vlad) tried to open/create the file 'E:\Program Files\3.2\bin\ht\host_list.htf' and the user was queried. The user responded by choosing 'Yes to All'. Details Rule 313
1168	3/25/2003 8:54:13 AM	sneezy	Alert	The process 'System' (as user NT AUTHORITY\SYSTEM) tried to accept a TCP connection from 10.20.10.104 on port 139 and this was prevented.

No rule changes pending [Generate rules](#) Logged in as: admin

The event log screen (see [Figure 8-1](#)) displays event messages within the time frame and severity level you specify and optionally by a specific host. These event messages explain the event that occurred and they provide a link to the rule that triggered the event. It also provides the exact time the event was recorded and a link to the registered host view for the host that generated the event.

Some Event Log messages contain a **Details** link you can click to view more information on the event that generated the message. (The details contained here can be useful to customer support.) Log messages also contain a **Rule number** link. Clicking a Rule number link takes you to the rule that was triggered when the message in question logged.

Use the **Find Similar** link to locate messages similar to the one from which you accessed the Find Similar box. You can check parameters you wish to search by and select a time frame greater or less than the time the event in question was logged.

Event Monitor

Similar to the Event Log, the Event Monitor, available from the Monitor category in the menu bar, lets you view system events provided by registered agents according to designated severity levels, and the host that generated the event. You can also enter the number of events to be displayed (default value is the last 50 events). Click the **Change** link to access a pop-up window from which you can edit these values and change the event filter. Refer back to [The Event Log, page 8-2](#) for more information on these fields.

Unlike the Event Log page, the Event Monitor page automatically refreshes itself every 15 seconds. The event list is updated with the latest events each time the page refreshes.



Note

The administrator inactivity timeout value is still in effect when you leave the Event Monitor screen displayed on your system. The automatic page refresh does not constitute activity.

The Event Monitor will continue to refresh even after the timeout expires. However, you will not be able to navigate to any other page. This allows you to leave the Event Monitor on screen without worrying about anyone being able to access CSA MC after the session timeout.

Event Log Management

The **Event Log Management** feature, available from the **Monitor** category in the menu bar, lets you create event database management tasks to manage the size of your event log. As your event log grows, specifying parameters for deleting events will help prevent this log from growing too large and from maintaining stale information.

To configure an event log management task, do the following.

-
- Step 1** Move the mouse over **Monitor** in the menu bar and select **Event Log Management** from the drop-down list that appears. See [Figure 8-2](#).
 - Step 2** Click the **New** button to create a new entry. This takes you to the configuration view.
 - Step 3** Enter a **Name** for the event managing task filter.
 - Step 4** Enter a **Description**. This is a useful line of text that is displayed in the list view and helps you to identify this particular configuration.
 - Step 5** In the available fields, enter one or more filtering parameters.

Delete Events:

- **Older than:** Enter a value (default is 30) for which events, once having been in the log for this number of days, are deleted.
- **With a severity level less than or equal to:** From the pulldown list, select a severity level. All events of the selected level, and events of lesser levels, are deleted according to the parameters you set.
- **Related to the following groups:** Select one or more groups here. All events generated by hosts in the selected groups are deleted according to the parameters you set.

To select multiple items in a list box, hold down the **Ctrl** key as you select each item. To unselect a single item, hold down the **Ctrl** key when you click on the item in question. Press the **Shift** key to select multiple successive items.

- Step 6** You can enter a specific **Deletion time** or you can accept the default time of midnight. Note that CSA MC checks once per hour to retrieve this configuration data. Therefore, if CSA MC is not scheduled to do a check for another 40 minutes and you have the deletion time set for 30 minutes from the present time, the deletion will not occur in 30 minutes. It will occur at the time set after the next time CSA MC checks this configuration data. (To delete events immediately, you can use the "Purge events" button in the Event Sets window. See [page 8-20](#) for details.)
- Step 7** Click the **Save** button.

Figure 8-2 Event Log Management Page

The screenshot shows a web browser window titled "Management Center for Cisco Security Agents - Microsoft Internet Explorer". The address bar shows "https://stormcenter/csamc/webadmin". The page content includes a navigation menu with "Monitor Systems Configuration Maintenance Reports Profiler Search Help". The main content area is titled "Monitor > Event Managing Tasks > Untitled_1". It contains a form with the following fields:

- Name: Event maintenance
- Description: Delete events according to these parameters
- Delete events section:
 - Older than: 30 day(s)
 - AND
 - With severity level less than or equal to: information
 - AND
 - Related to the following groups: Apache Servers - Dedicated, Default Servers, IDS Mode Servers, Default Servers, Mission Critical Systems
 - Deletion time (hh:mm:ss): 12:00:00
 - Last deletion recorded on: 2003-03-09 23:52:12

At the bottom of the form, there are "Save" and "Delete" buttons, a status bar indicating "No rule changes pending", and a "Generate rules" button. The user is logged in as "admin".

How Logging Works

The CSA MC Event Log does not contain every occurrence of an event from a system. Duplicate events are not logged for an hour after the first occurrence. See the [“Caching Query Responses” section on page 4-12](#).

**Caution**

In some cases, when an event is logging continuously, the server will suppress this logging for a time (10 minutes, unless verbose logging is enabled). Before it does this, a log message informing you of this suppression appears in the event log.

The following information is logged for each rule type.

- File access control logging: Process path and file names are logged.
- Network access control logging: Process path and network address are logged.

**Note**

No network access control rule denial events are logged for any TCP or UDP port resulting from multicast packet signals.

- Registry access control logging: Process path and registry key are logged.
- COM component access control logging: Process path and COM component PROGID/CLSID are logged.

A duplicate event is defined as follows:

- For file access controls, the name of the application and the file being accessed are the same.
- For network access controls, the name of the application, the remote address, and the network service port are the same.
- For registry access controls, the name of the application and the registry key name and value name are the same.
- For COM component access controls, the name of the application and the COM component PROGID or CLSID are the same.

Verbose Logging

Enable Verbose Logging Mode in the Group configuration view to change the event log timer to log *all* recurring events rather than only logging recurring events once every hour. Verbose logging applies to all policies that are attached to the group that have logging turned on.

For normal operations, you would not want to enable Verbose logging. Verbose logging is useful for troubleshooting and for analyzing how applications work with rule sets, i.e. related processes and subprocesses. In the latter case, using Verbose logging with Test Mode can be very useful for monitoring how a rule set would work before deploying it.

Logging and Query User Rules

When a user responds to a Query User box (by pressing Yes, Yes to all, No, No to all), the agent remembers the response and caches it for an hour. This way, if the same rule is triggered again within that hour, the action is allowed or denied based on what the user answered previously, with no pop-up query box appearing again. When the user responds to a triggered Query User pop-up box, the system action that triggered the pop-up, as well as the user's response, are logged in the CSA MC event log. With Verbose logging turned on, all subsequent automatic allows or denies are logged as well. Otherwise, the one hour logging timer prevents agents from logging the automatic allowed or denied system action if it occurs again within the hour.

About the Event Management Wizard

Use the Event Management Wizard to accomplish the following:

- To change the action of a rule that triggered a specific event. If an action is being denied on end user systems and you want to allow this action, you can automatically generate an "exception" allow rule which takes the application class and resource information in the event and creates an allow rule to counteract the rule that caused the deny.
- To create an exception rule that stops a specific event from logging. The Wizard makes use of the **Take precedence over similar <action type> rules** feature to manipulate rule precedence and prevent logging of an event.

- To create a Profiler analysis job for the application that caused the event. The Event Management Wizard is available for events triggered by Deny rules and Query User rules of the following types: (See event #80464 in [Figure 8-3](#)):

You launch the Wizard from a **Wizard** link which appears with certain event log message types as follows.

- Application control
- Buffer overflow
- COM component access control
- File access control
- Network access control
- Registry access control
- Rootkit/kernel protection
- Trojan detection

Figure 8-3 Event Management Wizard Link

Management Center for Cisco Security Agents - Microsoft Internet Explorer

Address: https://stormcenter/cssamc/webadmin

Management Center for Cisco Security Agents

Monitor Systems Configuration Maintenance Reports Profiler Search Help

Monitor > Event Log

Viewing 80473 - 80424 of 80473 events [change filter](#)

Event log generation time : 3/10/2003 9:25:25 AM
 Severity : Information - Emergency
 Host : All
 Policy : All
 Events per page : 50

[Latest](#) [Earliest](#)

#	Date	Host	Severity	Event
80473	3/10/2003 9:10:00 AM	-	Information	Administrator 'JSmith88' logged in from 10.20.10.34 (S1456).
80472	3/10/2003 8:00:39 AM	-	Information	'Log StormWatch Database Backup' completed successfully.
80471	3/10/2003 12:00:39 AM	-	Information	'Full StormWatch Database Backup' completed successfully.
80470	3/10/2003 12:00:02 AM	-	Warning	Product license for stormwatch will expire in 18 days.
80469	3/9/2003 4:00:23 PM	-	Information	'Log StormWatch Database Backup' completed successfully.
80468	3/9/2003 8:00:23 AM	-	Information	'Log StormWatch Database Backup' completed successfully.
80467	3/9/2003 12:00:22 AM	-	Information	'Differential StormWatch Database Backup' completed successfully.
80466	3/9/2003 12:00:01 AM	-	Warning	Product license for stormwatch will expire in 19 days.
80465	3/8/2003 4:00:20 PM	-	Information	'Log StormWatch Database Backup' completed successfully.
80464	3/8/2003 8:00:19 AM	-	Warning	The process 'D:\WINNT\Explorer.EXE' (as user System\Administrator) tried to open/write the file 'F:\Program Files\Office\EXCEL.EXE' and the user was queried by rule 67. The user responded by choosing 'No'. Details Rule 31 Wizard Find Similar

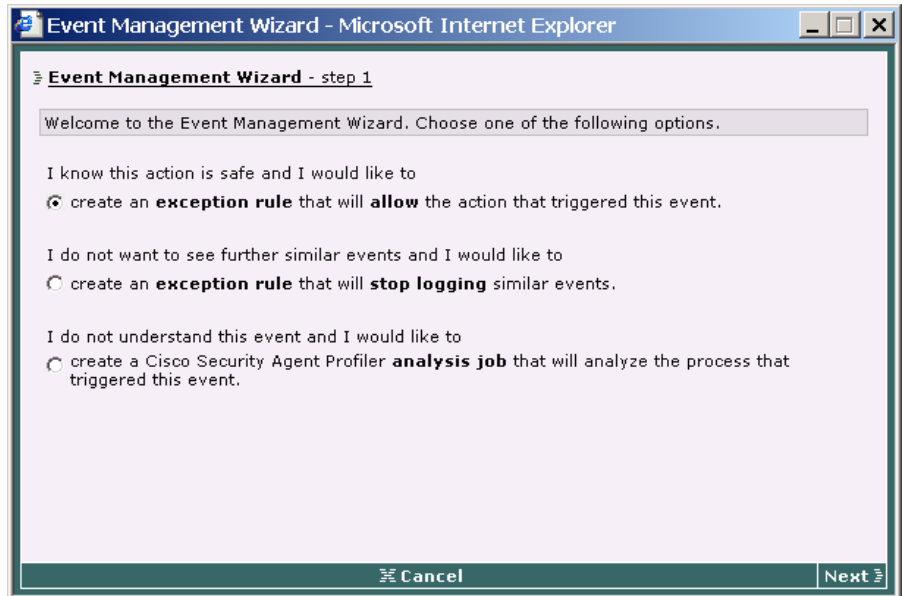
No rule changes pending [Generate rules](#) Logged in as: admin

Creating an Exception Rule

When you click the **Wizard** link from the Event Log page, you can choose to create an exception rule, an exception logging rule, or to configure an analysis job (see [Figure 8-4](#)). If you select to create an exception rule, when you click **Next** you are shown a summary of the rule that triggered the event. You click **Next** to continue to create the rule. The exception rule is an Allow rule that will be added to an existing policy or to a new policy and it will take precedence over the Deny or Query User rule that caused the event.

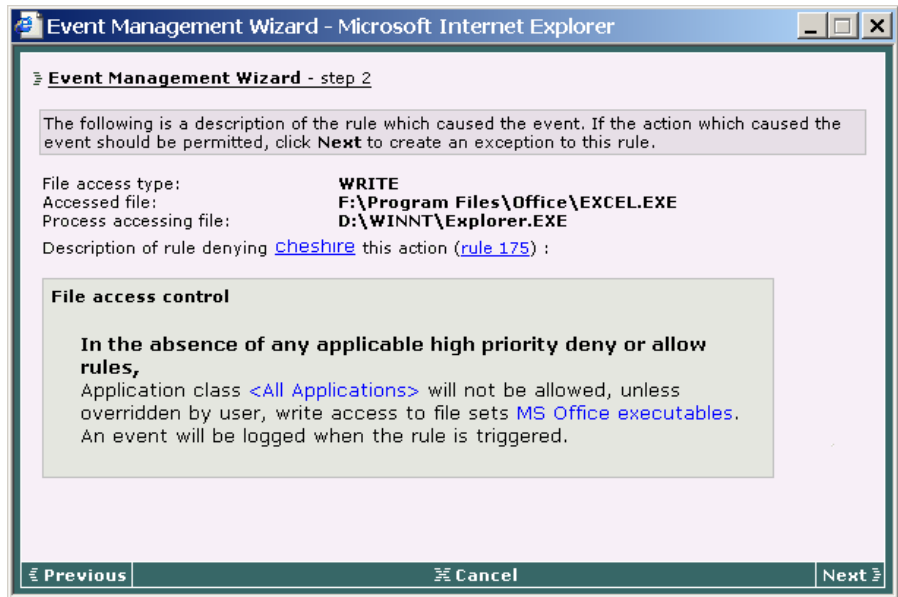
The wizard then takes you through a step by step procedure for affecting the changes you wish to make.

Figure 8-4 Exception Wizard Step 1



The second wizard page provides a summary of the rule for which you are creating an exception.

Figure 8-5 Exception Wizard Step 2

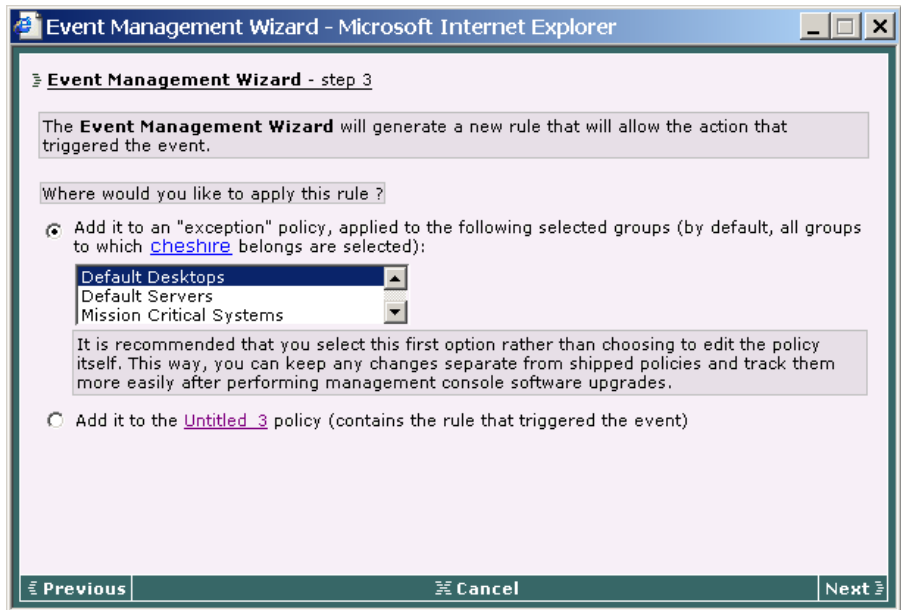


In the third exception rule wizard screen (see [Figure 8-6](#)), you are given two choices as to where you would like to apply this new rule. You can create a new policy (an "exception policy") which would contain the new exception rule. (This is the default and recommended choice.)

This new policy would be attached to the group(s) containing the host from which the event was received. If you choose to create this exception policy, all subsequent exception rules you create through the wizard will be added to the same exception policy if the group it is to be applied to is also the same. Therefore, a group could only have one exception policy, but contain any number of exception allow rules created through the wizard.

Your second choice is to add this new exception rule directly to the policy which contains the rule that caused the event. This would be a change to the policy itself. All groups that have this policy would receive the exception rule.

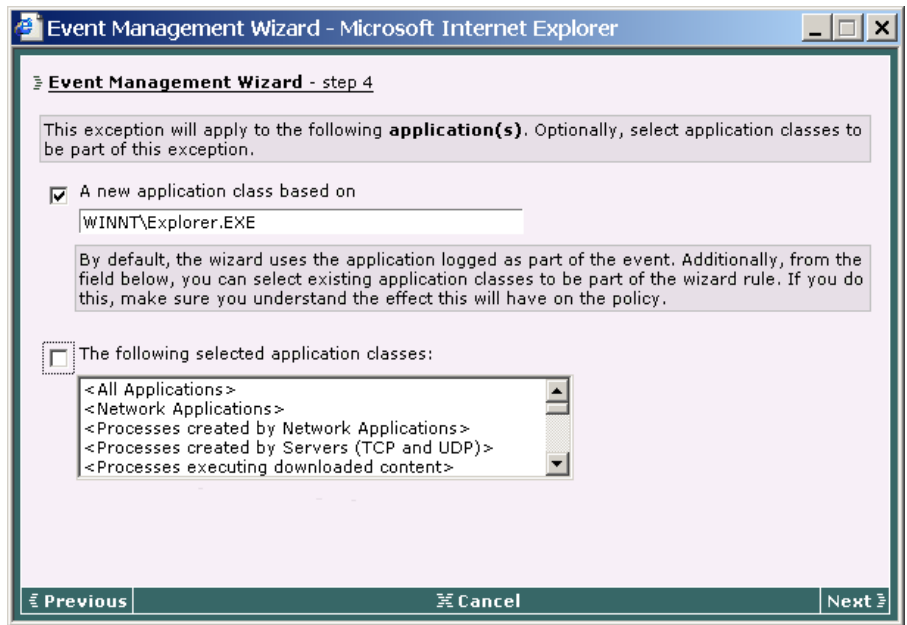
Figure 8-6 Exception Wizard Step 3



As you continue through the wizard, you can simply click Next and accept the defaults which create an allow rule for the exact application and resource named in the event. You can also select to include more groups to receive the exception and edit the process path of the application and/or include more application classes in the rule (see [Figure 8-7](#)).

When the wizard completes, it takes you to the new rule as it appears in CSA MC.

Figure 8-7 Exception Wizard Step 4

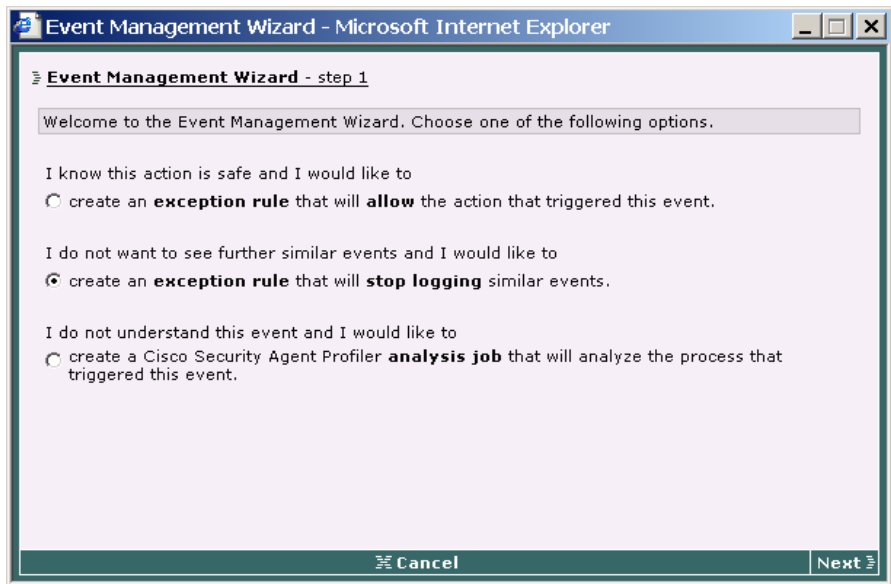


Creating a Logging Exception Rule

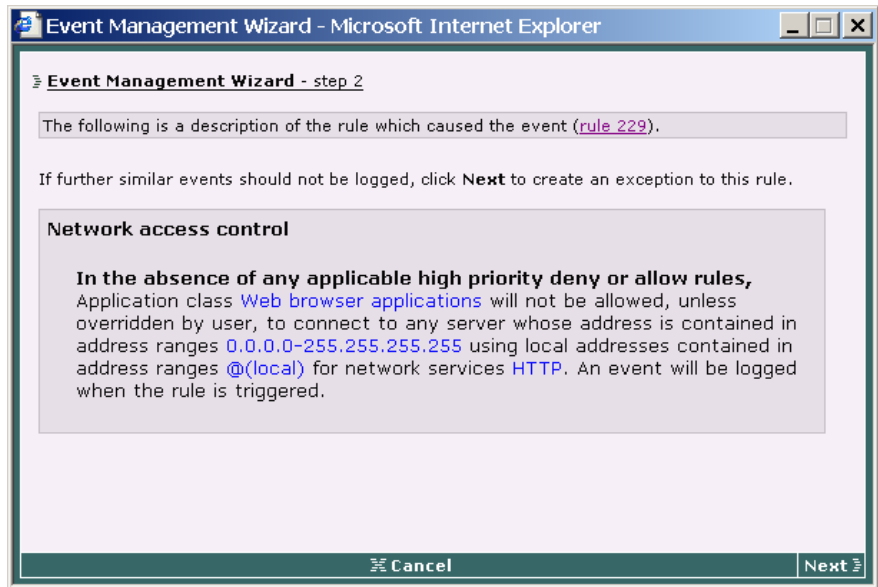
The Wizard makes use of the **Take precedence over similar <action type> rules** feature available in some rule types to manipulate rule precedence and prevent the logging of an event. The following rule types make use of precedence manipulation: File access control, Network access control, Registry access control, COM component access control, and Application control.

See the [“Writing Rules: Manipulating Precedence”](#) section on page 4-8 for more information on the manipulating precedence feature.

If you select to create an exception logging rule (see [Figure 8-8](#)), when you click **Next** you are shown a summary of the rule that triggered the event. You click **Next** to continue to create the exception logging rule.

Figure 8-8 Exception Logging Wizard Step 1

The exception logging rule is a rule that is added to an existing policy or to a new policy. This rule is an exact copy of the rule that triggered the event. The one difference is that the rule created by the wizard has the **Take precedence over similar <action type> rules** checkbox selected and the **Log** checkbox is unselected. This causes the rule created by the wizard to remain in effect, in the correct precedence within the policy, but not log an event when triggered.

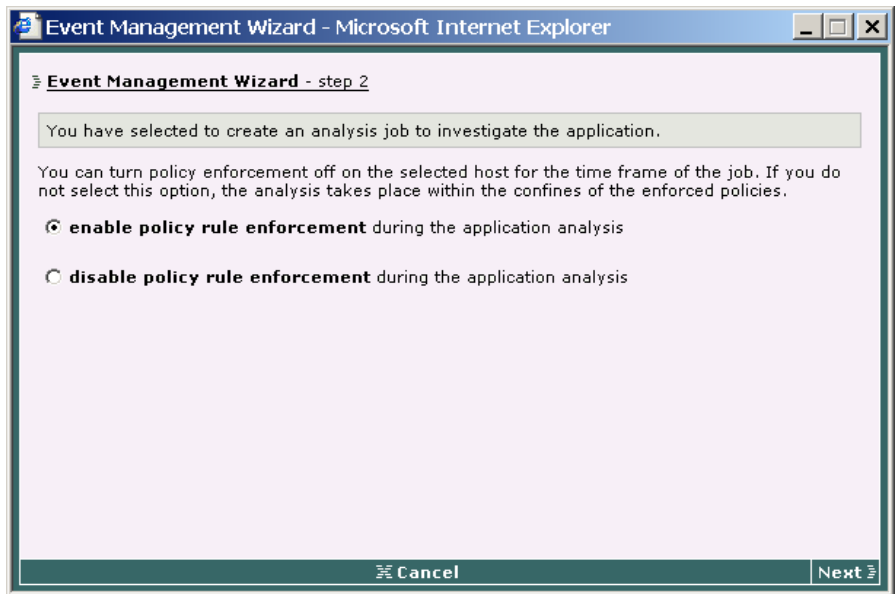
Figure 8-9 Exception Logging Wizard Step 2

Creating an Analysis Job

When you click the **Wizard** link from the Event Log page, you can choose to configure a Cisco Security Agent Profiler analysis job to investigate the application that triggered the event (see [Figure 8-4](#)).

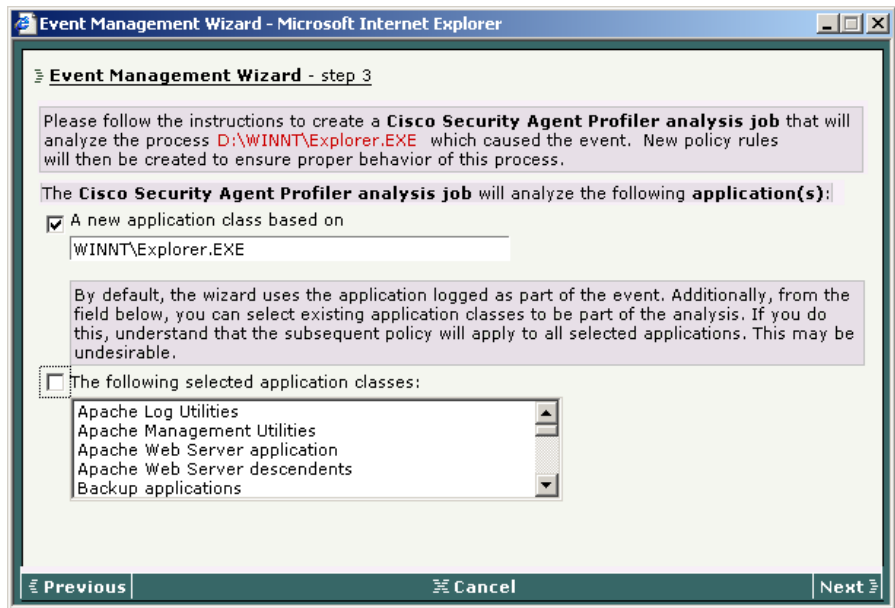
If you select to create an analysis job, optionally you can choose to **Disable policy rule enforcement** for the time frame of the job. Otherwise, the analysis takes place only within the confines of enforced policies. In that case, some events may be denied by rules during the analysis and therefore the analysis may not be complete.

If you select the Disable policy rule enforcement checkbox, when the logging agent receives an analysis job, any policies relevant to the application being analyzed are disabled on the selected host until the job is completed. You should understand that if the application being analyzed is untrusted or potentially a virus, you will allow to run unimpeded during the analysis if you disable policy rule enforcement.

Figure 8-10 Analysis Job Wizard Step 2

If you decide that the application is not dangerous and it can run without any policy restrictions, you can begin to configure the analysis job.

Figure 8-11 Analysis Job Wizard Step 3



The next analysis job wizard page (see [Figure 8-11](#)) displays the application that triggered the event. This is the application the analysis job will investigate. Optionally, you can select other application classes to be analyzed. But in that case, the policy created by Profiler would apply equally to all applications included in the analysis. For example, if the application class you are analyzing contains both Microsoft Word and Microsoft Outlook, the policy created by Profiler would be a combination of the resources required by both applications.

Continuing to click the **Next** button through the analysis job wizard configures the analysis job with chosen defaults for analysis workstation and job time frame. You can choose to edit these defaults or to accept them by making no changes.

When the wizard completes, it takes you to the new analysis job as it appears in CSA MC. You can edit it at this time or you can deploy the job by doing the following:

- **Generate rule programs** to distribute the analysis job to the host.
- Wait for the logging process to stop or click the **Stop logging** button to force the stop.

- Click the **Start analysis** button to start the analysis of the logged data.
- Optionally, use the **Import** button to import the policy, examine it and, if appropriate, deploy it to hosts.

Event Sets

Configure event sets for use in alerts, reports, and event logs. When configuring alerts, event sets cause CSA MC to trigger alerts based on specified events. Once configured, these event set configurations become available in corresponding alert selection fields.

**Note**

CSA MC ships with several preconfigured event sets you can use. If the included event sets do not suit your needs, use the instructions in the following pages to configure new event sets or to edit existing ones.

When creating your event sets, it's a good idea to adopt a naming convention that lets you quickly recognize event sets in your Alert configuration view.

**Note**

To learn more about how event sets are used for generating reports, see [Chapter 9, “Generating Reports”](#).

To configure event sets, do the following.

-
- Step 1** Move the mouse over **Monitor** in the menu bar of CSA MC. Select **Event Sets** from the drop-down list that appears. All existing event set configurations are shown.
 - Step 2** Click the **New** button to create a new event set. This takes you to the configuration view.

- Step 3** In the available edit fields, enter the following information (see [Figure 8-12](#)):
- **Name:** This is a unique name for this event set. Generally, it's a good idea to adopt a naming convention that lets you quickly recognize Event Sets in Alert configuration fields.
 - **Description:** This is a line of text that is displayed in the list view and helps you to identify this particular Event Set configuration in the event set list view.
- Under the **Event Specification** section, enter optional filtering parameters.



Note To select multiple items in a list box, hold down the Ctrl key as you select each item. To unselect a single item, hold down the Ctrl key when you click on the item in question. Press the Shift key to select multiple successive items.

- Step 4** Select **Filter by event** specifications.
- Leave the **Include all event types** radio button selected to have events of all types included or select the **Include only the following selected event types** radio button. If you select the second radio button, then you must also select specific event log messages to filter by. These messages represent the spectrum of generated events that appear in the Event Log view.
- Step 5** Select **Filter by severity** specifications.
- Leave the **Include all severity levels** radio button selected to have events of all severity levels included or select the **Include only the following selected severity levels** radio button. If you select the second radio button, then you must also select the severity level(s) that will trigger an alert for this event set. Available levels are: Information, Notice, Warning, Error, Alert, Critical, Emergency.
- Step 6** Select **Filter by group** specifications.
- Leave the **Include all hosts** radio button selected to have events generated by all hosts included or select the **Include only hosts in the following selected groups** radio button. If you select the second radio button, then you must select the group(s) that trigger an alert for this event set. Any groups selected here that log the event in question will trigger an alert.

Step 7 Select **Filter by policy** specifications.

Leave the **Include all policy rules** radio button selected to have events generated by all rules included or select the **Include only rules in the following selected policies** radio button. If you select the second radio button, then you must select the policy(ies) that trigger an alert for this event set. Any policies selected here that log the event in question will trigger an alert.

Step 8 Select **Filter by time** specifications.

Leave the **Include all timestamps** radio button selected to have events generated at all times included or select the **Include only these timestamps** radio button. If you select the second radio button, then you can create a custom time here or select from available times, Today, Last 24 hours, Last 7 days, and Last 30 days to trigger an alert when an event occurs with the specified time range.

You can also enter **Custom start** and **Custom end** times in the following manner:

- Specify a relative time using any of the following terms: tomorrow, yesterday, today, now, last, this, next, ago, year, month, week, day, hour, minute, and second.
- Enter a specific time using any of the following time formats: hh:mm:ss. If no meridian (AM or PM) is specified, hh is interpreted on a 24 hour clock (0-23). Note that entering minutes and/or seconds is optional.
- Enter a specific month and day with optional year in the formats: mm/dd?/yy?, monthname dd ?, yy? (Note that the question marks here indicate that the information inside the question marks is optional.) The default year is the current year.



Note When you select multiple categories to filter by, all selections have to match.

Step 9 When all required information is entered, click the **Save** button to enter and save your event set in the CSA MC database.

In the Event Sets configuration page, the CSA MC frame at the bottom of the page provides a **View** button and a **Purge events** button.

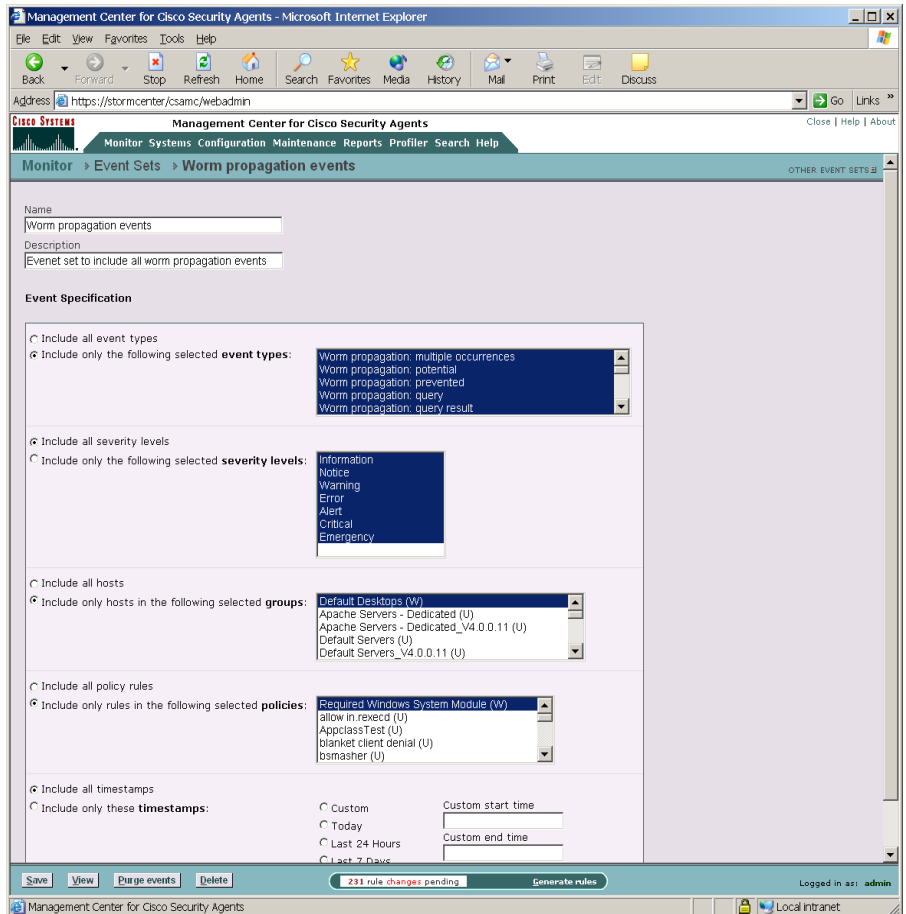
- When you click the **View** button, all events that match the configured event set are displayed.



Caution

When you click the **Purge events** button, all events that match the configured event set are deleted from the event log. If you make changes to an existing Event Set and click the Purge events button without saving those changes, all edits are saved and events are purged.

Figure 8-12 Event Set Configuration View



Third Party Access to Events

To access events in the database for exporting to a different format (or for your own reports), connect to the database using ODBC DSN "CSCODSN."

You can access events through the database view EventListView. (This is a SQL server view.) The columns defined in this view are as follows:

**Note**

SNMP and Log file alert types can be used by third party event management applications. See [Table 8-2](#) for more details on those alert fields. (Note that the fields in the SNMP and Log file alerts are the same as those described in [Table 8-1](#).)

Table 8-1 EventList View Fields

Field	Description
EventId	An ID uniquely identifying the event. Increasing, in order of event arrival at CSA MC.
EventTime	The time at which the event occurred, using the clock of the host that generated the event.
HostId	An integer uniquely identifying the host that generated the event. This is NULL for events generated by CSA MC.
HostName	A non-unique string name for the host that generated the event.
HostOSType	The OS type for the host that generated the event, 'W' for Windows, 'U' for UNIX
CurrentHostIPAddress	The most recently recorded IP address for the host that generated the event.
SeverityCode	An integer, as follows in increasing severity -- Information (1), Notice (2), Warning (3), Error (4), Alert (5), Critical (6), Emergency (7)
SeverityName	The string representation of SeverityCode.
ProcessName	When applicable, the full path of the process that generated the event.
FileName	When applicable, the name (not path) of the relevant file from a file event.
SourceIPAddress	When applicable, the source IP address of a network event.
DestinationIPAddress	When applicable, the destination IP address of a network event.
SourcePort	When applicable, the port used by the source of a network event.

Table 8-1 *EventList View Fields (continued)*

Field	Description
DestinationPort	When applicable, the port used by the destination of a network event.
RuleId	An integer uniquely identifying the rule that caused the event.
RuleDescription	The user-specified string description for the rule that caused the event.
PolicyId	An integer uniquely identifying the policy which contains the rule that caused the event.
PolicyName	The string name of the policy which contains the rule that caused the event.
EventCode	An integer which uniquely defines the event code.
EventCodeTag	A short string representing the event code.
EventText	The complete formatted text of the event. (A Test Mode event is preceded by the string "TESTMODE".)
EventType	A string representing the type of the rule that caused the event, as discussed in Chapters 4 and 5. This field can be used as a broad-level categorization of CSA MC events. Possible values are as follows: File access control, Network access control, Network shield, Registry access control, Network worm protection, Trojan detection, Sniffer and protocol detection, File version control, COM component access control, File monitor, Service restart, NT Event log, Application control, Agent service control, Data access control, Connection rate limit, Profiler, Kernel protection, Network interface control, Rootkit / kernel protection, Buffer overflow, Syslog control, Resource access control, Downloaded content, Global virus scan, Global event log, Global network scan, Global email worm, Self-protection, Administrative.

Table 8-1 *EventList View Fields (continued)*

Field	Description
ButtonCode	The bottom 16 bits of this field represent the button that was pressed, with short integer values as follows, Yes (1), Yes to All (2), No(3), No to All (4), Terminate Process (5), OK (6). The upper 16 bits of this field represent whether the button was selected by default. A zero value indicates that the user actually pressed the button, while a non-zero value indicates that the default was chosen, e.g. because the query timed out.
Username	The name of the logged-in user at the time of the event.

Configuring Alerts

You can configure CSA MC to send various types of alerts to specified recipients when a policy triggers an event. Available alert types include: Email, Pager, SNMP, Log to file, and a Custom program that you provide.

Each alert type requires you to enter specific information.

To configure CSA MC to issue alerts when specified system events occur, do the following.

-
- Step 1** Move the mouse over **Monitor** in the menu bar and select **Alerts** from the drop-down list that appears. The list of Alerts (if any) appears.
 - Step 2** Click the **New** button to create a new alert. This takes you to the configuration view.
 - Step 3** In the Alert configuration view (see [Figure 8-13](#)), enter a **Name** and a useful **Description**. This information is displayed in the list view and helps you to identify this particular alert.

- Step 4** From the **Send alerts for the following event set** list box, select the event set(s) you want to trigger the alert you're creating. Configuring Event Sets provides flexibility in selecting the events for which you want to be alerted.



Note The "time" filter in an event set is ignored for alerts. Alerts are generated as events are logged.

To select multiple items in a list box, hold down the **Ctrl** key as you select each item. To unselect a single item, hold down the Ctrl key when you click on the item in question. Press the **Shift** key to select multiple successive items.

If the available options here do not meet your needs, you can configure event set variables which become selectable in this field.

- Step 5** In the available alert configuration fields, enter data for *one or more* of the following alert types: Email, Pager, SNMP, Log, Custom (for alert configuration information, refer to [Table 8-2](#)).

For each alert type you want to send, select the corresponding **checkbox** and enter the required alert-specific information.



Note Although you can enter data into all available alert edit fields, if you do not check the corresponding checkbox, the alert in question is not enabled; however, the information you've entered is stored in the database. You can enable the alert type at a later time.

- Step 6** When your information is entered, click the **Save** button to save your new alert(s).



Note Use the **Clear Pending Alerts** button to clear all alerts that have been triggered by events but not yet sent. You might want to do this if several events are occurring simultaneously or continuously, you have already disabled the alert, and you have no further need for the continual notifications that are pending.

Figure 8-13 Alert Configuration View

The screenshot shows the Management Center for Cisco Security Agents web interface in Microsoft Internet Explorer. The browser address bar shows the URL: `https://stormcenter/csamc/webadmin`. The page title is "Management Center for Cisco Security Agents". The navigation menu includes "Monitor", "Systems", "Configuration", "Maintenance", "Reports", "Profiler", "Search", and "Help". The current view is "Alerts" > "Worm propagation alert".

The configuration form includes the following fields:

- Name:** Worm propagation alert
- Description:** Send an alert when a worm propagation event occurs
- Send Alerts:** A dropdown menu showing "Worm propagation events" selected. Other options include "Application and COM invocation", "Configuration errors", "Events from critical systems", and "Events from critical systems_V4.0.0.11".
- Email:** Checked. Recipient(s) email address(es): <jsmith88@example.com>. Sender address to use: <stormcenter@example.com>. Address of mail server: 209.165.201.0.
- Pager:** Unchecked. Fields for Telephone, PIN, Modem Init String, Modem Dial String, Max. characters/block, Port, and Baud.
- SNMP:** Unchecked. Fields for Community name and Manager IP address.
- Log:** Unchecked. Field for Log file.
- Custom:** Unchecked. Field for Custom program.

At the bottom of the form, there are "Save" and "Delete" buttons, a status bar indicating "48 rule changes pending", and a "Generate rules" button. The user is logged in as "admin".

Table 8-2 Alert Type Descriptions

Alert Type	Information	Description
Email	Recipient	Enter the email address of the mail recipient. Using brackets is optional. CSA MC will automatically enter them if you do not. You can enter multiple addresses separated by semi-colons: <dpaul@example.com>
	Sender address to use	Enter the mail sender in brackets. Some mail servers require this to be specified: <jsmith@example.com >
	Address of SMTP server	Enter the IP address or DNS name of the SMTP server.

Table 8-2 Alert Type Descriptions (continued)

Alert Type	Information	Description
Pager	Telephone	Enter the recipient's pager number: 5550112 Pager must support TAPI. Refer to your paging service provider for further information. (Note that if your company requires a number be dialed, such as 9, to access an outside line, you must enter that number here as part of the telephone number.)
	PIN number	PIN number: 10799 (This is a string that identifies the pager. Obtain this from the service provider.)
	Modem Init String	Modem Init String: ATQ0VI~ (You must check your modem instruction manual to locate this value. In some cases, it can be left blank. Not all modems require this value or a modem dial string.)
	Max. characters/block	Max. characters/block: 80 (Some paging services only allow a certain number of characters in a single message block. Check with your paging service. You can leave this field blank to use a default value of 80.)
	Modem Dial String	Modem Dial String: TO=20, ATDT (See comment above.)
	Port	Port: COM1 (This is the serial port the modem is attached to.)
	Baud	Baud rate: 2400 (This is the baud rate supported by the paging service. If you leave this field blank, CSA MC enters a default value of 2400 when you save the alert.)

Table 8-2 Alert Type Descriptions (continued)

Alert Type	Information	Description
SNMP	<p>Community Name</p> <p>Manager IP Address</p>	<p>Enter the community name. This is a text string agreed upon by the SNMP manager: <code>public</code></p> <p>Enter IP address of the system where the SNMP trap should be sent. Optionally, you can put a colon and a port number ("<code>:<port number></code>") after the IP address if you are using a non-standard port. (Standard port is 162.)</p> <p>Refer to the <code>CSAMC-SNMPv2.mib</code> document in the <code>CSCOPx\CSAMC\doc</code> directory for SNMP-MIB definitions for Cisco specific objects.</p> <p>Also see Third Party Access to Events, page 8-24 for third party event management details.</p>
Log	Log file name (using full path)	<p>Enter a name for the flat logging file that events will be written to.</p> <p><code>c:\alerts\logfile.txt</code></p> <p>This file can then be used by third party event management applications. See Third Party Access to Events, page 8-24 for details.</p>

Table 8-2 Alert Type Descriptions (continued)

Alert Type	Information	Description
Custom	Custom Program	<p>Enter a custom alert program name here.</p> <p>The server calls the program as it appears in this field. You must enter the full pathname so that CSA MC can locate the program.</p> <p>Your custom program must be an executable file. <code>c:\Program Files\CSCOpX\program.exe</code></p> <p>The program passes the event message in a file whose name is passed to the program as its first parameter. Alternately, the program can also read the event message from its standard input. The file containing the event is automatically deleted when the program exits or closes its standard input.</p> <p>FEATURE NOTES:</p> <ul style="list-style-type: none"> * The custom program must exist on the same system as CSA MC in the CSCOpX directory or subdirectory. If it is located elsewhere, the VMS policy will not allow it to run. * Custom programs cannot require any user input. * If a custom program is triggered and fails for some reason, it could take several minutes before the program closes itself and attempts to launch again. (If you are testing custom program alerts, one way to tell if the program has launched and is running, is to watch for it in the Task Manager.)
Named Pipe	Named Pipe	<p>A named pipe is a form of internal communication. This alert type allows the integration of third party software for the purpose of receiving alerts over Windows named pipes. Consult your third party documentation for further configuration details.</p> <p>Note that this feature is for use with third party vendors that support alerts over Windows named pipes.</p>

Generate an Alert Log File for Third Party Applications

Using the **Log** checkbox and the **Log file** edit field in the Alerts configuration page (see [Figure 8-13](#)), you can have CSA MC generate a flat logging file to which events are written. Third party event management applications can then parse the information found in this file.

To generate this file, select the Log checkbox and enter the Log file name, using the full path that you want to write event data to. For example, enter `c:\alerts\logfile.txt`

Event data is written to this file as follows:

```
EventId,EventTime,HostId,HostName,  
CurrentHostIPAddress,HostOSType,Severity,EventType,  
EventText,EventCodeTag,FileName,ProcessName,  
SourceIPAddress,DestinationIPAddress,SourcePort,  
DestinationPort,RuleId,RuleDescription,RulePriority,  
PolicyId,PolicyName,ButtonCode,UserName
```

Entry fields are separated by a delimiter of a comma. Event entries themselves are separated by a carriage return/line feed (ASCII Hex 0D 0A).

Once a log file exceeds 1 MB, it is closed and its name is suffixed with a time stamp. A new file, using the same file name entered in the CSA MC Alerts Log file field, is then created. Events continue to be written to this new file until it reaches 1 MB.



Generating Reports

Overview

You can configure the Cisco Security Agent to log an event each time a system action triggers a rule.

You can use the event logging data received from agents to generate reports that indicate overall network health. Using these reports, you can monitor how your current rule sets are working and adjust them, if necessary.

You can also generate reports related to configuration information.

This section contains the following topics.

- [Types of Reports, page 9-2](#)
- [Viewing Reports, page 9-2](#)
- [Generating Reports, page 9-2](#)
- [About the ActiveX Crystal Report Viewer, page 9-8](#)

Types of Reports

CSA MC lets you generate reports using various criteria. For example, you can create reports based on event severity level, on the group that generated the event, and on the individual host systems producing events. You can sort by other parameters such as time frame, host, and event code that you configure separately.

Viewing Reports

When you generate your reports, you're given the option of selecting the type of viewer through which to display the report. From the Viewer type pulldown menu, you can select the following.

- **ActiveX:** The report viewer for ActiveX uses an ActiveX control that can be placed inside an HTML page and viewed through any browser that supports ActiveX. (Supported by Internet Explorer 3.02 and higher. Not supported by Netscape.) See [About the ActiveX Crystal Report Viewer, page 9-8](#) for more information.
- **HTML Frame:** This view is selected by default if you do not select a viewer type. Using this viewer, you can display reports in HTML using frames to illustrate category data in a left frame. (Supported by Internet Explorer 3.02 and higher and Netscape Navigator 4.7 and higher.)

When you print reports, the formatting will vary depending on which view type you have selected and the printer settings on the printer you're using.

Generating Reports

You generate reports by selecting various sorting options in the CSA MC report configuration views. When you are finished selecting sorting parameters, you can generate your report. The report opens in a new browser window.

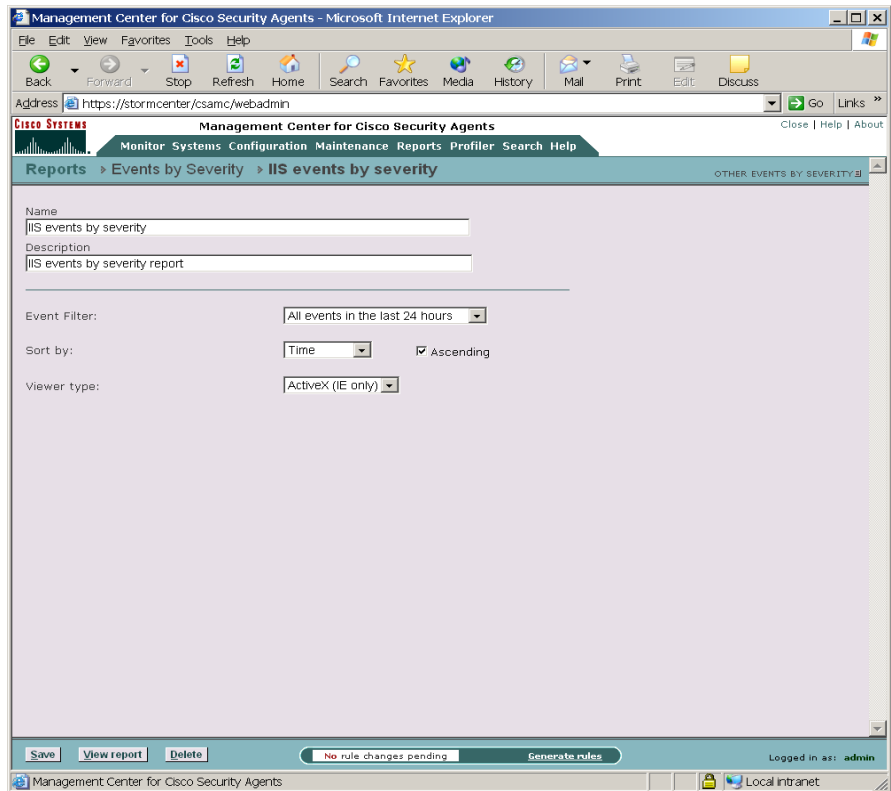
Events by Severity

You can generate reports using various selection and sorting criteria. In this case, you are creating a report based on event severity levels.

To generate an Events by Severity report, do the following.

-
- Step 1** Move the mouse over **Reports** in the menu bar of CSA MC. Select **Events by Severity** from the drop-down list that appears. Any existing reports are shown.
 - Step 2** Click the **New** button to create a new report. This takes you to the configuration view.
 - Step 3** In the Events by Severity report configuration view, enter a **Name** and a **Description** for the report.
 - Step 4** From the pulldown list, select an **Event Filter**. This is an Event Set you create from the Monitor> Event Sets configuration view (see the “[Event Sets](#)” section on page 8-20).
 - Step 5** From the **Sort by** pulldown list, select a parameter for sorting this report's contents (see [Figure 9-1](#)).
 - Step 6** Enable or Disable the **Ascending** checkbox depending on the order in which you want to view your reports.
 - Step 7** Select a **Viewer type**. By default, ActiveX is selected. This is the recommended viewer. For more information, see [Viewing Reports, page 9-2](#).
 - Step 8** Click the **Save** button to save the parameters you've just configured for generating this report.
 - Step 9** Click the **View Report** button and the report is automatically displayed in a new browser window.

Figure 9-1 Events by Severity Report Configuration



Events by Group

You can generate reports using various selection and sorting criteria. In this case, you are creating a report based on the groups that have generated the events.

To generate an Events by Group report, do the following.

- Step 1** Move the mouse over **Reports** in the menu bar of CSA MC. Select **Events by Group** from the drop-down list that appears. Any existing reports are shown.
- Step 2** Click the **New** button to create a new report. This takes you to the configuration view.

- Step 3** In the Events by Group report configuration view, enter a **Name** and a **Description** for the report.
- Step 4** From the pulldown list, select an **Event Filter**. This is an Event Set you create from the Monitor> Event Sets configuration view (see the “[Event Sets](#)” section on [page 8-20](#)).
- Step 5** From the **Sort by** pulldown list, select a parameter for sorting this report's contents.
- Step 6** Enable or Disable the **Ascending** checkbox depending on the order in which you want to view your reports.
- Step 7** Select a **Viewer type**. By default, ActiveX is selected. This is the recommended viewer. For more information, see [Viewing Reports, page 9-2](#).
- Step 8** Click the **Save** button to save the parameters you've just configured for generating this report.
- Step 9** Click the **View Report** button and the report is automatically displayed in a new browser window.

Host Detail

You can generate reports based on hosts in specific groups you select as part of the report. A host detail report provides in-depth information on the hosts in the groups you select for the report.

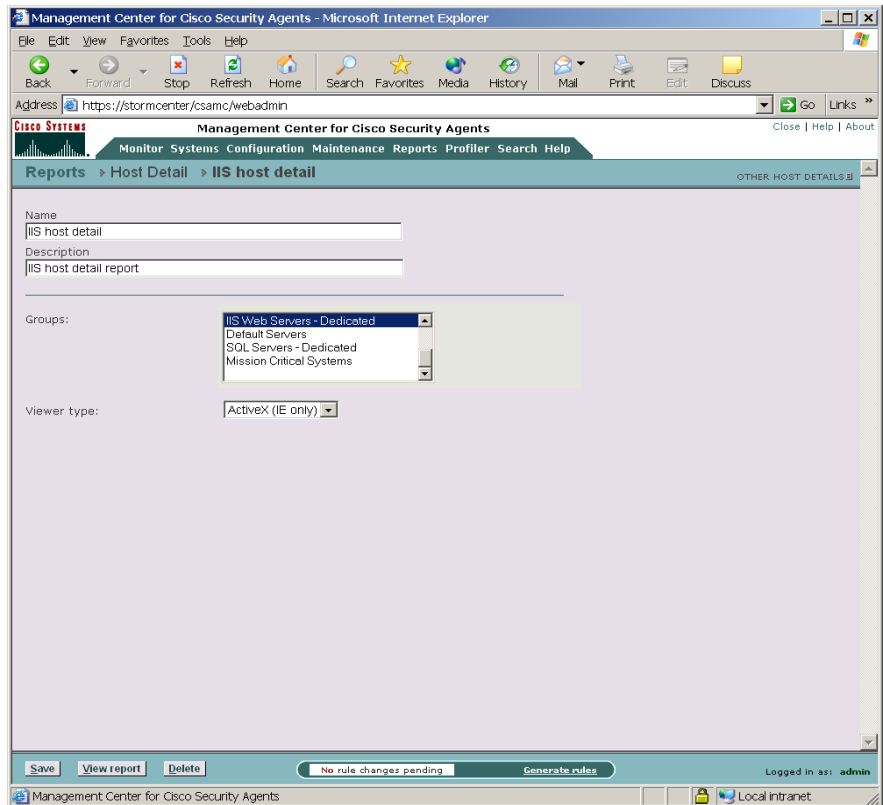
To generate a host detail report, do the following.

-
- Step 1** Move the mouse over **Reports** in the menu bar of CSA MC. Select **Host Detail** from the drop-down list that appears. Any existing reports are shown.
 - Step 2** Click the **New** button to create a new report. This takes you to the configuration view.
 - Step 3** In the Host Detail report configuration view (see [Figure 9-2](#)), enter a **Name** and a **Description** for the report.
 - Step 4** Select the **Groups** for which you want to generate a report. To select multiple items in a list box, hold down the **Ctrl** key as you select each item. To unselect a single item, hold down the Ctrl key when you click on the item in question. Press the **Shift** key to select multiple successive items. You can also select All Hosts here to generate a report for all registered hosts.

Generating Reports

- Step 5** By default, ActiveX is selected as the **Viewer type**. This is the recommended viewer. For more information, see [Viewing Reports, page 9-2](#).
- Step 6** Click the **Save** button to save the parameters you've just configured for generating this report.
- Step 7** Click the **View Report** button and the report is automatically displayed in a new browser window.

Figure 9-2 Host Detail Report Configuration



Policy Detail

You can generate reports by selected policies. A policy report provides in-depth information on the policies you select for the report.

To generate a policy detail report, do the following.

-
- Step 1** Move the mouse over **Reports** in the menu bar and select **Policy Detail** from the drop-down list that appears.
 - Step 2** Click the **New** button to create a new report. This takes you to the configuration view.
 - Step 3** In the Policy Detail report configuration view, enter a **Name** and a **Description** for the report.
 - Step 4** Select the **Policies** for which you want to generate a report. To select multiple items in a list box, hold down the **Ctrl** key as you select each item. To unselect a single item, hold down the **Ctrl** key when you click on the item in question. Press and hold the **Shift** key to select multiple successive items.
 - Step 5** By default, ActiveX is selected as the **Viewer type**. This is the recommended viewer. For more information, see [Viewing Reports, page 9-2](#).
 - Step 6** Click the **Save** button to save the parameters you've just configured for generating this report.
 - Step 7** Click the **View Report** button. Your report is created and is automatically displayed in a new browser window.

Group Detail

You can generate reports by selected group or groups. A group report provides in-depth information on the groups you select for the report.

To generate a group detail report, do the following.

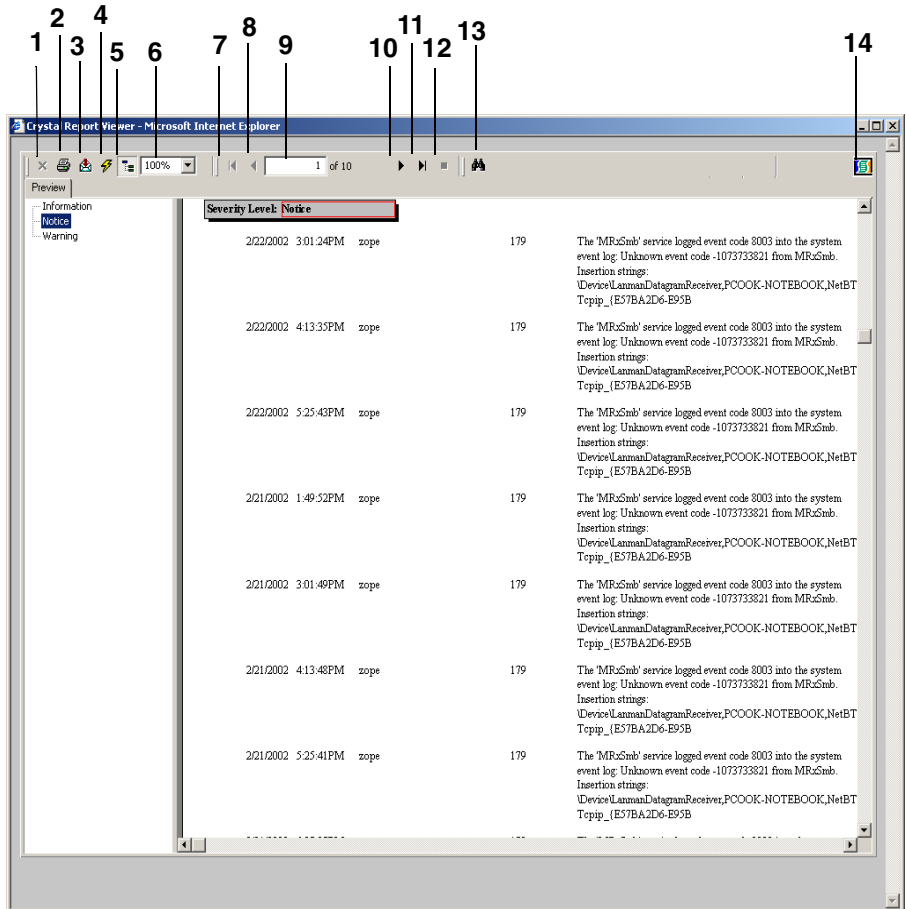
-
- Step 1** Move the mouse over **Reports** in the menu bar and select **Group Detail** from the drop-down list that appears.
 - Step 2** Click the **New** button to create a new report. This takes you to the configuration view.

- Step 3** In the Group Detail report configuration view, enter a **Name** and a **Description** for the report.
- Step 4** Select the **Groups** for which you want to generate a report. To select multiple items in a list box, hold down the **Ctrl** key as you select each item. To unselect a single item, hold down the Ctrl key when you click on the item in question. Press and hold the **Shift** key to select multiple successive items.
- Step 5** By default, ActiveX is selected as the **Viewer type**. This is the recommended viewer. For more information, see [Viewing Reports, page 9-2](#).
- Step 6** Click the **Save** button to save the parameters you've just configured for generating this report.
- Step 7** Click the **View Report** button. Your report is created and is automatically displayed in a new browser window.

About the ActiveX Crystal Report Viewer

The report viewer for Active X reports contains elements that allow you to print, export, and search reports.

Figure 9-3 Crystal Report Viewer



Descriptions of the viewer button controls are as follows:

1. Close current view: This X button is generally grayed out unless a drill down item has been viewed. When active, this button allows you to close the drill down preview.
2. Print: This button allows you to print the report. When you click the button, a print setup window appears. From there, you can change the page setup of the report and select the print range.

3. Export report: The envelope button with the red arrow line allows you to export reports. This export utility lets you export reports to four file format choices (providing the applications for those formats are installed on the client system).

To export a report, click the Export button and the "Crystal Smart Viewer Export" window appears. Select where you want to save the exported file to from the **Save in** field pulldown arrow. Then enter a name for the exported file in the **File name** field. From the **Save as type** pulldown field, select the format to save the report as. The choices are: Crystal Report, Rich Text Format, Word Document, and Excel Document.

When you've made your format selection, click the **Save** button. A progress box lets you know the report is being exported.

4. Refresh: The refresh button is the one with the lightning bolt on it. When you click this button, a message is sent to the server to repopulate the report with fresh data.
5. Toggle group tree: The toggle group tree (navigation button) appears as a miniature directory tree. Clicking this button causes the left pane of the viewer to appear or disappear.
6. Zoom control: This control lets you set the magnification control of the viewer. You can select levels from the pulldown menu or you can enter your own numbers here.
7. Go to first page: Clicking the "Go to first page" control arrow takes you to the first page of the report. (When you are on the first page, this arrow is grayed out.)
8. Go to previous page: Clicking the "Go to previous page" control arrow takes you to the page of the report that is directly before the page you are currently viewing.
9. Current page out of total pages viewed: The number that appears in this field is the number of the page that is currently displayed. If your report has several pages, this number is followed by the word "of" and another number indicating the total number of pages in the report. This field can also function as a "go to" control. Entering a number in this field and pressing the Enter key takes you to that page in the report.
10. Go to next page: This right facing arrow button, when clicked, takes the viewer to the page immediately following the current page being viewed.

11. Go to last page: This right facing arrow button touching a vertical line, when clicked, takes the viewer to the last page of the report.
12. Stop loading: This button is generally grayed out. It contains a solid square in the middle and is only active when a report is loading. Clicking this button stops the report from loading in the window but it does not stop the report from continuing to be processed on the CSA MC system.
13. Search text: Clicking this button (with binoculars on it) allows you to search for any specific text, number, or character in the report. After entering the information you're searching for and clicking the "Find next" button, any matching items in the report are enclosed in a red-lined box.
14. Logo box: This box in the upper right corner of the viewer contains the viewer logo. This logo appears as a graphical rotating image when the viewer is busy loading.



Using Management Center for Cisco Security Agents Utilities

Overview

The Management Center for Cisco Security Agents provides various utilities for advanced product maintenance tasks that extend beyond the administrator configuration and policy generation tasks done through the CSA MC user interface. Those utilities are documented here.

This section contains the following topics.

- [Start and Stop Server Service, page 10-2](#)
- [Start and Stop Agent Service, page 10-2](#)
- [Backing Up Configurations, page 10-3](#)
- [Restoring Backup Configurations, page 10-6](#)
- [Free Up Disk Space on CSA MC \(Insufficient Disk Space Event\), page 10-7](#)
- [Using the COM Extract Utility, page 10-9](#)
- [Manual Agent Data Filter Installation, page 10-10](#)
- [Exporting and Importing Configurations, page 10-12](#)

Start and Stop Server Service

Stop and start the Management Center for Cisco Security Agents service on a host by running the following commands from a command prompt window on the server host system:

```
net stop "Cisco Security Agent MC"  
net start "Cisco Security Agent MC"
```

Start and Stop Agent Service

Administrators can stop and start the Cisco Security Agent service on a Windows host by running the following commands from a command prompt window on the agent host system:

```
net stop "Cisco Security Agent"  
net start "Cisco Security Agent"
```

Administrators can stop and start the Cisco Security Agent service on a UNIX host by running the following commands from a command prompt window on the agent host system:

```
/etc/init.d/csa stop  
/etc/init.d/csa start
```

**Note**

Running this stop command to stop agent services on a system disables all rules on that system. Running a start csa command starts the agent service and reinstates all rules.

The shipped mandatory UNIX policy, "Secure Management Module," allows secured management applications to stop the agent service. For example, after having logged in by selecting Command Line Login from the login screen in the options menu, you can issue the command `/etc/init.d/csa stop`. Refer to the policy in CSA MC to see how these secured management applications are already defined and may be modified using application builder rules.

**Note**

The UNIX agent has a utility (csactl) to provide capabilities that the Windows agent provides in its user interface. See [Appendix A, “Cisco Security Agent Overview”](#) for details.

If an agent has a policy containing an Agent service control rule that denies the stopping of the agent, administrators cannot stop the agent service on the system in question. See the [“Agent Service Control”](#) section on page 4-24.

Backing Up Configurations

It is a good idea to back up your management server configurations at regular intervals. If your server system fails for any reason, and a copy of your configuration database is not stored elsewhere, you could lose your policy information.

Note that if you have a disk crash or some other type of failure on your management server, to recover, in addition to restoring your backup configuration, you must also restore the original server certificate information. If you do not do this, your existing agents cannot register with the restored CSA MC. All this information is saved as part of the backup process.

The **Backup Configuration** feature, available from the **Maintenance** category in the menu bar, lets you backup your database and your certificate files together at regularly schedules intervals or as needed.

To backup your CSA MC configuration, do the following.

Step 1 Move the mouse over **Maintenance** in the menu bar and select **Backup Configuration** from the drop-down list that appears.

Step 2 In the Backup Configuration window you can select the following radio buttons:

- **No database backup:** Select this option if you do not want backups to occur automatically at scheduled intervals but want to perform them manually. After selecting this radio button, enter the **directory path** (including drive letter) to which you want to save your backup configuration. Then click the **Backup now** button.
- **Scheduled database backup:** Select this option to schedule regular backups and then choose one of the scheduled backup options: low frequency, medium frequency, and high frequency. Enter the directory path (including drive letter) to which you want to save your backup configuration and click the Save button. Backups will now occur as scheduled.

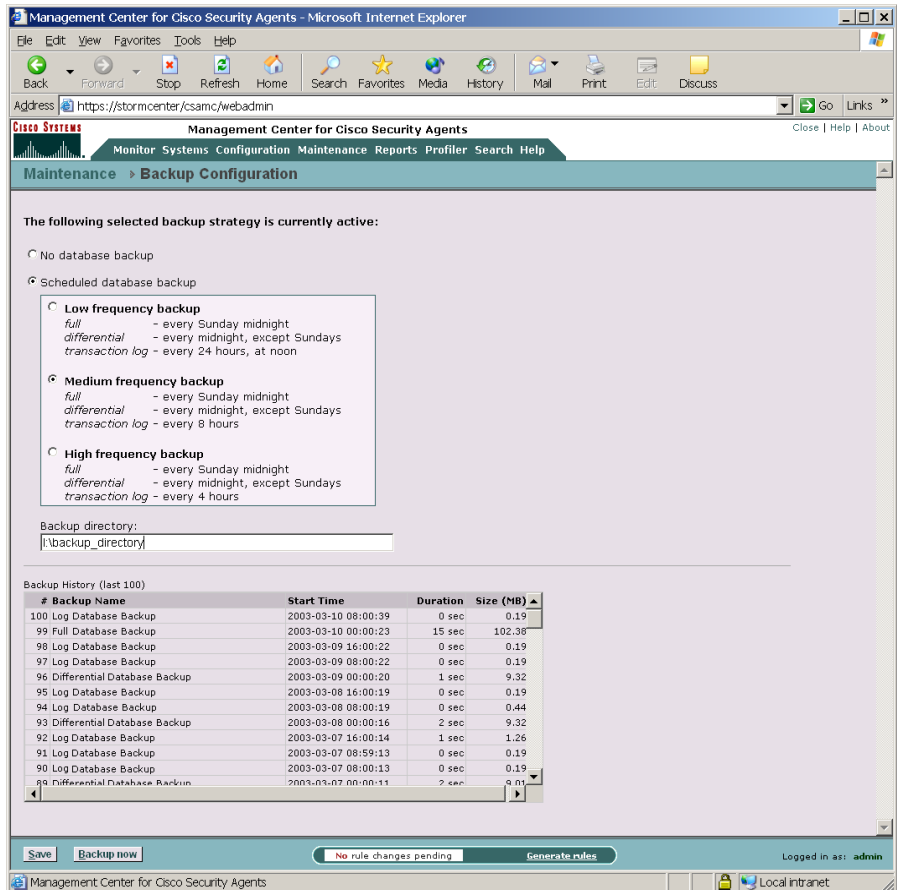
Backup types are categorized as follows:

full—A full backup occurs every Sunday night at midnight. This full backup includes the entire database with license and certificate information.

differential—This type of backup occurs every night at midnight (except Sunday nights when a full backup occurs). A differential backup includes only data that has changed since the last backup (full or differential) occurred.

transaction log—This backup occurs every 24 hours (low), 8 hours (medium), or 4 hours (high) depending on the frequency you select. The presence of this transaction log allows administrators to back out configuration changes to a certain point. Please refer to Microsoft documentation for details about the transaction log.

Figure 10-1 Backup Configuration Window



Backup Files appears as follows in the directory you select:

- full_backup_[db_name].bak—for full backups
- diff_backup_[db_name].bak—for differential backups
- log_backup_[db_name]_[x].bak—for log backups, where x is an integer from 1 to 23 (backup hour)
- crt_log_backup_[db_name].bak—for current transaction log backup

Restoring Backup Configurations

Restore backup CSA MC configurations, including database, license, certificate information, and transaction logs by running the Restore utility, called **Restore Configuration**, located in the default CSAMC\bin directory:

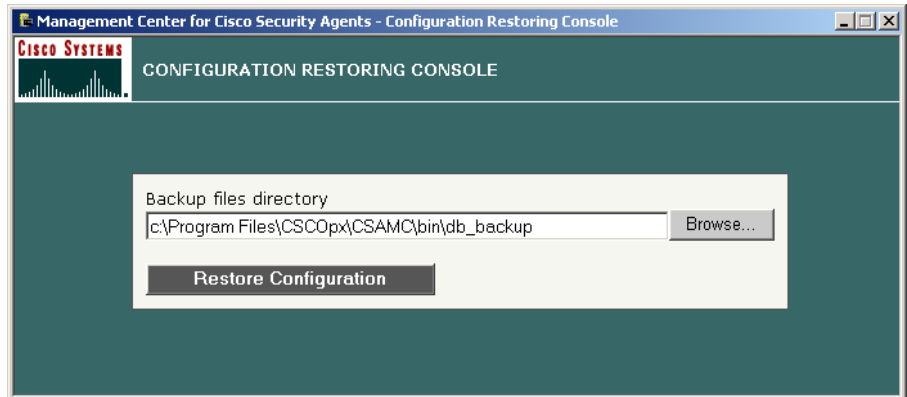


Note

If you are restoring a backup configuration due to a disk failure, after you re-install CSA MC and then restore the backup configuration, you may find that the final set of uncommitted transactions were lost.

- Step 1** Double-click the **Restore Configuration** file (located in the CSCOPx\CSAMC\bin directory) to display the CSA MC restore user interface. See [Figure 10-2](#).
- Step 2** Enter the directory path where the backup files are stored. (The default storage directory appears here automatically.)

Figure 10-2 Database Restoring Console



- Step 3** Click the **Restore database** button. When you click the Restore database button, you are asked if you want to restore the backup configuration. Click **Yes** to do so. The restore process now takes place. Once the restore is complete, a log file, the Database Restoring Log, is displayed.

**Note**

When you restore backup configurations, you cannot select to restore only the transaction log, or only a differential backup. All files are automatically restored from the most recent backup that exists in the directory.

Free Up Disk Space on CSA MC (Insufficient Disk Space Event)

Use the following procedure to shrink your database files and log files and increase the amount of free disk space on your CSA MC system. If your CSA MC event log contains an "insufficient disk space" message, this is an appropriate procedure for freeing up space.

Delete old log files and old CSA MC database files

Step 1 Stop the CSA MC and agent services.

From a command prompt window, type the following commands:

```
net stop "Cisco Security Agent"  
net stop "Cisco Security Agent MC"
```

Step 2 Access the `CSCOPx\CSAMC\db` and `CSCOPx\CSAMC\cfg` folders and delete old database files with names such as:

```
csamc.mdf.Vm.n.0.ddd, csamc_log.ldf.Vm.n.0.ddd,  
csamc_volatile_data.ndf.Vm.n.0.ddd.
```

**Caution**

DO NOT delete files `csamc.mdf`, `csamc_log.ldf` and `csamc_volatile_data.ndf`. (Files with no old version number suffixes.)

Step 3 Access the `CSCOPx\CSAMC\log` directory and delete old log files of the form `csamc.nnn`.

Step 4 Start the CSA services. From a command prompt window, type the following commands:

```
net start "Cisco Security Agent"  
net start "Cisco Security Agent MC"
```

Purge old CSA MC Events

-
- Step 1** Log in to CSA MC and select **Monitor->Event Log**. Note the total number of events. This number will help you gauge how many events are available for purging.
- Step 2** Now go to **Monitor->Event Sets** and choose an existing event set or create a new one. An event set always provides a purge option. At this point it is up to you as to which events to purge. For example,
- "Choose "All Events In the Last 7 Days"
 - "Choose "Include only the following selected severity levels" and choose "Information", "Notice" & "Warning"
 - "Optionally choose a time range, other than "Last 7 Days", that will clean up a lot of old events
- Step 3** Then click the **Purge Events** button at the bottom of the screen, and wait for the task to complete.
- Step 4** Repeat steps 3 and 4 until you have purged all events that are not needed.

Shrink the CSA MC database

-
- Step 1** Stop the CSA MC and agent services. From a command prompt window, type the following commands:
- ```
net stop "Cisco Security Agent"
net stop "Cisco Security Agent MC"
```
- Step 2** If do not have a scheduled log backup configured (from CSA MC **Maintenance->Backup Configuration** window) type the following command (If you have a scheduled backup, skip this step and proceed to step 3.):
- ```
osql -E -Q "backup log csa with no_log"
```
- Step 3** From a command prompt window, type the following
- ```
osql -E -Q "dbcc shrinkdatabase (csa)"
```
- Step 4** Start the CSA MC and agent services. From a command prompt window, type the following commands:
- ```
net start "Cisco Security Agent"
net start "Cisco Security Agent MC"
```

Win2K Disk Cleanup

-
- Step 1** From the Windows **Start** menu, select **Programs -> Accessories -> System Tools->Disk Cleanup**.
 - Step 2** Choose the disk drive where the CSA MC directories reside.
 - Step 3** Check all the boxes and click **OK**. At this point the agent may issue some queries concerning dll's and ocx's that are being modified. Answer **Yes** to all queries.
The Windows Disk Cleanup utility deletes the files and exits.

Using the COM Extract Utility

CSA MC provides a COM component extraction utility, called `extract_com`, which installs in the `\CSCOpX\CSAgent\bin` directory with each Cisco Security Agent. Running this utility extracts all COM component PROGID's and CLSID's for software running on the system in question and places this data into a text file. You can cut and paste these ID's from the text file into your COM component sets and access rules.

Run the `extract_com` utility on an agent system in the following manner:

-
- Step 1** Open a command prompt window.
 - Step 2** From the `\Cisco\CSAgent\bin` directory type in `extract_com filename` "filename" is the name of the text file you want the utility to create. It is into this file that all COM PROGID and CLSID data is placed.

For example, enter:

```
\Cisco\CSAgent\bin>extract_com foo.txt
```

The Cisco Security Agent creates the "foo.txt" file in the same `\bin` directory as the `extract` utility. You can access it from there.

**Caution**

Both COM Component access control rule fields and Variable COM Component set fields require a very specific syntax for entering PROGID's and CLSID's. The COM component file created by the extract_com utility may display PROGID's and CLSID's without the proper syntax in the output file. Despite this, when you enter these ID's into text fields for rules or variables you **MUST** use the correct syntax detailed on [page 7-22](#).

Manual Agent Data Filter Installation

On Windows platforms, if you install Web server software (IIS or Apache) after you install the Cisco Security Agent on the server system, or if you have installed the Web server in a directory other than the default, you must manually install the Cisco Security Agent data filter in order to use Data access control rules on the system in question.

**Note**

If your Web server software is already installed (in its default directory) when you install the agent on Windows, the server software is detected by the agent and the data filter capability is automatically installed with the agent.

On Solaris, in order to use Data access control rules (on Apache or IPlanet servers) you must install the data filter manually after you install the Cisco Security Agent. Unlike Windows, the Solaris installation does not detect Web server software and does not install the data filter with the agent. You must always manually install it.

Install Data Filter on Windows

If you have installed Web server software (IIS or Apache) after you install the Cisco Security Agent on the server system, or if you have installed the Web server in a directory other than the default, run the following command(s) to manually install the CSA data filter on the server system making use of Data access control.

For a Microsoft IIS Web server, run the following command:

```
csa_datafilter -i iis
```

For an Apache Web server, run one of the following Apache version appropriate commands:

```
csa_datafilter -i apache 1.3 <.conf file with full path name>  
<modules directory path>  
csa_datafilter -i apache 2.0 <.conf file with full path name>  
<modules directory path>
```

For example, if Apache 2.0 was installed with its default settings after the agent is installed, you would run the following command to install the data filter.

```
csa_datafilter -i apache 2.0 "c:\Program  
Files\Apache\conf\httpd.conf" "c:\Program Files\Apache\modules"
```

**Note**

If there are spaces in the directory path, you must put quotations around the pathname.

**Caution**

You must restart the web server service after the data filter is installed for data access control rules to take effect.

Uninstall Data Filter on Windows

For a Microsoft IIS Web server, run the following command to uninstall the data filter:

```
csa_datafilter -u iis
```

For an Apache Web server, run one of the following Apache version appropriate commands to uninstall the data filter:

```
csa_datafilter -u apache 1.3 <.conf file with full path name>  
<modules directory path>  
csa_datafilter -u apache 2.0 <.conf file with full path name>  
<modules directory path>
```

Install Data Filter on Solaris

Run the following command to manually install the CSA data filter on the server system making using of Data access control.

```
webserver:root>./csa_datafilter -i
```

[output:]

```
CSA web server filter installation:
Should I install filters for IPlanet and/or Apache [No] y
Enter the path of the IPlanet config directory (null for none):
/usr/iplanet/servers/https-webserver/config
Enter the path of the Apache root (null for none):
webserver:root>
```



Caution

You must restart the web server service after the data filter is installed for data access control rules to take effect.

Uninstall Data Filter on Solaris

```
webserver:root>./csa_datafilter -u
```

[output:]

```
SA web server filter removal:
Should I uninstall filters for IPlanet and/or Apache [No] y
Enter the path of the IPlanet config directory (null for none):
/usr/iplanet/servers/https-webserver/config
magnus.conf saved as magnus.conf.sav
obj.conf saved as obj.conf.sav
The CSA agent filter for IPlanet 6.0 has been removed
Enter the path of the Apache root (null for none):
webserver:root>
```

Exporting and Importing Configurations

Under the Maintenance category in the menu bar, use the Export utility to export your policies to other CSA MCs. If you have multiple CSA MCs, you might want to export some basic policies to those servers for deployment. Likewise, using the Import utility, you can download and import those policies as well as preconfigured policies that Cisco provides.

**Note**

The Export utility exports policies (not individual rules), including the accompanying application classes and configuration variables. Because of communication channels established in the original configuration, some site-specific imported configuration information (IP addresses) may not work on another server. Exporting an item will also export related data. In particular, exporting policies will export application classes and configuration variables referenced in rules within the policy. Exporting a group will export associated policies but not hosts.

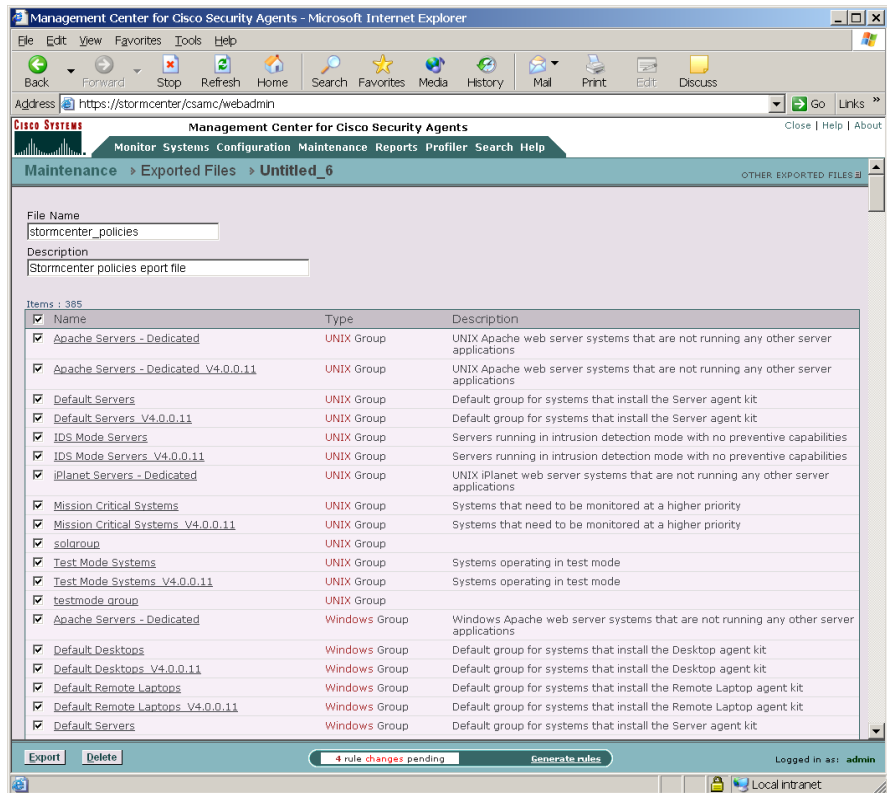
**Caution**

The Export/Import functions are not intended to be used as a backup/restore mechanism as they do not preserve system specific information such as group-host memberships.

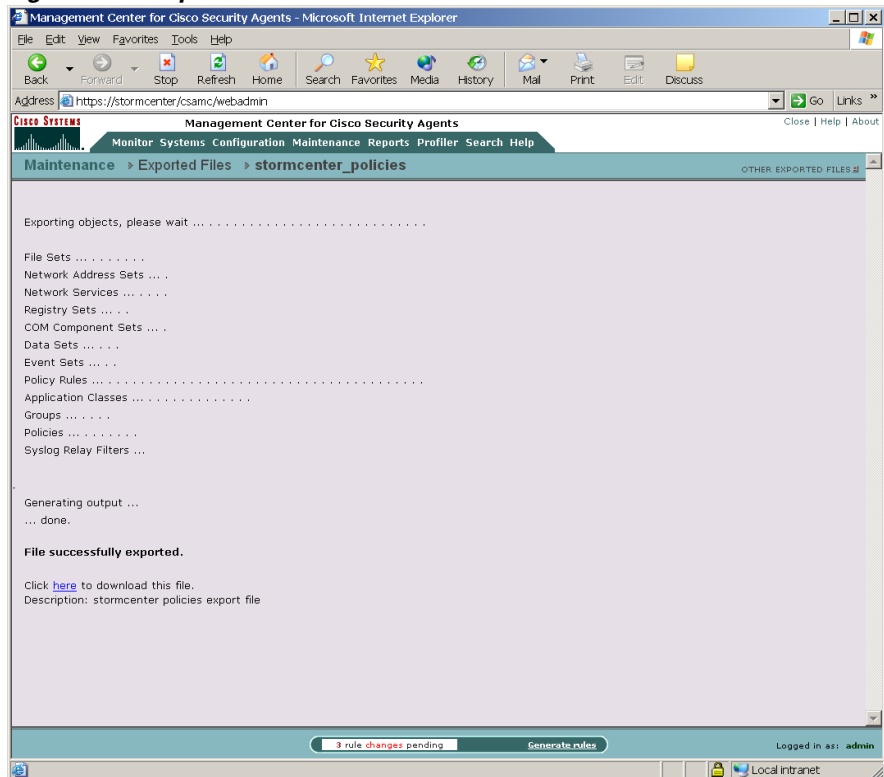
To Export configurations, do the following.

- Step 1** From the menu bar **Maintenance** drop-down list, move the mouse over **Export/Import**. A cascading menu with further selections appears. Select **Export Files** from the drop-down list that appears. Any previously exported files are shown.
- Step 2** Click the **New** button to create a new exported file. This takes you to a checkbox list of all configuration items.
- Step 3** **Check** the box beside the configurations you want to export. (See [Figure 10-3](#)).

Figure 10-3 Export Configurations



- Step 4** At the top of the page, enter a **File Name** for the exported file you are creating. CSA MC will append an ".export" extension to the file name you enter.
- Step 5** Click the **Export** button. The files are exported under the file name you create. Now you must save the file to the system.
- Step 6** Once the export has completed, a link is displayed that allows you to save the exported file. The link reads "Click [here](#) to download this file." **Click** on the "here" link to save the file to a directory you specify (see [Figure 10-4](#)).
Once you save the file, you can import it to any server.

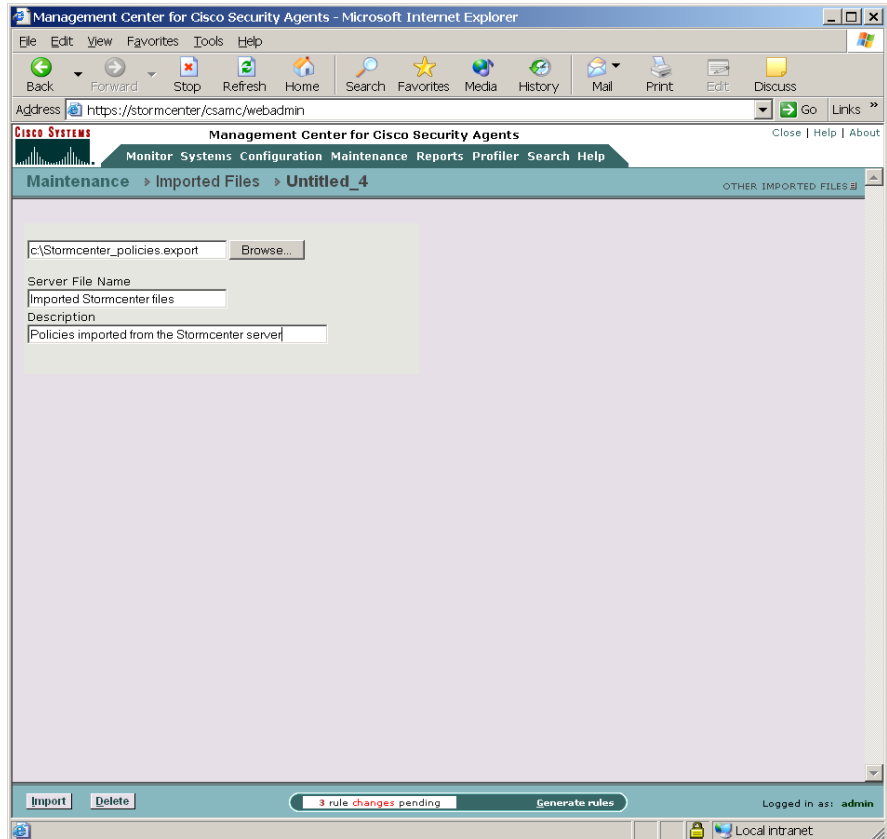
Figure 10-4 Export Download View

To Import configurations, do the following.

- Step 1** From the menu bar **Maintenance** drop-down list, move the mouse over **Export/Import**. A cascading menu with further selections appears. Select **Import Files** from the drop-down list that appears. Any previously imported files are shown.
- Step 2** Click the **New** button to create a new imported file. This takes you to the configuration Browse view (see [Figure 10-5](#)).
- Step 3** Click the **Browse** button to locate the exported file you want to import.
- Step 4** Once you have located the file to import, enter a **Server File Name** by which you can identify the configuration file you're importing. This is the name that will appear under the Imported Files link once the import is complete.

Step 5 Enter a Description and then click the **Import** button to import the configuration.

Figure 10-5 Import Configurations



Imported files are automatically entered into the CSA MC database of the server you're importing them to. You don't have to do anything beyond the import function to unpack the exported file.

**Note**

When you import configuration items, if it is found that there is already an existing exact match for an item, the new configuration data is not copied over. Instead, the existing item will be reused. But if the import process finds that there is an existing item with the same name as a new one, but with different configuration components (variables, etc.), the existing item is renamed by appending the version number (V3.2.0.105, for example) to the name. The imported item is then copied into the database with no version number so that both items can co-exist in the database.

Therefore, for any partial configuration item duplications that may exist after an import, the item with no version number appended to its name is always the most recently added item.

If you already have two existing similar configuration items in the database (one already renamed with the version number) then when the third data item is imported, the original item is renamed with the version number and with a "_0" notation appended to differentiate it.



Using Cisco Security Agent Profiler

What is Profiler

Cisco Security Agent Profiler software works with CSA MC and the Cisco Security Agent, serving as a data analysis and policy creation tool for administrators who are deploying policies across systems and networks.

Because the rules that comprise CSA MC policies are application-centric, understanding the resources applications require for normal operations is integral to building effective policies. Profiler does that by analyzing applications as they operate in a normal environment and generating useful policies based on that analysis.

This section contains the following topics.

- [How Profiler Works, page 11-2](#)
- [The Analysis Process, page 11-3](#)
- [Configure an Analysis Job, page 11-6](#)
- [Start Analysis, page 11-11](#)
- [Importing the Policy, page 11-12](#)
- [The Profiler Policy, page 11-14](#)
- [Profiler Reports, page 11-18](#)

How Profiler Works

When deployed on a system running a Cisco Security Agent, Profiler monitors the actions of designated applications on that system, logging all resource access attempts made by the application. Profiler then analyzes the logging data it collects and develops a policy for the application in question. This policy enforces what is determined to be normal application behavior while restricting all other behaviors. These other behaviors could now be construed as abnormal or suspicious based on the analysis.

**Note**

If you are creating your own policies and not using Profiler, refer to [Chapter 12, “Policy Definition Guidelines”](#) for information.

The Analysis Process

The application analysis and policy creation process is performed by three different contributing components: CSA MC, the agent (logging agent), and the Profiler Analysis Job.

- Through *CSA MC*, you designate which application you want to analyze. You also select an agent host on which the analysis is to take place and a time frame within which the analysis will be completed. This analysis configuration is then sent to the agent on the selected host in the same way policies are sent to agents.

Profiler examines all the logged data it receives from the logging agent. When the analysis is complete, it creates a policy for the application and generates reports containing information on all resources accessed by the application. The policy enforces the normal operations seen in the log file and will deny any operation attempts by the application that do not align with this normal behavior.

- The *agent* receives the analysis configuration information when it next polls in to *CSA MC*. This agent now becomes the "logging agent" in this process. It logs all operations performed by the designated application. As this logging takes place, it is assumed that the application is being thoroughly exercised in a normal operating environment. When the analysis is complete, the logged data is sent to the Profiler Analysis Workstation.

CSA MC imports the policy created by Profiler.

Profiler Analysis Jobs

When the Profiler license is copied to CSA MC, a new menu item is added to the CSA MC menu bar. This item is called **Profiler**. It is by accessing this **Analysis Jobs** window from the Profiler menu item that you can configure parameters for analyzing a particular application.



**Note**

Users of StormWatch 3.2 and StormFront 2.5 should note that Profiler (formerly StormFront) now installs automatically with CSA MC and becomes usable when the appropriate Profiler license key is copied to CSA MC.

When you are ready to configure an analysis job for an application, you must have the following information:

- What application you want to analyze for this particular job: You should have an appropriate application class configured for the analysis. (You can leverage existing application classes, but it is recommended that you analyze only one application at time. See [page 11-8](#) for more information.)
- Which host you want to select for application analysis: You should have an appropriate host chosen for the analysis job.

Creating, Saving, and Cancelling Job Data

Management Center Button Frame

Similar to most CSA MC windows, analysis job action items appear in a frame at the bottom of CSA MC.

**Note**

The available buttons in the bottom frame change in accordance with the actions available for the page you're viewing. With Profiler, several actions are performed from the same page as the analysis job progresses. You may have to refresh the analysis jobs page for the buttons to change appropriately.

Available buttons and links are as follows.

- **New:** Use the New button to create new a configuration item within the list view you have selected. Click the New button and a new item appears in the list view. Click the new item link to access the configuration view for that item.
- **Delete:** Use the Delete button in conjunction with the checkboxes beside each list view item. To delete a configuration, select its checkbox (you can select several at once) and click the Delete button. All checked items are deleted. To quickly select all checkboxes, click the very top checkbox in the list view heading bar. Clicking the Delete button then deletes all items.
- **Clone:** Use the Clone button in conjunction with the checkboxes beside each list view item. To clone a particular configuration, select its checkbox and click the Clone button. You can clone one item at a time. New links to the cloned configurations appear in the list view.

**Note**

When you clone an item that contains variable items like application classes, the cloned item uses the same variables used in the original item. The variables themselves are not cloned.

- **Save:** When you enter configuration information, whether you are entering new data or editing existing data, you must click the Save button once you are finished to save your configuration in the CSA MC database. If you do not click Save before moving to another page in CSA MC, your data is lost.
- **Stop Logging:** If you want to stop the job early and send collected data to the workstation for analysis, click this Stop logging button.
- **Start analysis:** When the logging for the analysis job is complete, a "Start analysis" button appears in the bottom frame of the Analysis Job page. Click this button to have the analysis workstation begin to analyze the logging data.
- **Import:** When the analysis of the logging data is complete, Profiler creates a policy which you can import into CSA MC. The "Import" button appears when the policy creation is complete.

Configure an Analysis Job

To configure an analysis job, do the following:

**Note**

In some cases, you can configure an analysis job using the Event Management Wizard accessible from particular event log entries. See the [“About the Event Management Wizard” section on page 8-9](#) for more information.

-
- Step 1** Move the mouse over **Profiler** in the menu bar and select **Analysis Jobs** from the drop-down list that appears. The list of existing jobs (if any) is displayed.
 - Step 2** Click the **New** button to create a new analysis job. This takes you the analysis job configuration page. (See [Figure 11-1](#).)
 - Step 3** Enter a **Name** for the analysis job you are creating.
 - Step 4** Enter a **Description** for your analysis job. This description becomes visible in the analysis job list view.
 - Step 5** **Verbose logging mode:** By default, Profiler filters its logging process so that duplicate events are not logged. You can turn this feature off by selecting this checkbox. If you do turn this filtering off, your logs will be a great deal larger, but the advantage is that you will be able to see how often the same resource is accessed when you view the Profiler reports.

**Note**

The **Target operating system** you selected is displayed in a read-only field. The **Job status** field is also a read-only field. It displays text, informing you of each stage of the analysis job. When you first configure your job, it displays "Not yet deployed."

Figure 11-1 Analysis Job Configuration Window

The screenshot shows the Management Center for Cisco Security Agents web interface in Microsoft Internet Explorer. The browser address bar shows `https://stormcenter/csamc/webadmin`. The page title is "Management Center for Cisco Security Agents" and the breadcrumb navigation is "Profiler > Analysis Jobs > Word_application".

The main content area displays "Saved changes." and a warning: "You must regenerate the rule program to begin the Profiler analysis job." Below this is the "General" section with the following fields:

- Name: Word_application
- Description: Analyze Microsoft Word
- Target operating system: Windows
- Job status: Not yet deployed
- Verbose logging mode

The "Profiler Analysis Job Configuration" section includes:

- A list box for "Perform an analysis of the selected application classes:" containing:
 - Word_analysis (selected)
 - <Network Applications>
 - <Processes created by Network Applications>
 - <Processes created by Servers (TCP and UDP)>
 - <Processes executing downloaded content>
- A dropdown for "For the selected host:" set to "cheshire".
- A checkbox for "Disable policy rule enforcement" which is unchecked.
- Start job at time: 03/07/2003 16:13
- End job at time: 03/07/2003 17:13
- Stop job when either of the following occurs:
 - Log file size exceeds 5 MB
 - Application is invoked 3 times

At the bottom, there is a "Save" button, a status bar indicating "66 rule changes pending", and a "Generate rules" button. The user is logged in as "admin".

- Step 6** In the **Perform an analysis of the selected application classes** list box, select the application class or classes you want to analyze. To select multiple items in a list box, hold down the **Ctrl** key as you select each item. To unselect a single item, hold down the **Ctrl** key when you click on the item in question. Press and hold the **Shift** key when you click on an item to select multiple successive items.

**Caution**

You can select an application class that contains more than one application for the analysis. But in that case, the policy created by Profiler would apply equally to all applications included in the analyzed application class. For example, if the application class you are analyzing contains both Microsoft Word and Microsoft Outlook, the policy created by Profiler would be a combination of the resources required by both applications.

Next you must assign the job to a specific host system.

Step 7

Select the host you are assigning the job to in the **For the selected host** list box. e.g. Note that you cannot have more than one analysis job running on a host at one time.

**Note**

Once the job begins, you can click the **Stop logging** button that appears in the bottom frame. The job stops automatically according to the parameters you enter on this analysis jobs page. But if you want to stop the job early and send collected data to the workstation for analysis, click this Stop logging button.

Step 8

Optionally, you can select to **Disable policy rule enforcement** for the time frame of the job. Otherwise, the analysis takes place only within the confines of enforced policies. Some events may be denied by rules and therefore the analysis may not be complete.

**Caution**

If you select the Disable policy rule enforcement checkbox, when the logging agent receives an analysis job, any policies relevant to the application being analyzed are disabled on the selected host until the job is completed. This may be undesired if the application in question is unknown or is in any way suspicious.

- Step 9** Next you must enter job time frames.
- **Start job at time:** From the pulldown options, select a time for the job to start once the host polls in and receives the analysis job. If you specify no time here, "now" is automatically entered. This means the analysis job will start immediately when the host receives it.
 - **End job at time:** You must enter a time for the job to end. Profiler will not allow you to save the job until you do. When you enter a *log size* parameter or an *application invocation number* in the fields below, they act as overrides of this end time.

- Step 10** Stop job when either of the following occurs:
- **Log file size exceeds __ MB:** You can enter a size restriction on the log file. When it reaches the size you indicate, the analysis is finished. (Note that the maximum log file size you can enter here is 256 MB. This is also the default value.)
 - **Application is invoked __ times:** You can specify an application invocation restriction. Once the application is invoked on the system the number of times you indicate, the analysis is finished.

**Caution**

It is not always appropriate to use an invocation number limit. For example, for server applications, time frame parameters might be a more appropriate criteria for ending a job.

**Note**

If you enter analysis completion parameters in more than one field, the parameter that is reached first is the one that applies.

- Step 11** Click the **Save job** link in the bottom frame of CSA MC to save this job.
- Step 12** Once your job is configured to your satisfaction, click the **Generate rules** link in the bottom frame and continue by clicking the subsequent **Generate** link to distribute the analysis job to the group hosts you've selected.

Depending on the job parameters you've configured, the selected host will begin the analysis job after it polls in to CSA MC and receive the new rules.

**Note**

Keep in mind that if you have configured your analysis job to begin immediately and your agents are configured to poll in to CSA MC once every hour, the analysis job will not begin until the agent next polls in. In this example case, that time frame could be up to one hour. Additionally, be careful not to designate the end time as a time frame that could occur before the agent polls in and receives the job. In this case, the job will not run at all.

Monitoring the Analysis Job

You can check your CSA MC **Event Log** to view the analysis job progression. An event is sent when the job begins and again when it finishes. When a policy is ready to be imported to CSA MC, an event log message appears indicating this.

Figure 11-2 Analysis Job Event Log Messages

The screenshot shows the Management Center for Cisco Security Agents interface in Microsoft Internet Explorer. The browser address bar shows <https://stormcenter/csamc/webadmin>. The page title is "Management Center for Cisco Security Agents". The navigation menu includes "Monitor", "Systems", "Configuration", "Maintenance", "Reports", "Profiler", "Search", and "Help". The "Monitor" section is active, and the "Event Log" is selected. The event log displays the following information:

Viewing 80460 - 80411 of 80460 events [change filter](#)

Event log generation time : 3/7/2003 4:25:17 PM
 Severity : Information - Emergency
 Host : All
 Policy : All
 Events per page : 50

Navigation: [Latest](#) | [Earliest](#)

#	Date	Host	Severity	Event
80460	3/7/2003 4:00:16 PM	cheshire	Information	Log files for analysis job 'Word_application' were sent to the analysis workstation.
80459	3/7/2003 3:53:11 PM	cheshire	Information	Logging for analysis job 'Word_application' has ended.
80458	3/7/2003 3:43:23 PM	cheshire	Information	Logging for analysis job 'Word_application' has begun.

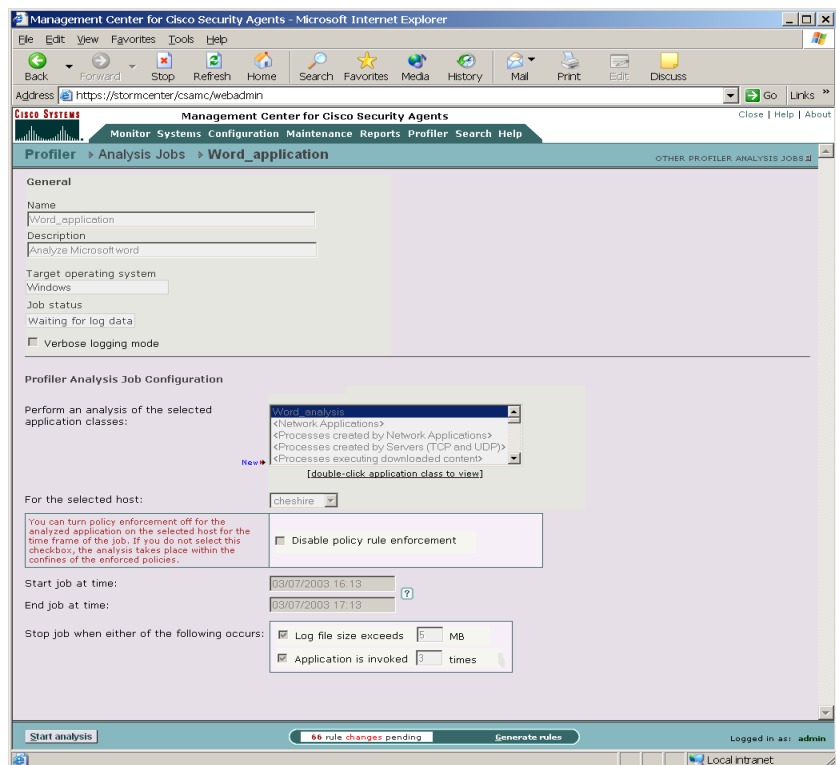
At the bottom of the page, there is a status bar showing "66 rule changes pending" and a "Generate rules" button. The user is logged in as "admin".

Start Analysis

When the Event Log in CSA MC displays "Log files for Profiler analysis job were sent to the analysis workstation", you can begin the data analysis of the logging information on the analysis workstation.

Begin this analysis by accessing the Analysis Job window for this particular job and clicking the **Start analysis** button in the bottom frame (see Figure 11-3). This begins the analysis. An Event Log message appears informing you that "Data analysis for policy creation has started."

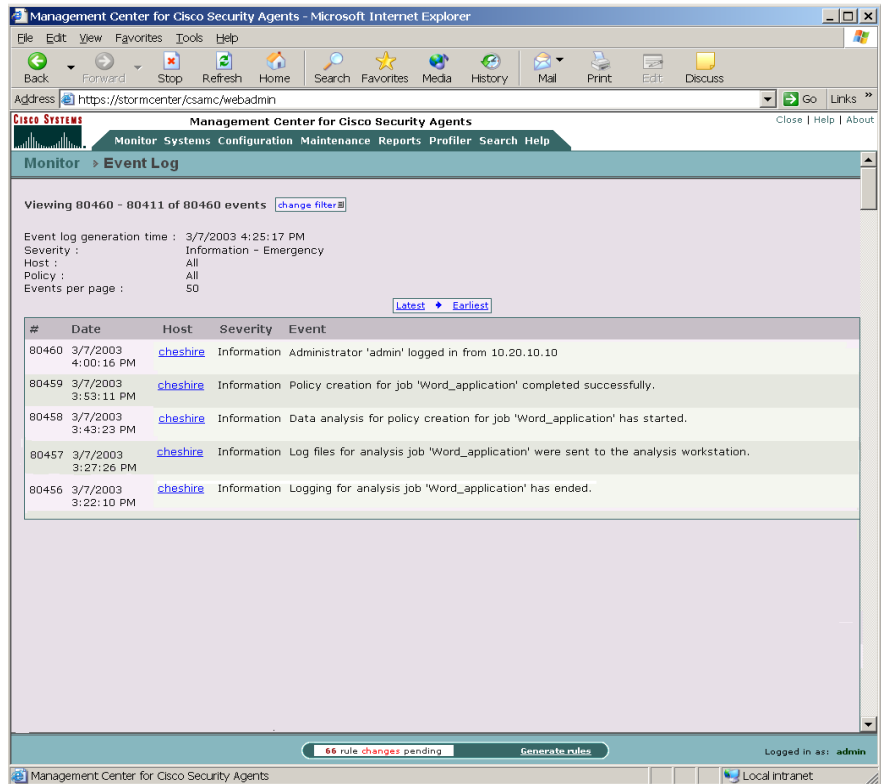
Figure 11-3 Start Analysis of Logged Data



When the analysis is complete, the Event Log file displays the message "Policy creation for Profiler analysis job completed successfully".

Once policy creation is complete, you can import the policy.

Figure 11-4 Event Log Messages for Job Completion



The screenshot shows the Management Center for Cisco Security Agents web interface. The browser window title is "Management Center for Cisco Security Agents - Microsoft Internet Explorer". The address bar shows "https://stormcenter/csamc/webadmin". The page title is "Management Center for Cisco Security Agents". The navigation menu includes "Monitor", "Systems", "Configuration", "Maintenance", "Reports", "Profiler", "Search", and "Help". The "Monitor" menu is expanded to show "Event Log".

The Event Log section displays the following information:

- Viewing 80460 - 80411 of 80460 events
- Event log generation time : 3/7/2003 4:25:17 PM
- Severity : Information - Emergency
- Host : All
- Policy : All
- Events per page : 50

The event log table shows the following entries:

#	Date	Host	Severity	Event
80460	3/7/2003 4:00:16 PM	cheshire	Information	Administrator 'admin' logged in from 10.20.10.10
80459	3/7/2003 3:53:11 PM	cheshire	Information	Policy creation for job 'Word_application' completed successfully.
80458	3/7/2003 3:43:23 PM	cheshire	Information	Data analysis for policy creation for job 'Word_application' has started.
80457	3/7/2003 3:27:26 PM	cheshire	Information	Log files for analysis job 'Word_application' were sent to the analysis workstation.
80456	3/7/2003 3:22:10 PM	cheshire	Information	Logging for analysis job 'Word_application' has ended.

At the bottom of the interface, there is a status bar showing "66 rule changes pending" and a "Generate rules" button. The user is logged in as "admin".

Importing the Policy

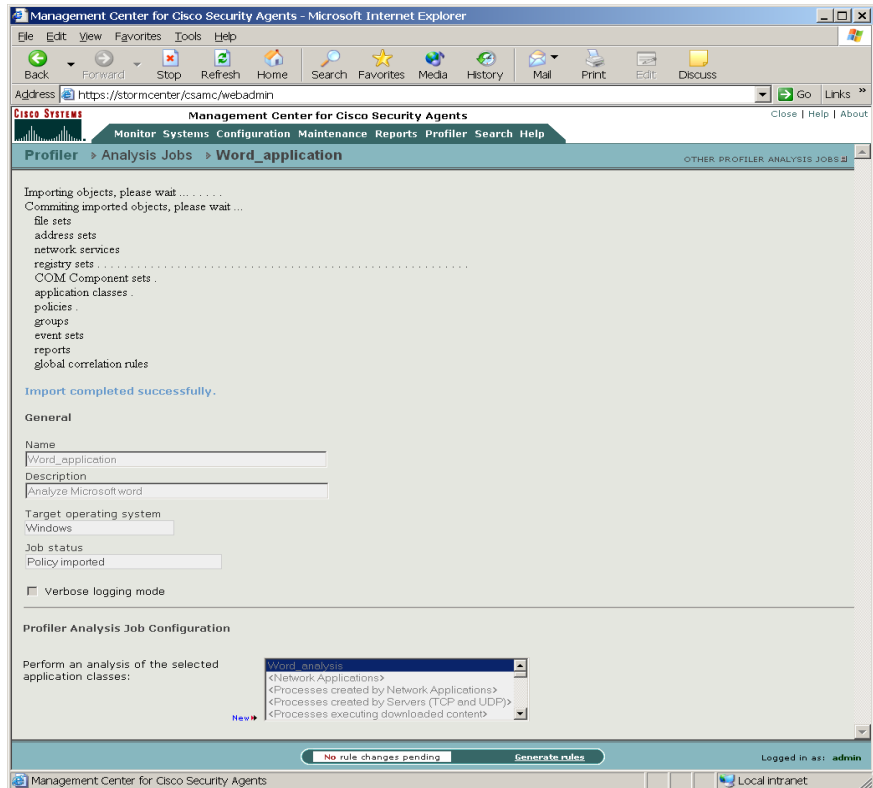
When Profiler has completed its analysis of the logging data, the policy it created is ready to be imported into CSA MC.

Import the policy by once again accessing the Analysis Job window for this particular job. Click the **Import** button in the bottom frame. (This button only appears when the policy is ready for importing.)

**Note**

The policy and its accompanying "variables" are imported into CSA MC. Profiler creates its own variables for use in the rules it also creates. See [Figure 11-5](#).

Figure 11-5 Import Process



The Profiler Policy

Once imported, the Profiler policy is added to your list of policies with the word "Job" appended to the original analysis job name. For example, if the analysis job name is "Word_application", the name of the policy would be "Job Word_application policy."

Figure 11-6 View the Policy

The screenshot shows the Management Center for Cisco Security Agents web interface in Microsoft Internet Explorer. The browser address bar shows <https://stormcenter/csamc/webadmin>. The page title is "Management Center for Cisco Security Agents". The navigation menu includes "Monitor Systems", "Configuration", "Maintenance", "Reports", "Profiler", and "Search Help". The current page is "Configuration > Policies".

The main content area displays a list of 39 items. The table below represents the data shown in the screenshot:

Name	Rules	Description	OS
<input type="checkbox"/> Common Security Module	8 rules	Base security policy module for all UNIX systems	UNIX
<input type="checkbox"/> Common Web Server Security Module	14 rules	Base web server request filter policy module for all UNIX systems	UNIX
<input type="checkbox"/> File System Lockdown Module	2 rules	Policy module to restrict ALL modifications of executable files	UNIX
<input type="checkbox"/> Insecure Management Module	4 rules	Policy module to manage UNIX systems in an insecure manner	UNIX
<input type="checkbox"/> Network Lockdown Module	2 rules	Policy module to restrict ALL network access	UNIX
<input type="checkbox"/> Network Quarantine Module	2 rules	Policy module to prevent ALL network access	UNIX
<input type="checkbox"/> Required System Module	10 rules	Policy module to allow critical UNIX functions	UNIX
<input type="checkbox"/> Restrictive Apache Module	11 rules	Restrictive policy module for UNIX Apache Web Server	UNIX
<input type="checkbox"/> Restrictive iPlanet Module	8 rules	Restrictive policy module for iPlanet Web Server	UNIX
<input type="checkbox"/> Restrictive SendMail Module	8 rules	Restrictive policy module for UNIX SendMail Server	UNIX
<input type="checkbox"/> Secure Management Module	8 rules	Policy module to manage UNIX systems in a secure manner	UNIX
<input type="checkbox"/> Server Module	3 rules	Base policy module for UNIX servers	UNIX
<input type="checkbox"/> Cisco VPN Client Module	6 rules	Policy module for Cisco VPN client	Windows
<input type="checkbox"/> CiscoWorks Common Security Module	15 rules	Base security policy module for all systems running CiscoWorks	Windows
<input type="checkbox"/> CiscoWorks Restrictive VMS Module	3 rules	Policy module for systems running only the VMS bundle	Windows
<input type="checkbox"/> CiscoWorks VMS Module	28 rules	Policy module for servers running CiscoWorks VMS product components	Windows
<input type="checkbox"/> Common Security Module	19 rules	Base security policy module for all Windows systems	Windows
<input type="checkbox"/> Common Web Server Security Module	15 rules	Base web server request filter policy module for all Windows systems	Windows
<input type="checkbox"/> Data Theft Prevention Module	10 rules	Policy module to prevent theft of sensitive data files	Windows
<input type="checkbox"/> Desktop Module	11 rules	Base policy module for desktops	Windows
<input type="checkbox"/> Distributed Firewall Module	10 rules	Policy module to restrict network services	Windows
<input type="checkbox"/> File Integrity Module	4 rules	Policy module to monitor access to key system files	Windows
<input type="checkbox"/> Inbound Port Blocking Module	4 rules	Policy module to block incoming connections	Windows
<input type="checkbox"/> Instant Messenger Module	7 rules	Policy module for Instant Messenger	Windows
<input type="checkbox"/> Job Word_application_policy	10 rules	Analysis job Word_application_policy	Windows
<input type="checkbox"/> Network Lockdown Module	2 rules	Policy module to restrict ALL network access	Windows
<input type="checkbox"/> Network Quarantine Module	2 rules	Policy module to prevent ALL network access	Windows
<input type="checkbox"/> Remote Laptop Module	12 rules	Base policy module for remote laptops	Windows
<input type="checkbox"/> Required Windows System Module	11 rules	Policy module to allow critical Windows functions	Windows
<input type="checkbox"/> Restrictive Apache Module	13 rules	Restrictive policy module for Windows Apache Web Server	Windows
<input type="checkbox"/> Restrictive DHCP Server Module	5 rules	Restrictive policy module for DHCP/BOOTP servers	Windows

At the bottom of the interface, there are buttons for "New", "Delete", "Clone", and "Compare". A status bar indicates "1 rule change pending" and "Generate rules". The user is logged in as "admin".

Reviewing the Policy

The policies created by Profiler enforce normal application behavior and maintain application and system integrity. To achieve this, the general strategy behind the creation of Profiler policies is to protect the application from the system and to protect the system from the application.

As with all new policies you create, you should review the policy generated by Profiler and run it in Test Mode for some period of time to ensure that it works as intended. You should also review the reports generated during the analysis as they are valuable resources for understanding the application as well as the policy.

**Note**

Profiler does not add system hardening or global correlation "built-in" rules to the policy. For example, you can add Trojan detection to the policy.

Profiler Policy Methodology

Protecting the application from the system

As part of the policy, Profiler creates file access control rules with the purpose of protecting the application data. These rules are left disabled by default as they restrict all other applications from accessing the analyzed application's data files. This is a fairly restrictive approach and, depending on the application itself, you may or may not want to enable these rules as part of the policy.

Protecting the system from the application

Resources accessed by the application are broken down into file, network, registry, and COM categories and then rules for each category are created by Profiler. Allow rules permit what was seen as normal application behavior while deny rules prevent access to all resources were not used by the application during the logging period.

Because policy requirements may vary from site to site, Profiler generates several rules that are disabled by default. The disabled rules are generally network and registry restrictions. Profiler creates these rules but keeps them disabled, leaving it up to the administrator to decide whether or not to impose these added restrictions. These rules are disabled by default because, generally, you should use the application-specific policies created by Profiler in combination with the Sample Network (Permissive, Selective, and Restrictive) policies shipped with the CSA MC.

If you decide to edit Profiler policies based on your site's requirements, the reports generated during the logging analysis process contain information on all the resources accessed by the application during the logging period. The "summary" reports generated for each resource type are particularly useful in helping to pinpoint what resources may require more or less restrictive rules. (See [Profiler Reports, page 11-18](#) for details.)

The general methodology behind the creation of rules for each resource type is as follows:

- File access control rules

Profiler creates file set variables that are combinations of file extension and directory pairs for accessed resources. These are used in allow file access control rules. Profiler then creates a deny file access control rule that prevents access to all other files and directories.

Use File Directory Summary and Individual File Summary reports to help refine these rules, if needed.

- COM component access control rules (Windows only)

Profiler creates COM component set variables which it then uses in a COM component access control rule to allow access to the required COM components. It then creates a COM component deny rule to deny all applications access to the COM components not used during the logging period.

Use COM Object Summary reports to help refine these rules as needed.

- Registry access control rules (Windows only)

Profiler creates these rule types but disables them by default. Registry access control rules are very powerful system control tools. Restricting access to a required registry key could produce undesired results.

Profiler creates Registry Set variables based on the registry resources accessed during the logging period. These registry variables are broken into those that should be allowed and those that can be denied. Those allowed are registry keys accessed during the logging period. All others fall in the deny range. This deny applies only to write access. All registry keys are still allowed read access. You can enable these rules, but you should understand the restrictions you are imposing.

Use Registry Key Summary reports to help refine these rules, if needed.

- Network access control rules

Profiler creates network access control rules but disables network deny rules by default. Network allow rules are created to allow network services for all addresses, both client and server, that were accessed during the logging period. The disabled deny rules then deny all services, client and server, on all ports for the analyzed application. These are fairly restrictive rules. If you intend to enable them or refine them (change port number restrictions or address information), you should refer to the Network Summary reports for information on network services used by the application.

Variable and Application Class Creation

When Profiler creates the rules for the policy, it also creates all the registry and COM component variables required by the rules. All Windows files are entered as literals. (Note that UNIX files are grouped into sets.)

Additionally, Profiler creates a new application class for the analyzed application and uses this new application class in all rules that make up the policy. You should note that if you select more than one application class for the analysis job, the application class created for the policy is an aggregate of all the analyzed applications.

If you decide that the application is not dangerous and it can run without any policy restrictions, you can begin to configure the analysis job.

Profiler Reports

During the analysis process, Profiler sorts the logging data it receives from the logging agent into categorized reports. You can view these reports on the CSA MC system by accessing the **Analysis Reports** item from the **Profiler** menu.



Note

The Analysis Reports menu item is *only* viewable on the CSA MC system and is not accessible remotely. You cannot view Profiler reports remotely.

Reports on specific jobs only become available once the job has successfully completed. The CSA MC Event Log displays a message to inform you that reports have been created.

Figure 11-7 Registry Events Report

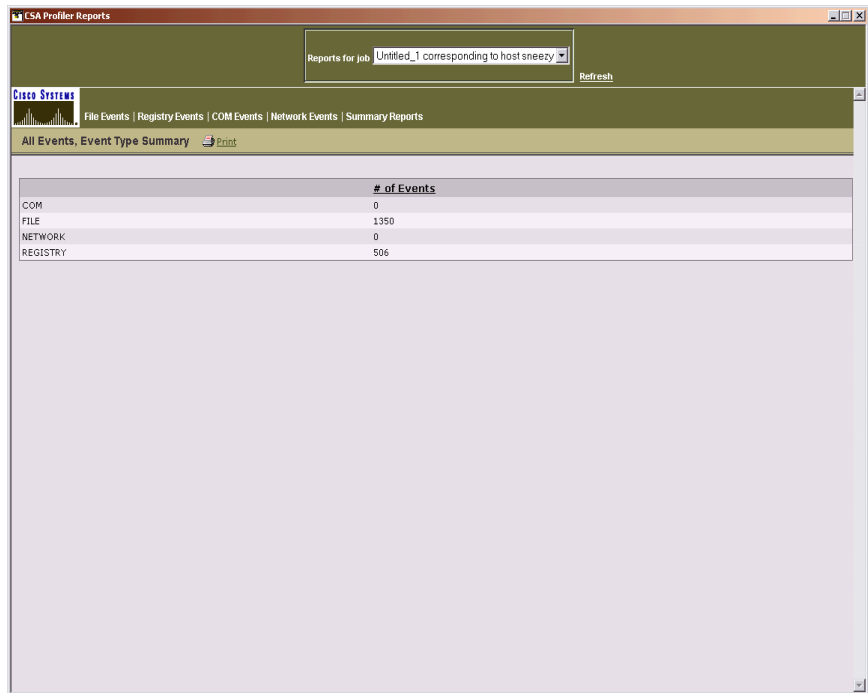
Time	Key Name	Value Name	PID	Process Name
Wed Apr 2 11:22:23 EST 2003	HKLM\Software\Microsoft\Windows\CurrentVersion\Explorer	null	3344	winzip32.exe
Wed Apr 2 11:22:23 EST 2003	HKU\S-1-5-21-1275210071-507921405-1060284298-1000\Software\Microsoft\Windows\CurrentVersion\Explorer\User Shell Folders	null	3344	winzip32.exe
Wed Apr 2 11:22:23 EST 2003	HKU\S-1-5-21-1275210071-507921405-1060284298-1000\Software\Microsoft\Windows\CurrentVersion\Explorer\Shell Folders	null	3344	winzip32.exe
Wed Apr 2 11:22:23 EST 2003	HKU\S-1-5-21-1275210071-507921405-1060284298-1000\SOFTWARE\MICROSOFT\Windows\CURRENTVERSION\Explorer\Shell Folders	Personal	3344	winzip32.exe
Wed Apr 2 11:22:23 EST 2003	HKLM\SOFTWARE\Microsoft\Cryptography\ RNG	null	3344	winzip32.exe
Wed Apr 2 11:22:23 EST 2003	HKLM\SOFTWARE\MICROSOFT\Cryptography\RNG	Seed	3344	winzip32.exe
Wed Apr 2 11:22:23 EST 2003	HKLM\SOFTWARE\MICROSOFT\Cryptography\RNG	Seed	3344	winzip32.exe
Wed Apr 2 11:22:23 EST 2003	HKLM\SOFTWARE\MICROSOFT\Cryptography\RNG	Seed	3344	winzip32.exe
Wed Apr 2 11:22:23 EST 2003	HKLM\SOFTWARE\MICROSOFT\Cryptography\RNG	Seed	3344	winzip32.exe
Wed Apr 2 11:22:23 EST 2003	HKLM\SOFTWARE\MICROSOFT\Cryptography\RNG	Seed	3344	winzip32.exe
Wed Apr 2 11:22:23 EST 2003	HKLM\SOFTWARE\MICROSOFT\Cryptography\RNG	Seed	3344	winzip32.exe
Wed Apr 2 11:22:23 EST 2003	HKU\S-1-5-21-1275210071-507921405-1060284298-1000\Software\Microsoft\Windows\CurrentVersion\Explorer\MountPoints\ D	null	3344	winzip32.exe
Wed Apr 2 11:22:23 EST 2003	HKU\S-1-5-21-1275210071-507921405-1060284298-1000\Software\Microsoft\Windows\CurrentVersion\Explorer\MountPoints\ D	null	3344	winzip32.exe
Wed Apr 2 11:22:23 EST 2003	HKU\S-1-5-21-1275210071-507921405-1060284298-1000\Software\Microsoft\Windows\CurrentVersion\Explorer\MountPoints\ C	null	3344	winzip32.exe
Wed Apr 2 11:22:23 EST 2003	HKU\S-1-5-21-1275210071-507921405-1060284298-1000\Software\Microsoft\Windows\CurrentVersion\Explorer\MountPoints\ C	null	3344	winzip32.exe

Report Components

When you access the Profiler reports window, you can view individual reports for all completed jobs from the same window by selecting a particular job from the **Reports for job** pulldown list at the top of the window.

Reports are broken down into the system and network resource types that were accessed by the application during the Profiler logging session. Each report category has several sub-topics you can select from for organizing information.

Each category drop-down menu provides an overall summary view. This view displays all the data of that particular category which was accessed during the analysis time frame. If you select to view **All Events** for a report category (see [Figure 11-8](#)), additional views further sort the information Profiler has collected by time frame, individual resource (e.g. single file or registry key), source and destination address in the case of network resources, and other criteria depending on the resource type in question.

Figure 11-8 All Events Report Sorting


The screenshot shows the 'All Events, Event Type Summary' report in the Cisco Security Agent Profiler Reports interface. The interface includes a header with the report title and a dropdown menu for job selection. Below the header, there are navigation tabs for 'File Events', 'Registry Events', 'COM Events', 'Network Events', and 'Summary Reports'. The main content area displays a table with the following data:

	# of Events
COM	0
FILE	1350
NETWORK	0
REGISTRY	506

Use the data from these reports to further refine your policies or to understand why particular rules were created for the policy.

You can view reports from the following categories:

File Event Reports

File reports display information such as the name of the file accessed, the application accessing the file, and the operation performed on the file. More specifically, they provide:

- Time: Useful for determining the time frame between events.
- Directory: This is the directory location (local or network share) of the file resource accessed in the event.
- File type: This is the individual file accessed in the event.

- Operation: This is the operation (read, write) performed on the accessed file.
- Process name: This is the application that accessed the resource.
- Number of events: This is the number of times the event in question occurred during the logging period.

Registry Event Reports (Windows only)

Registry reports provide details such as the name and value of the registry key that was accessed and the process that accessed it. More specifically, they provide:

- Time: Useful for determining the time frame between events.
- Key name: This is the name of the registry key accessed during the event.
- Value name: This is the registry value accessed during the event.
- PID: This is the process ID of the event. This is useful for distinguishing between different invocations of the same process.
- Process name: This is the application that accessed the resource.
- Number of events: This is the number of times the event in question occurred during the logging period.

COM Event Reports (Windows only)

COM reports display information on the COM Class ID that was accessed and the process that made the request. More specifically, they provide:

- Time: Useful for determining the time frame between events.
- Object name: This is the unique identifier for the COM object accessed during the event.
- PID: This is the process ID of the event. This is useful for distinguishing between different invocations of the same process.
- Process name: This is the application that accessed the resource.
- Number of events: This is the number of times the event in question occurred during the logging period.

Network Event Reports

Network reports display details such as the protocol accessing the network, the source and destination addresses of the connection, and the source and destination ports. More specifically, they provide:

- **Time:** Useful for determining the time frame between events.
- **Role:** This indicates whether the system in question was acting as a client or server during the network event.
- **Protocol:** This indicates whether this event was a TCP or UDP network connection.
- **Source address:** This is address where the connection originated from during the event.
- **Source port:** This is the port used during the event.
- **Destination address:** This is the destination address of the network connection for the event.
- **Destination port:** This is the destination port used for the connection. (Note that this port is used for the associated network rule that is generated as part of the policy.)
- **PID:** This is the process ID of the event. This is useful for distinguishing between different invocations of the same process.
- **Process name:** This is the application that accessed the resource.
- **Number of events:** This is the number of times the event in question occurred during the logging period.

Summary Reports

Summary reports display the number of times each resource type was accessed during the logging time frame.

Working with Reports

Profiler reports contain a great deal of application information. You can search through this data using the browser window's own search capabilities. From the report page you want to search on, press and hold the **Ctrl** button and press the **F** key. The browser search window appears.

You can also highlight, copy and paste report text into an application such as Microsoft Excel. From Excel, you can then organize the data in any manner you choose.



Policy Definition Guidelines

Overview

The policies you create on the Management Center for Cisco Security Agents should reflect a well-planned, enterprise-wide security policy. If your corporate enterprise already has security guidelines in place, the rules that form your policies should reflect those guidelines.

When you begin to configure policies, there is a common methodology you can use to successfully form the rules that will provide the security and the flexibility you require. This appendix provides a general approach you should take when creating your CSAMC policies.

You can use the Cisco Security Agent Profiler to automate the process of creating policies. Refer to [Chapter 11, “Using Cisco Security Agent Profiler”](#) for information.

This section contains the following topics.

- [Analyzing Applications, page 12-2](#)
- [Configuring Policies—The Methodology, page 12-3](#)

Analyzing Applications

The access control rules you create as part of your policies are application-centric. The application classes, those shipped with CSAMC and the ones you configure yourself, are the key to the rules you build as part of your security policies. Understanding how those applications work is necessary for configuring rules that adequately address the needs of a secure yet unobtrusive policy for that application.

There are three specific areas to consider when determining the type of security required by the application in question. There are overall *generic types of protection* that stop malicious code such as SYN flood protection, Network worm protection, Port scan detection (Network shield), and Trojan detection. There are *application-specific types of protection* you can put in place to allow the application to operate normally while insulating it from any undesired access. Then there are *environment-specific types of protection* that control access to the application in question and its data over various network channels. It is the latter two, application-specific and environment-specific protection requirements that this appendix concentrates on.

When analyzing an application for the purpose of writing a policy, consider the following questions.

- What resources does the application own (file, network, and registry resources)?
- Can the application access other resources?
- Can other applications access this application's resources?
- How is the application administered? (e.g. configuration tools used, accessed locally or remotely)
- Does the application interact with other applications as part of its normal operation?
- Does the application spawn processes and if so, what resources do those processes access?
- What application-based rules vs. environmental rules are necessary?

Determining the answers to these types of questions will help you target the resources you want to control as part of your policy for protecting the application.

For example, asking the questions above when analyzing how a Web server application operates would first lead you to determine which files are installed and used by the Web server application itself. What network resources are accessed and what registry keys are owned by the application? How is the Web server administered? Are html files FTP'ed to the server or is Front Page used locally on the system? These are questions targeted at producing application specific rules for a policy.

You would also note how the Web server is used and who can access it. Is it an intranet or Internet server? Does it act as a standalone server or does it access other resources? If there are forms users fill out on the Web server, does it use a backend SQL Server to store data? If so, which applications must be able to communicate with each other and what services, other than HTTP, are required for this communication? These are questions targeted at producing environment specific rules for a policy.

Ultimately, you want your policy to secure both the application and the environment it operates within.

Configuring Policies—The Methodology

Once you understand how an application works, you can begin forming a policy to protect it. There are five general areas you want to address for each resource you are protecting. By addressing the security needs of these five areas, you can configure a well-formed policy to protect the resources you are targeting.

When building a policy to protect a designated resource, refer to the following steps to help you address each resource area.

Step 1 Protect the application executables.

You must prevent writing to the application executables themselves. This maintains the integrity of the executable. The only time the executable should change is if you're upgrading the application.

This type of rule would prevent a Trojan from naming itself "Netscape.exe" to disguise itself as the real Netscape executable.

Step 2 Restrict the application processes.

Dictate what the applications in question can and cannot do. Likely, you'll want specific applications to write only to their own file types. To restrict an application, you must look at the files the application needs to read and write to and then restrict it to only those files.

This type of rule would prevent a buffer overrun from compromising a running application, and damaging other components on the system.

When applications are invoked, they often spawn other processes as part of the action they are performing. It may be desirable to place different restrictions on spawned processes. Therefore, when you analyze an application in preparation for writing rules, CSAMC gives you the option of including or excluding child processes created by the original application. You can also restrict the child processes of an application and create a rule to address only those processes.

Step 3 Protect application specific data.

Restrict access to specified data by other applications. For server policies, you'll want to protect information in certain directories on the server in question, allowing restricted access to specific files and block all outside access to other files.

In order to correctly formulate this rule, you must examine what other applications (if any) need to access the application data.

This type of rule would protect another application from retrieving sensitive data from a server, such as credit card information or a password file.

Step 4 Permit network access as required.

If an application requires network connectivity, you should specify what required network services must be enabled. Components that are “network visible” are especially vulnerable to attacks. It is important to control what these network-accessible applications (and their spawned processes) can do.

Step 5 Protect the application registry keys.

Restrict access to sensitive application-specific registry keys. You want to allow the specific application to write to its own registry keys, but prevent all other applications from writing to those registry keys.

Depending on the application you are writing a policy for, you may skip one or more of these steps. For example, if there are no registry settings requiring protection, you can skip step 5.

General Server Policy

The General-Server policy described here uses the applicable steps mentioned previously to secure common server resources. This is a generic server policy that can be applied to any server. Depending on the type of server you’re protecting, you’ll want to apply this General-Server policy and then create an additional policy, which more specifically targets the resources you want protected, to augment this general one. Meanwhile, here is an overview of a General-Server policy.

Table 12-1 *General Server Policy*

Rule Type	Description
File access control	Allow, all applications read system dll’s
Network access control	Deny, lockdown network access client
Network access control	Deny, lockdown network access server
File access control	Deny, protect system executables
Network shield	Detect network port scans, detect and protect against network SYN flood attacks
Trojan detection	Detect and terminate potential application Trojans

Note that the rules in these tables are ordered (top to bottom) according to their priority. High priority deny rules take precedence over all others. Allow rules take precedence over deny rules. This General-Server policy now locks down the server machine protecting the system directory and protecting network access.

Sample Web Server Policy

Once you have a general server policy to protect basic server resources, you can write a policy that actually targets the resources used by the particular server application you want to protect. For the purposes of this example, the application is a Web server. The executable is “WEB.EXE.”

This targeted server policy builds on the General-Server policy restrictions, allowing the services required for WEB.EXE to operate securely. Once we explain the components of this policy, we will combine both the General-Server and Sample Web Server policies and implement them together to provide the overall protection the Web server application requires.

Table 12-2 Sample Web Server Policy

Rule Type	Description
File access control	High Priority Deny, protect Web server data
File access control	Allow, let WEB.EXE write to temp files and log files
Network access control	Allow, let WEB.EXE talk to network
File access control	Query user, protect Web server directories from others
Registry access control	Deny, protect sensitive Web server keys
File access control	Deny, prevent WEB.EXE all file write access

Here is how the methodology detailed in the first section of this document was applied to the creation of this policy. The Description, appearing in italics below, given for each rule in the Web server policy table is listed here with the “methodology” step that applies to it.

Step 1 Protect the application executables.

Protect Web server directories from others: Here we have denied all applications from writing to the directories that contain the Web server application executables.

Step 2 Restrict the application processes.

For a general purpose policy, you want to protect the system from the application in question. Therefore, you can allow the application (ex. WEB.EXE) to read all system files, but restrict writes to system files. (If you are concerned about the application reading certain system files, you can restrict reads to those files specifically, if necessary.)

Prevent WEB.EXE all file write access: This rule denies the Web server application access to all files on the system.

Let WEB.EXE write to temp files and log files: This rule allows the Web server application to write to temp and log files used by the application.

Note that restricting access to a resource should always be done in the policy that owns that resource.

Step 3 Protect application specific data.

Protect Web server data: This rule prevents anyone from writing to html files and defacing web pages.

Step 4 Permit network access as required.

Let the WEB.EXE talk to network: This allows WEB.EXE to act as a server for the http service.

Step 5 Protect the application registry keys.

Protect sensitive Web server keys: This would protect, for example, keys controlling user authentication settings.

Combined General Server and Sample Web Server Policies

To fully protect the Web server, we apply our base General-Server policy and our targeted Sample Web Server policy to the agent running on the Web server system. When applied to the Web server, the combined policies work as displayed in the table below (in order of rule precedence).

Table 12-3 Combined Policies

Rule Type	Description
File access control	High Priority Deny, protect Web server data
File access control	Allow, let WEB.EXE write to temp files and log files
File access control	Allow, all applications read system dll's
Network access control	Allow, let WEB.EXE talk to network
File access control	Query user, protect Web server directories from others
Network access control	Deny, lockdown network access client
Network access control	Deny, lockdown network access server
Registry access control	Deny, protect sensitive Web server keys
File access control	Deny, prevent WEB.EXE all file write access
File access control	Deny, protect system executables
Network shield	Detect network port scans, detect and protect against network SYN flood attacks
Trojan detection	Detect and terminate potential application Trojans

Reference

“Vulnerable applications” defined in various rules are network-aware applications. These application types are much more vulnerable than others. They are as follows:

- TCP and UDP servers and processes created by them are vulnerable because they are susceptible to buffer overflow attacks.
- Processes that read downloaded content are vulnerable because they may be interpreting and taking action based on downloaded data.
- Remote clients are applications running on another machine and are therefore vulnerable because CSA does not know what these applications are when they attempt to access resources.
- Removable media, in some cases, is categorized as vulnerable. This includes media accessed from CD-ROM, floppy, USB drives, or any other peripheral device.



Third Party Product Integration

Overview

The Management Center for Cisco Security Agents provides integration with other third party products. This section provides information on supported third party integration applications.

In most cases, you are referred to the third party documentation for configuration information.

This section contains the following topics.

- [Cisco VPN Client Support, page 13-1](#)
- [Cisco Security Monitor Integration Support, page 13-2](#)
- [netForensics Integration Support, page 13-2](#)
- [Check Point™ OPSEC™ Integration, page 13-2](#)

Cisco VPN Client Support

The Cisco Security Agent is a supported configuration for the "Are You There?" feature of the Cisco VPN Client, Release 4.0. For configuration details, please refer to Chapter 1 of the Cisco VPN Client Administrator Guide, in the section entitled "Configuring VPN Client Firewall Policy -- Windows Only."

Cisco Security Monitor Integration Support

Cisco Security Monitor is a Security Information Management application that can receive security events from multiple devices. Security Monitor presents the information in a real-time, web-based console so that these events can be managed across the network. Security Monitor also provides event notification, event reporting, and event correlation.

To integrate events generated by the Cisco Security Agent with the Security Monitor application, refer to Chapter 3 of your Security Monitor documentation, “Configuring Devices to Monitor.”

netForensics Integration Support

netForensics is a Security Information Management application that can receive security events from multiple devices. This gives the administrator the convenience of having a single point from which to manage events from heterogeneous sources. netForensics presents the information in a real-time, web-based console so that these events can be managed across the network.

To integrate events generated by the Cisco Security Agent with the netForensics application, refer to your netForensics documentation.

Check Point™ OPSEC™ Integration

The Check Point™ OPSEC™ (Open Platform for Security) provides a set of API's (Application Programming Interfaces) which allow integration of various network security components. The SCV (Secure Configuration Verification) API provides a mechanism by which the configuration of a machine running the VPN-1® SecureClient™ can be verified.

With its Cisco Security Agent product, Cisco provides an “SCV Check” which can be used to verify that the agent is running on machines connecting via the VPN-1 SecureClient. With such a configuration, machines which fail the “SCV Check” are not allowed to establish connections through the Firewall.

Configuration Prerequisites

The following components are required to integrate the Cisco Security Agent as an SCV check within the OPSEC framework:

- On Machine A, an installation of Management Center for Cisco Security Agents, version 4.0 or greater.
- On Machine B, an installation of the Check Point VPN-1 & Firewall-1®, along with the Management Client and Policy Server, all of which are components of Check Point NG FP1 (Next Generation Feature Pack 1). The Firewall should be configured for VPN-1 SecureClient use.
- On Machine C, an installation of the Check Point VPN-1 SecureClient which points to the Firewall on Machine B. Also on Machine C, an installation of the Cisco Security Agent, installed from the Management Center for Cisco Security Agents on Machine A. (See the *Caution* below.)



Caution

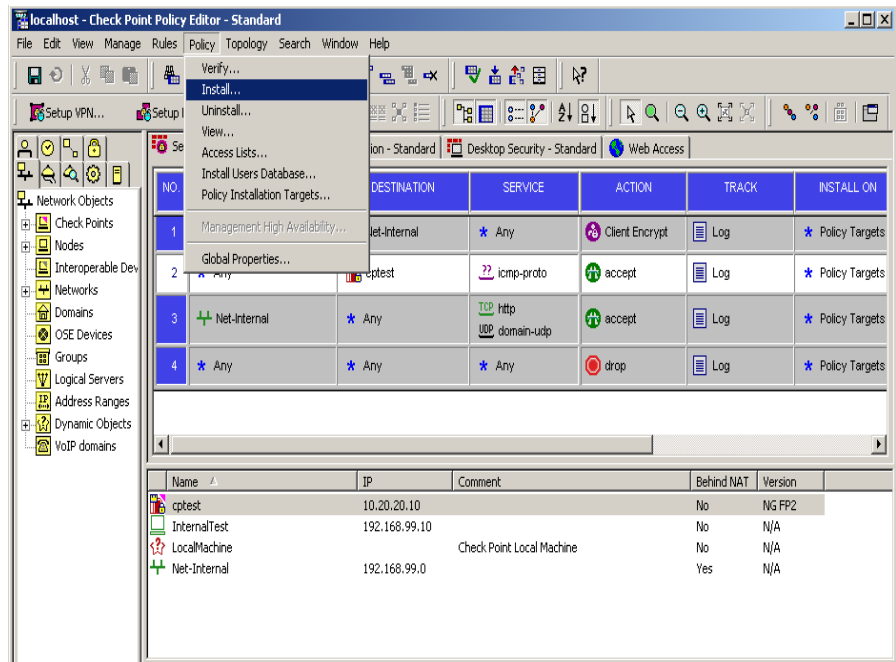
On Machine C, it is important that you install the SecureClient software before you install the Cisco Security Agent.

Integration Configuration

This section provides the procedure for deploying the SCV check. The following instructions assume the existence of (and refer to) the prerequisites described in the previous section. These instructions also refer to a file called `LOCAL.SCV`, which is accessible from the self-extracting executable located `ThirdParty\OpSec\SCV.exe` on the CSA MC product CD.

-
- Step 1** On Machine B, copy the `LOCAL.SCV` file from the CD to the `\winnt\fw1\ng\conf` directory. Note that any pre-existing versions of `LOCAL.SCV` should be renamed so that they are not overwritten.
- Step 2** Using the Check Point™ Policy Editor, perform a Policy->Install onto Machine C and on to any other SecureClient machines for which the SCV check “CSAgent” is to be enforced. (Configuration for enforcing SCV checks varies across Check Point™ Feature Packs. Please refer to the “Desktop Security Guide” for VPN-1 and SecureClient configuration details.)

Figure 13-1 Check Point Policy Editor



- Step 3** On Machine C (and other relevant SecureClient machines), the new policy will automatically be downloaded. With the SCV check now enforced, only machines with an installed (and running) Cisco Security Agents are allowed to establish connections through the Firewall. Otherwise, the user receives a message box stating “Cisco Security Agent SCV Check Failed.”



Note No configuration is required on the client side. The Cisco Security Agent installation automatically installs and registers the relevant files.



Cisco Security Agent Overview

Overview

This chapter describes the agent and provides information on the agent user interface. There is no configuration necessary on the part of the end user in order to run the agent software. Optionally, as the administrator, you can ask users to enter individualized contact information into the fields provided. If required, the agent user interface makes it easy for the user to enter this data and send it to CSA MC.

If you have configured Query User rules for file access, users should know how to respond to query pop-up boxes. This information is included in the HTML help provided with the agent user interface. You may want to refer end users to this agent help.

This section contains the following topics.

- [Downloading and Installing, page A-2](#)
- [The Agent User Interface, page A-7](#)
- [Responding to Pop-up Query Boxes, page A-12](#)
- [Suspend Agent Security, page A-13](#)
- [Installing Software Updates on Agents, page A-14](#)
- [Installing the UNIX Agent, page A-16](#)
- [UNIX Agent csactl Utility, page A-18](#)

Downloading and Installing

Once you build an agent kit on CSA MC, you deliver the generated URL, via email for example, to end users so that they can download and install the Cisco Security Agent. They access the URL to download and then install the kit. This is the recommended method of agent kit distribution.

But you may also point users to a URL for the CiscoWorks system. This URL will allow them to see all kits that are available. That URL is:

```
http://<ciscoworks system name>/csamc/kits
```

If you are pointing users to the “kits” URL and you have multiple agent kits listed here, be sure to tell users which kits to download.

End users must have administrator privileges on their systems to install the agent. Systems to which agents are installed must meet the following requirements:

Table A-1 Agent Requirements (Windows)

System Component	Requirement
Processor	Intel Pentium 200 MHz or higher Note Uni-processor and dual processor systems are supported
Operating Systems	<ul style="list-style-type: none"> Windows XP (Professional English 128 bit) Service Pack 0 or 1 Windows 2000 (Professional, Server or Advanced Server) with Service Pack 0, 1, 2, or 3 Windows NT (Workstation, Server or Enterprise Server) with Service Pack 5 or higher Note Terminal Services are supported on Windows XP and Windows 2000 (Terminal Services are not supported on Windows NT.)
Memory	128 minimum—all supported Windows platforms

Table A-1 Agent Requirements (Windows) (continued)

System Component	Requirement
Hard Drive Space	15 MB or higher Note This included program and data.
Network	Ethernet or Dial up Note Maximum of 64 IP addresses supported on a system.

**Note**

The Cisco Security Agent uses approximately 20 MB of memory. This applies to agents running on all supported Microsoft platforms.

To run the Cisco Security Agent on your Solaris server systems, the requirements are as follows:

Table A-2 Agent Requirements (UNIX)

System Component	Requirement
Processor	UltraSPARC 500 MHz or higher Note Uni-processor and dual processor systems are supported
Operating Systems	Solaris 8, 64 bit Note If you have the minimal Sun Solaris 8 installation (Core group) on the system to which you are installing the agent, the Solaris machine will be missing certain libraries and utilities the agent requires. Before you install the agent, you must install the "SUNWlibCx" library which can be found on the Solaris 8 Software disc (1 of 2) in the /Solaris_8/Product directory. Install using the pkgadd -d . SUNWlibCx command.
Memory	256 minimum

Table A-2 Agent Requirements (UNIX) (continued)

System Component	Requirement
Hard Drive Space	15 MB or higher Note This included program and data.
Network	Ethernet Note Maximum of 64 IP addresses supported on a system.

**Note**

Agents systems must be able to communicate with CSA MC over HTTPS.

**Caution**

When upgrading or changing operating systems, uninstall the agent first. When the new operating system is in place, you can install a new agent kit. Because the agent installation examines the operating system at install time and copies components accordingly, existing agent components may not be compatible with operating system changes.

Network Shim Optional

In some circumstances, you may not want users to enable the network shim on their systems as part of the agent installation. For example, if users have VPN software or a personal firewall installed on their systems, the network shim's Portscan detection, SYN flood protection, and malformed packet detection capabilities may be in conflict with VPNs and personal firewalls. (There are no conflicts with the Cisco VPN client, Release 4.0.)

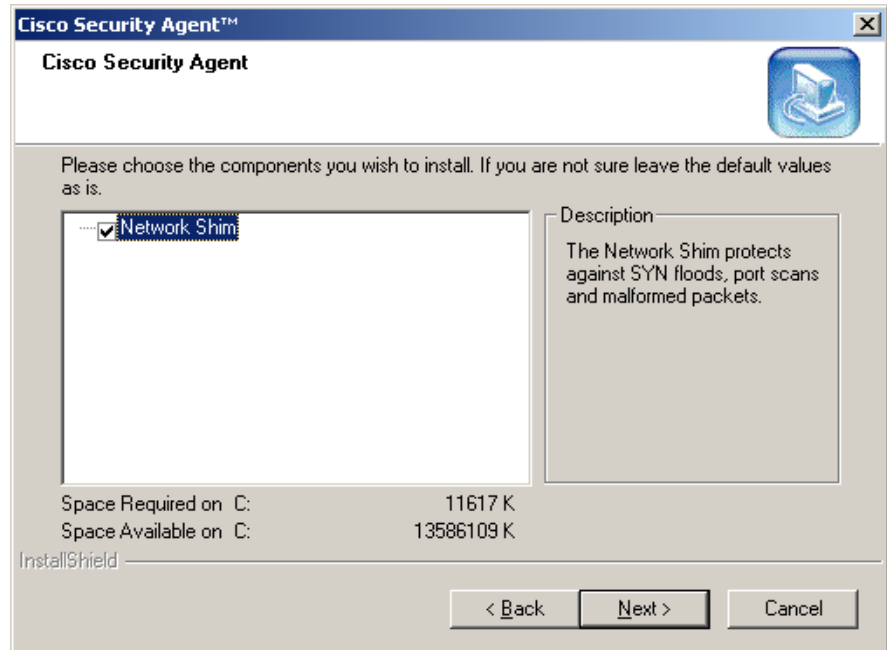
If you check the Quiet install checkbox when you make kits, you can also select whether the network shim is installed as part of the quiet install process.

To allow users to select whether or not to install the network shim themselves, you would create kits as not-quiet installations. (Do not select the Quiet install checkbox.) This way, users are prompted to enable the network shim during the agent installation. See [Figure A-1](#).

**Note**

Not enabling the network shim does not mean that Network Access Control rules won't work. It only means that the system hardening features mentioned in the previous paragraph are not enabled.

Figure A-1 Optional Network Shim



Once users install agents on their systems, they can optionally perform a reboot (if Automatic reboot is not selected at kit creation time). See [Figure A-2](#). Whether a system is rebooted or not, the agent service starts immediately and the system is protected.

Figure A-2 *Optional Agent Reboot*

If a system is not rebooted following the agent installation, the following functionality is not immediately available. (This functionality becomes available the next time the system is rebooted.)

Windows agents

- Network Shield rules are not applied until the system is rebooted.
- Buffer overflow protection (located on the Trojan page for Windows) is only enforced for new processes.
- Data access control rules are not applied until the web server service is restarted.
- COM component access control rules are not applied until the system is rebooted.

UNIX agents, when no reboot occurs after install, the following caveats exist:

- Buffer overflow protection is only enforced for new processes
- Network access control rules only apply to new socket connections

- File access control rules only apply to newly opened files.
- Data access control rules are not applied until the web server service is restarted.

At this time, the agent automatically and transparently registers with CSA MC.

You can see which hosts have successfully registered by clicking the **Hosts** link available from the **Systems** category in the menu bar. This displays the hosts list view. All registered host system names appear here. Agents are now ready to receive policies.

**Note**

By default, agents poll in to the management server every 10 minutes for policy updates (unless you change this value in the Groups configuration view. See [Chapter 3, “Configuring Groups and Managing Hosts”](#) for details.)

The Agent User Interface

**Note**

The Cisco Security Agent user interface does not run on UNIX systems. The UNIX agent has a utility (csactl) to provide capabilities that the Windows agent provides in its user interface. See [for details](#).

**Note**

If **No user interaction** (available on Windows groups only) is enabled for the system group, no agent UI appears on the end user system. See the [“Configuring Groups” section on page 3-3](#) for details.

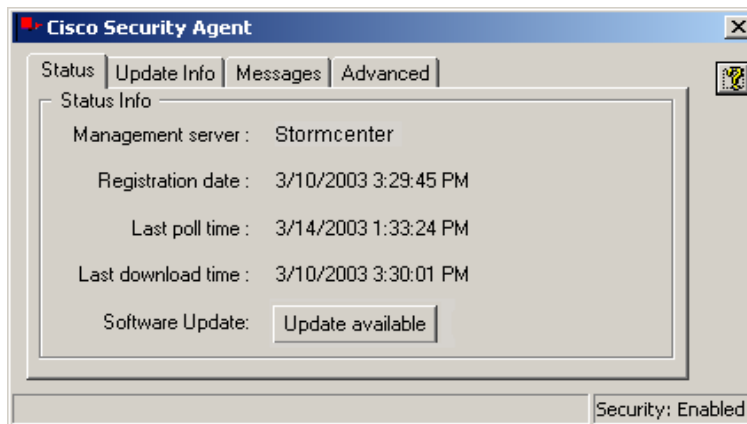
To open the agent user interface, users can double-click on the agent icon in their system trays. The user interface opens on their desktop. It contains four tabs. Most fields are read-only.

Status tab: This tab provides the following information (see [Figure A-3](#)).

- The name of the CSA MC with which this agent is registered.
- The date and time the agent registered with CSA MC.
- The date and time when the agent last polled in to CSA MC (data is not downloaded each time the agent polls).

- The date and time the agent last downloaded data from CSA MC.
- Lets users know if there is a software version update available for their agent (see [Installing Software Updates on Agents, page A-14](#)).

Figure A-3 Agent Status Tab



Note

Viewable from all tabs, on the bottom of the agent UI, is a Security status field. This status lets the user know if security is enabled or disabled. They can change this status (if allowed by the administrator) from the agent pulldown menu. See [Suspend Agent Security, page A-13](#).

Update Info tab: This tab provides the following information (see [Figure A-4](#)).

- Contact information, including user name, telephone number, location, and email address. Users enter this information here and click the Update button. CSA MC receives this contact data and you can now quickly locate a user if the agent indicates that there is a problem.

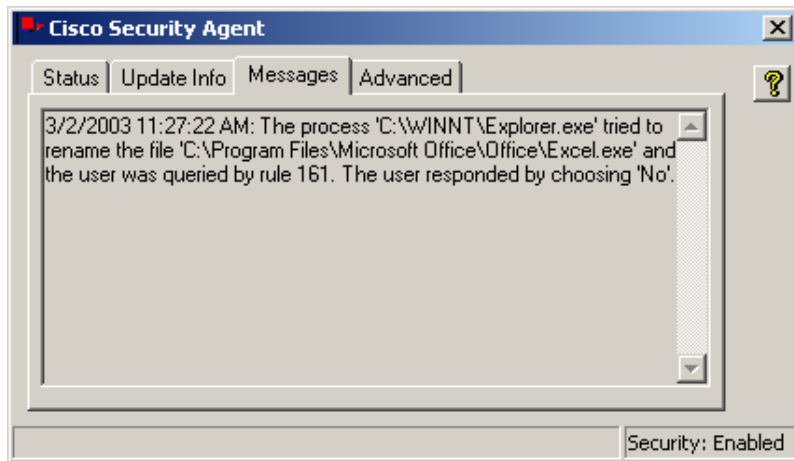
Figure A-4 Agent Update Info Tab

The screenshot shows a window titled "Cisco Security Agent" with four tabs: "Status", "Update Info", "Messages", and "Advanced". The "Update Info" tab is active. Below the tabs is a "Contact Info" section with five text input fields: "First Name" (Jo), "Last Name" (Smith), "Telephone" (555-0125 x416), "Location" (Waltham), and "Email Address" (jsmith@example.com). An "Update" button is located to the right of the "Email Address" field. At the bottom right of the window, there is a "Security: Enabled" indicator.

Messages tab: This tab provides the following information (see [Figure A-5](#)).

- When an agent denies a system action, a message informing the user of this event is placed in the Messages field. Note that a line of text in the Status tab informs the user that there are messages present. Click the Messages tab to view them.

Figure A-5 Agent Messages Tab



Events are also stored in the NT event log on the agent system.

**Note**

When a policy is triggered on an agent system and a message appears in the Messages tab, the flag icon in the system tray *waves*. This waving continues until the user opens the agent GUI and clicks on the Messages tab.

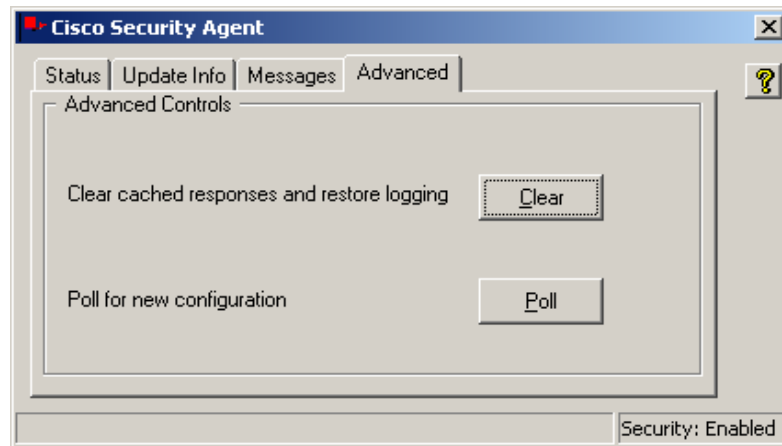
Advanced tab: This tab provides the following information (see [Figure A-6](#)).

When the agent logs an event with CSA MC, it remembers that event for an hour and does not log it again (even if the event occurs again) until that hour time frame has expired. This is to prevent the logfile from filling too quickly. The same applies to Query User pop-up messages (see [Responding to Pop-up Query Boxes, page A-12](#)). Once the user has answered a pop-up query, the system remembers the answer and responds automatically, not prompting the user with another query pop-up.

The Advanced tab on the agent lets the user clear the cache and re-enable logging.

- Clicking the **Clear** button tells the system to clear all cached responses and display a Query User pop-up box when the event in question occurs again. Clicking Clear also tells the system to clear its memory of all logged events (causing events to once again produce log messages if they occur).
- Clicking the **Poll** button forces the agent to poll the management server. This way, the agent receives any rule changes immediately. It will stop fast polling after the first successful configuration request. If the first attempt is unsuccessful, the agent will attempt to poll 2 more times. This is useful if new rules are being deployed and tested.

Figure A-6 Agent Advanced Tab



Responding to Pop-up Query Boxes

You can create rules that prompt users to either allow or deny the action when an attempt is made to access protected resources. If the rule in question is triggered, a pop-up box appears and requires the user to select one of the following responses.

**Note**

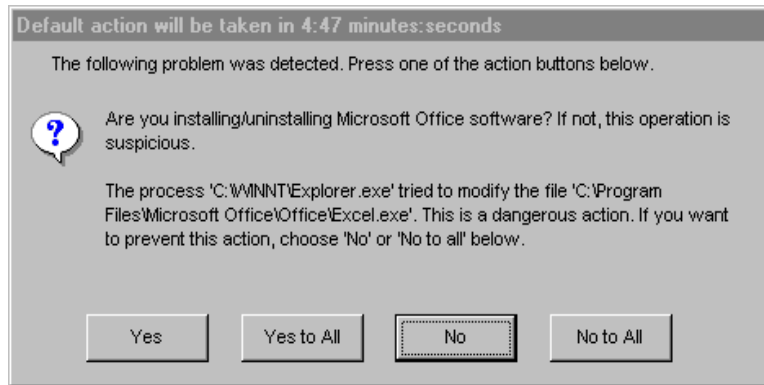
Query user options are not available for UNIX agents.

- Yes: Allows the application access to the resource in question.
- Yes to all: Allows the application access to all related query user protected resources, with no further queries appearing.
- No: Denies the application access to the resource in question.
- No to all: Denies the application access to all related query user protected resources, with no further queries appearing.

If the user does not make a selection within 5 minutes, the action is either allowed or denied by default depending on what you have configured for the rule. See the [“Querying the User” section on page 4-10](#) for more details.

**Note**

When users move their mouse over a button in a pop-up query box, a "tool tips" text box appears over the button explaining what the button does. Reading these tool tips can help if users are unsure how to respond to a query.

Figure A-7 Query User Pop-up Box**Note**

The HTML help provided with agent explains this pop-up dialog box to the end user. You may want to refer them to it.

Suspend Agent Security

Provided there is not an Agent service control rule (See Chapter 4 for rule details) that denies this action, all users can Suspend the security the agent provides on a Windows host by accessing the agent UI and clicking on the flag in the menu bar. A pulldown menu appears and users can select to **Suspend Security** and then **Resume Security** from there.

**Note**

If there is no agent UI on a system (no user interaction feature), the ability to suspend security is not available to users.

Provided there is not an Agent service control rule that denies this action, Windows administrators can run the following commands from a command prompt window on the agent host system to stop and start the agent service:

```
net stop "Cisco Security Agent"  
net start "Cisco Security Agent"
```

Provided there is not an Agent service control rule that denies this action, administrators can stop and start the agent service on a UNIX host by running the following commands from a command prompt window on the agent host system:

```
/etc/init.d/csa stop  
/etc/init.d/csa start
```

**Caution**

Suspending agent security and/or stopping the agent service on any system disables all rules on that system. Starting the agent service and resuming security reinstates all rules.

Installing Software Updates on Agents

Cisco occasionally provides software updates for Cisco Security Agents. You configure CSA MC to distribute the appropriate software updates to specified agents across the network. When agents poll in to check for new rules, if there is an update available for the agent in question, it receives the update at that time. See the [“Distributing Software Updates” section on page 3-26](#) for CSA MC configuration details.

There are two types of software updates: automatic and not-automatic.

- **Automatic updates:** For automatic updates, when agents poll in to the server, they transparently receive and install the software update. They are not required to do anything. It all occurs automatically in much the same way rules are updated without user participation. See [page 3-26](#) for automatic update configuration details.
- **Not-automatic updates:** When a software update installation is not automatic, agent systems receive a prompt (see [Figure A-8](#)) to either Update now (install) or Postpone the installation from 1-10 days at which time they will receive another prompt to update. If users select to postpone updating, an "Update available" button appears on the first tab of the agent GUI (see [Figure A-3](#)). Users can click this button to install the update when they are ready.

**Note**

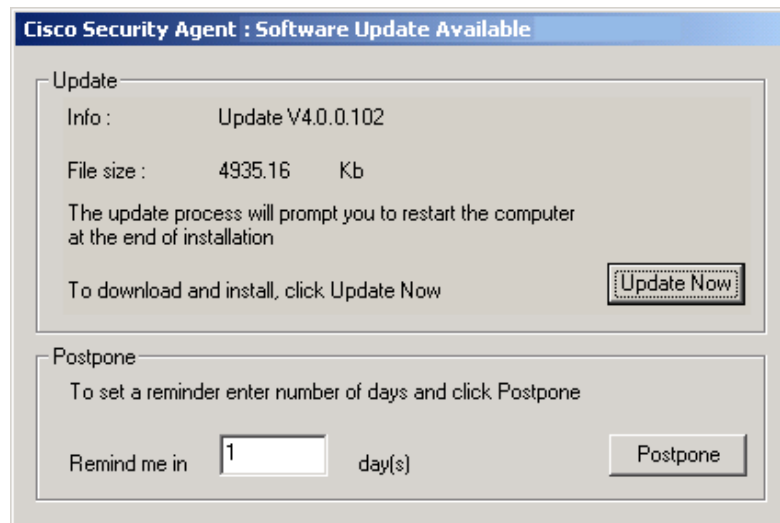
UNIX agents receive no prompt when software updates are available. Use the `csactl` utility (see [page A-18](#)) on UNIX systems to check for updates and install them.

If the Update time frame (a configurable parameter in the Pending software updates page) is narrow and agents postpone installing the update, they cannot try again until the same time the following day.

In some cases, agent systems must be rebooted after installing a patch. They are prompted if this is necessary.

Agent systems contain online help explaining how to install software updates. You may want to refer users to it.

Figure A-8 Agent Software Update Popup Window



Installing the UNIX Agent

This section details the commands you enter and the subsequent output that is displayed when you install the Cisco Security Agent on Solaris systems. When you download the agent kit from CSA MC, do the following to unpack and install it. (Note that you can put the downloaded tar file in any temp directory. Do not put it in the opt directory, for example, as you may then experience problems with the installation.)

Step 1 You must be super user on the system to install the agent package.

```
$ su
```

Step 2 Untar the agent kit.

```
# tar xf CSA-Server_4.0.0.15-setup.tar
```

Step 3 Install the agent package. (Use the command listed below when you install. This command forces the installation to use a package administration file to check the system for the required OS software agent dependencies. If the required dependencies are not present, such as the "SUNWlibC" library, the install aborts.)

```
# pkgadd -a CSCOcsa/reloc/cfg/admin -d .
```

```
[Output:]
```

```
The following packages are available:
```

```
1 CSCOcsa CSAagent
   (sun4u) 4.0.0.15
```

Step 4 Select the correct package or press enter to unpack all current packages.

```
Select package(s) you wish to process (or 'all' to process all
packages). (default: all) [?,??,q]:
```

```
[Output:]
```

```
Processing package instance <CSCOcsa> from </space/user>
```

The install now displays the Cisco copyright and prompts you to continue the installation.

Step 5 Answer yes (y) to continue the installation.

```
This package contains scripts which will be executed with
super-user permission during the process of installing this
package.
Do you want to continue with the installation of <CSCOcsa>
[y,n,?]y
```

Step 6 Installing CSAgent as <CSCOcsa>

Step 7 The installation continues to copy and install files. When the install is complete, the following is displayed:

```
The agent installed cleanly, but has not yet been started. The
command: /etc/init.d/csamanager start
will start the agent. The agent will also start automatically
upon reboot. A reboot is recommended to ensure complete system
protection.
The following packages are available:
  1 CSCOcsa CSAgent
    (sun4u) 4.0.0.15
```

Step 8 Quit (q) when installation is finished.

```
Select package(s) you wish to process (or 'all' to process all
packages). (default: all) [?,??,q]: q
```

Step 9 Optionally, reboot the system by entering the following.

```
# shutdown -y -i6 -g0
```

**Caution**

If a UNIX system is not rebooted following the agent installation, the following functionality is not immediately available: Buffer overflow protection is only enforced for new processes, network access control rules only apply to new socket connections, and file access control rules only apply to newly opened files.(This functionality becomes available the next time the system is rebooted.)

The agent installs into the following directory:

```
/opt/CSCOcsa
```

**Caution**

If you are upgrading the UNIX agent and you encounter the following error, "There is already an instance of the package and you cannot install due to administrator rules", you must edit the file `/var/sadm/install/admin/default`. Change "instance=unique" to "instance=overwrite" and then proceed with the upgrade.

Uninstall UNIX Agent

To uninstall the agent, enter the following command:

```
# pkgrm CSCOcsa
```

**Note**

If an agent is running a policy which contains an Agent service control rule, the agent cannot be uninstalled unless this rule is disabled. (Administrators can generally do this through a remote management session if the default policies applied to the CSA MC system are not changed to restrict this access.) See the [“Agent Service Control” section on page 4-24](#) for details on this rule type.

UNIX Agent csactl Utility

Because the UNIX Cisco Security Agent has no user interface, a utility is provided which allows you to check the UNIX agent status, poll in to CSA MC and re-enable logging. The command you enter on your UNIX system to perform these functions is **csactl**. Enter the csactl command as follows:

```
# /opt/CSCOcsa/bin/csactl <command>
```

Available commands are:

poll	Triggers an immediate poll of the management server. (Also lets you know if there is a software update available.)
resetlog	Resets the logging holdback -- allows all log messages.

status	Displays a small amount of status information. (Also lets you know if there is a software update available.)
swupdate	Updates agent software.
info <text>	This is a mechanism for directly sending custom (informational) textual events to CSA MC. Once the message reaches the CSA MC, it can be viewed or a notification can be sent to an administrator.
warning <text>	This is a mechanism for directly sending custom (warning) textual events to CSA MC. Once the message reaches CSA MC, it can be viewed or a notification can be sent to an administrator.
alert <text>	This is a mechanism for directly sending custom (alert) textual events to CSA MC. Once the message reaches CSA MC, it can be viewed or a notification can be sent to an administrator.
about	Displays agent software version number.

The commands listed above are only available to root.

For example, poll in to CSA MC by entering the following:

```
# /opt/CSCOcsa/bin/csactl poll
Poll of management center succeeded
```

For example, check the status of the agent by entering the following:

```
# /opt/CSCOcsa/bin/csactl status
Status:
Management center: stormcenter
Registration time: 2002-03-20 15:19:16
Host id: {FG9DA858-6131-45E9-18BD-EE32BA2D0676}
Last download time: 2002-03-20 15:19:23
Last poll time: 2002-03-20 15:20:42
Software update: newer version is available
```

For example, to perform a software update:

```
# /opt/CSCOcsa/bin/csactl swupdate
```

**Note**

You must reboot the system after performing a software update.

For example, re-enable logging if duplicate messages are being throttled:

```
# /opt/CSCOcsa/bin/csactl resetlog
Reset Log throttle sent to kernel
```



System Components

Overview

This appendix contains information on CSA MC and agent core components, explaining how these components relate to each other, including details on various CSA MC and agent services.

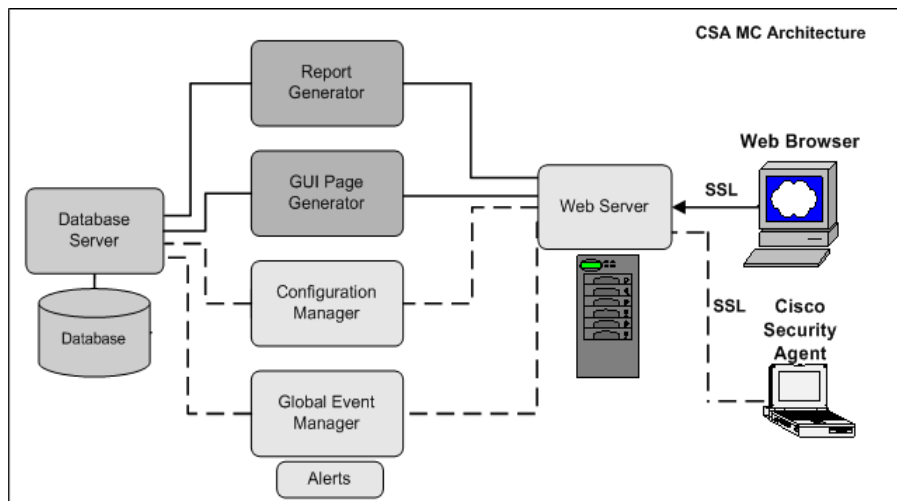
This section contains the following topics.

- [CSA MC Components, page B-2](#)
- [Agent Components, page B-3](#)

CSA MC Components

CSA MC architecture is displayed in this appendix. Note that although the agent is mentioned often here, it is only in terms of CSA MC's relation to the agent. Agent software does have its own system components which are described in this chapter. It is CSA MC that pushes security policies to the agents and coordinates the events it receives back from the agents. The mechanisms that are required to perform those tasks are described here as part of the CSA MC architecture.

Figure B-1 CSA MC Components



The **web browser**, shown on the right in the diagram, represents any web browser on any system across an enterprise from which administrators can securely access the CSA MC web-based interface. Communications between the web browser and the web server occur over SSL, allowing administrators to securely access the database of rule configurations from any location.

The *web server* provides the means of communication between the web browser and all other CSA MC system components. The web server displays reporting information, configuration version data, and event logging data.

It is through the **web server** that the agents installed on systems across an enterprise can exchange data with the CSA MC **configuration manager** and the **global event manager**. When agents poll in to CSA MC for rule set updates, it is

the configuration manager that pulls the rules from the database and distributes them to the particular agents for which they are intended. Agents also send events to the global event manager which stores this information in the central SQL server database.

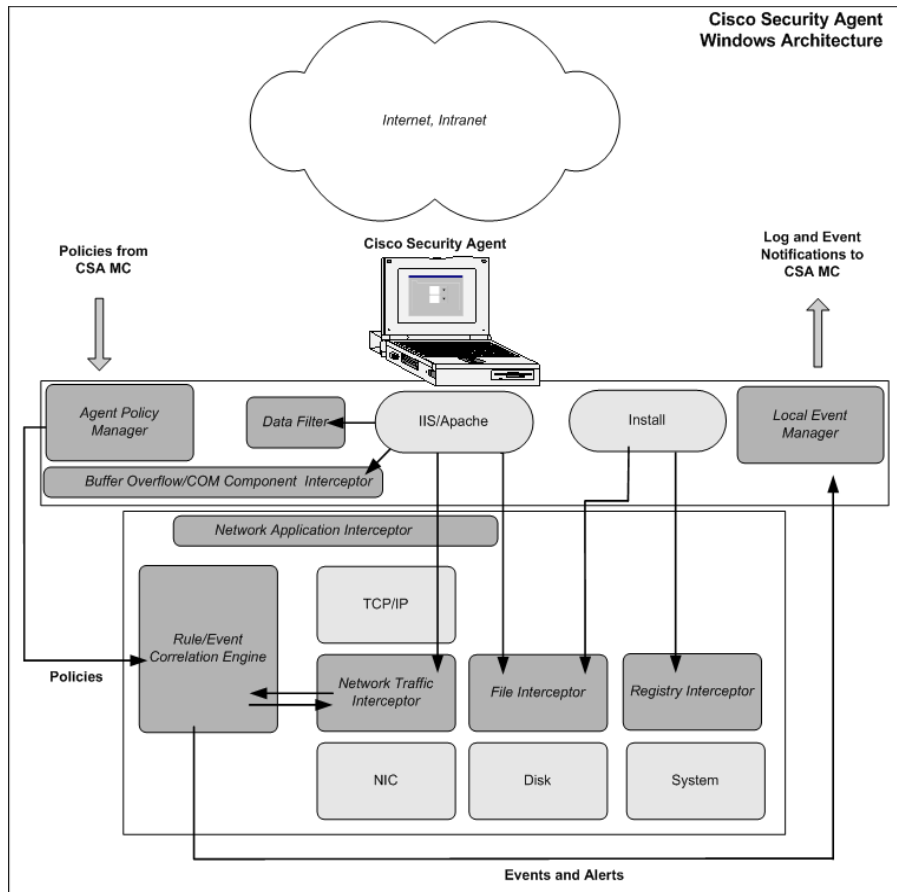
The **SQL server database** is the central repository for configuration data (host agents, groups, file rules, network rules, registry rules, etc.) created by the administrator and for the system event information provided by the agents. It is in this database that rules and information on system groupings are stored when the administrator generates rules and policies through the web-based interface. When reports are requested by the administrator, the **report generator** component gathers rule and event data kept in the database and produces reports using this information.

All information (rule configurations, event logs, etc.) passed between CSA MC and the agents distributed across your enterprise is encrypted providing a secure communication channel for the exchange of data.

Agent Components

[Figure B-2](#) shows the agent in terms of its system components, displaying where those components operate in relation to general system functions. For example, the interceptors shown in the diagram install and work at the kernel level.

Figure B-2 Cisco Security Agent Components (Windows)



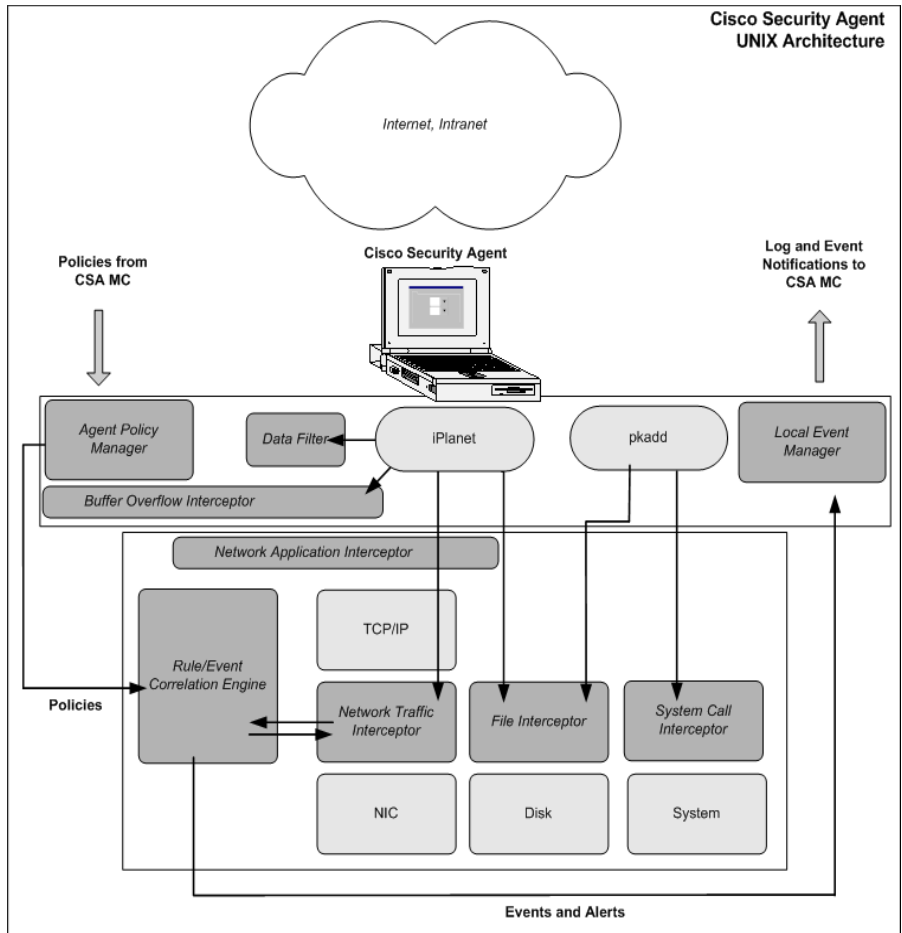
Starting from the left side of the diagram, the agent **policy manager** receives the rules configured by the administrator from CSA MC. These rules are sent to the agent's **rule/event correlation engine**. If a rule set already exists there, those rules are updated or replaced with the newest rule set.

The **interceptors** do as their name indicates, they intercept key actions that are attempted on the system and check the action in question against the rule correlation engine to determine if a rule set allows or denies it. Based on the information the interceptors receive, they either allow the action to take place or they stop it cold.

Actions are stopped based on certain criteria that are part of each rule and consequently each interceptor acts based on a component-targeted set of criteria. For example, the **network application interceptor** controls which applications are allowed to communicate with the network, while the **network traffic interceptor** provides system hardening features such as SYN flood protection and port scan detection. The **file interceptor** controls which applications can read and/or write to specified system files and directories. The **registry interceptor** controls system behavior, preventing applications from writing to particular registry keys. All of these controls can be as broad or as granular as necessary.

As the interceptors are allowing or denying actions, they produce an event each time a rule set is triggered by a system action. These events are stored in the rule/event correlation engine which forwards them on to the **local event manager** and **global event manager**. Events are also stored in the NT event log or W2K event viewer on the agent system.

Figure B-3 Cisco Security Agent Agent Components (UNIX)





Third Party Copyright Notices

Management Center for Cisco Security Agents utilizes third party software from various sources. Portions of this software are copyrighted by their respective owners as indicated in the copyright notices below.

Openssl license

Copyright © 1998-2000 The OpenSSL Project. All rights reserved.

Redistribution and use in source and binary forms, with our without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
3. All advertising materials mentioning features or use of this software must display the following acknowledgement:
“This product includes software developed by the OpenSSL Project for use in the OpenSSL Toolkit. (<http://www.openssl.org/>)”
4. The names “OpenSSL Toolkit” and “OpenSSL Project” must not be used to endorse or promote products derived from this software without prior written permission. For written permission, please contact openssl-core@openssl.org.

5. Products derived from this software may not be called “OpenSSL” nor may “OpenSSL” appear in their names without prior written permission of the OpenSSL Project.
6. Redistributions of any form whatsoever must retain the following acknowledgement:
“This product includes software developed by the OpenSSL Project for use in the OpenSSL Toolkit (<http://www.openssl.org/>)”

THIS SOFTWARE IS PROVIDED BY THE OpenSSL PROJECT “AS IS” AND ANY EXPRESSED OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE OpenSSL PROJECT OR ITS CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, BEEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

SSLEAY license

Copyright © 1995-1198 Eric Young (eay@cryptsoft.com)

All rights reserved.

This package is an SSL implementation written by Eric Young (eay@cryptsoft.com). The implementation was written so as to conform with Netscape’s SSL.

This library is free for commercial and non-commercial use as long as the following conditions are adhered to. The following conditions apply to all code found in this distribution, be it the RC4, RSA, lhash, DES, etc., code; not just the SSL code. The SSL documentation included with this distribution is covered by the same copyright terms except that the holder is Tim Hudson (tjh@cryptsoft.com).

Copyright remains Eric Young's, and as such any Copyright notices in the code are not to be removed.

If this package is used in a product, Eric Young should be given attribution as the author of the parts of the library used.

This can be in the form of a textual message at program startup or in documentation (online or textual) provided with the package.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
3. All advertising materials mentioning features or use of this software must display the following acknowledgement:
"This product includes cryptographic software written by Eric Young (ey@cryptsoft.com)"
The word 'cryptographic' can be left out if the routines from the library being used are not cryptographic related ;-).
4. If you include any Windows specific code (or a derivative thereof) from the apps directory (application code) you must include an acknowledgement:
"This product includes software written by Tim Hudson (tjh@cryptsoft.com)"

THIS SOFTWARE IS PROVIDED BY ERIC YOUNG "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE AUTHOR OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

Apache license

Copyright © 1995-1999 The Apache Group. All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
3. All advertising materials mentioning features or use of this software must display the following acknowledgment:
“This product includes software developed by the Apache Group for use in the Apache HTTP server project (<http://www.apache.org/>).”
4. The names “Apache Server” and “Apache Group” must not be used to endorse or promote products derived from this software without prior written permission. For written permission, please contact apache@apache.org.
5. Products derived from this software may not be called “Apache” nor may “Apache” appear in their names without prior written permission of the Apache Group.
6. Redistributions of any form whatsoever must retain the following acknowledgement:
“This product includes software developed by the Apache Group for use in the Apache HTTP server project (<http://www.apache.org/>).”

THIS SOFTWARE IS PROVIDED BY THE APACHE GROUP “AS IS” AND ANY EXPRESSED OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE APACHE GROUP OR ITS CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT

LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

TCL license

This software is copyrighted by the Regents of the University of California, Sun Microsystems, Inc., Scriptics Corporation, and other parties. The following terms apply to all files associated with the software unless explicitly disclaimed in individual files.

The authors hereby grant permission to use, copy, modify, distribute, and license this software and its documentation for any purpose, provided that the existing copyright notices are retained in all copies and that this notice is included verbatim in any distributions. No written agreement, license, or royalty fee is required for any of the authorized uses. Modifications to this software may be copyrighted by their authors and need not follow the licensing terms described here, provided that the new terms are clearly indicated on the first page of each file where they apply.

IN NO EVENT SHALL THE AUTHORS OR DISTRIBUTORS BE LIABLE TO ANY PARTY FOR DIRECT, INDIRECT, SPECIAL, INCIDENTAL, OR CONSEQUENTIAL DAMAGES ARISING OUT OF THE USE OF THIS SOFTWARE, ITS DOCUMENTATION, OR ANY DERIVATIVES THEREOF, EVEN IF THE AUTHORS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

THE AUTHORS AND DISTRIBUTORS SPECIFICALLY DISCLAIM ANY WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, AND NON-INFRINGEMENT. THIS SOFTWARE IS PROVIDED ON AN "AS IS" BASIS, AND THE AUTHORS AND DISTRIBUTORS HAVE NO OBLIGATION TO PROVIDE MAINTENANCE, SUPPORT, UPDATES, ENHANCEMENTS, OR MODIFICATIONS.

GOVERNMENT USE: If you are acquiring this software on behalf of the U.S. Government, the Government shall have only "Restricted Rights" in the software and related documentation as defined in the Federal Acquisition Regulations (FARs) in Clause 52.227.19 (c) (2). If you are acquiring the software on behalf of the Department of Defense, the software shall be classified as "Commercial Computer Software" and the Government shall have only "Restricted Rights" as

Perl License:

defined in Clause 252.227-7013 (c) (1) of DFARs. Notwithstanding the foregoing, the authors grant the U.S. Government and others acting in its behalf permission to use and distribute the software in accordance with the terms specified in this license.

Perl License:

Larry Wall's Copyright Notice Distributed with Perl. Copyright © 1989, 1990, 1991, Larry Wall. All rights reserved. This program is distributed in the hope That it will be useful, but WITHOUT ANY WARRANTY; without even the implied warranty of MERCHANTABILITY or FITNESS FOR PARTICULAR PURPOSE.

To get the standard perl source distribution, go to <http://www.cpan.org>.

libwww License

This product contains software developed by libwww: W3C's implementation of HTTP can be found at: <http://www.w3.org/Library/> Copyright © 1994-2000 World Wide Web Consortium, (Massachusetts Institute of Technology, Institut National De Recherche en Informatique et en Automatique, Keio University). All Rights Reserved. This program is distributed under the W3C's Software Intellectual Property License. This program is distributed in the hope that it will be useful, but WITHOUT WARRANTY; without even an implied warranty of MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See W3C License <http://www.w3.org/Consortium/Legal/> for more details.

This product includes computer software created and made available by CERN. Copyright © 1995 CERN.

libpcap

This product contains software derived from libpcap.

Copyright (c) 1988, 1989, 1990, 1991, 1993, 1994, 1995, 1996
The Regents of the University of California. All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that: (1) source opcode distributions retain the above copyright notice and this paragraph in its entirety, (2) distributions including binary opcode include the above copyright notice and this paragraph in its entirety in the documentation or other materials provided with the distribution, and (3) all advertising materials mentioning features or use of this software display the following acknowledgement: "This product includes software developed by the University of California, Lawrence Berkeley Laboratory and its contributors." Neither the name of the University nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, WITHOUT LIMITATION, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE.

CMU-SNMP Libraries

This product contains software developed by Carnegie Mellon University. Copyright 1998 by Carnegie Mellon University. All Rights Reserved

Permission to use, copy, modify, and distribute this software and its documentation for any purpose and without fee is hereby granted, provided that the above copyright notice appear in all copies and that both that copyright notice and this permission notice appear in supporting documentation, and that the name of CMU not be used in advertising or publicity pertaining to distribution of the software without specific, written prior permission.

CMU DISCLAIMS ALL WARRANTIES WITH REGARD TO THIS SOFTWARE, INCLUDING ALL IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS, IN NO EVENT SHALL CMU BE LIABLE FOR ANY SPECIAL, INDIRECT OR CONSEQUENTIAL DAMAGES OR ANY DAMAGES WHATSOEVER RESULTING FROM LOSS OF USE, DATA OR PROFITS, WHETHER IN AN ACTION OF CONTRACT, NEGLIGENCE OR OTHER TORTIOUS ACTION, ARISING OUT OF OR IN CONNECTION WITH THE USE OR PERFORMANCE OF THIS SOFTWARE.

Open Market Inc., Fastcgi license

This product contains software developed by Open Market Inc.

THIS FASTCGI APPLICATION LIBRARY SOURCE AND OBJECT CODE (THE "SOFTWARE") AND ITS DOCUMENTATION (THE "DOCUMENTATION") ARE COPYRIGHTED BY OPEN MARKET, INC ("OPEN MARKET"). THE FOLLOWING TERMS APPLY TO ALL FILES ASSOCIATED WITH THE SOFTWARE AND DOCUMENTATION UNLESS EXPLICITLY DISCLAIMED IN INDIVIDUAL FILES.

OPEN MARKET PERMITS YOU TO USE, COPY, MODIFY, DISTRIBUTE, AND LICENSE THIS SOFTWARE AND THE DOCUMENTATION SOLELY FOR THE PURPOSE OF IMPLEMENTING THE FASTCGI SPECIFICATION DEFINED BY OPEN MARKET OR DERIVATIVE SPECIFICATIONS PUBLICLY ENDORSED BY OPEN MARKET AND PROMULGATED BY AN OPEN STANDARDS ORGANIZATION AND FOR NO OTHER PURPOSE, PROVIDED THAT EXISTING COPYRIGHT NOTICES ARE RETAINED IN ALL COPIES AND THAT THIS NOTICE IS INCLUDED VERBATIM IN ANY DISTRIBUTIONS.

NO WRITTEN AGREEMENT, LICENSE, OR ROYALTY FEE IS REQUIRED FOR ANY OF THE AUTHORIZED USES. MODIFICATIONS TO THIS SOFTWARE AND DOCUMENTATION MAY BE COPYRIGHTED BY THEIR AUTHORS AND NEED NOT FOLLOW THE LICENSING TERMS DESCRIBED HERE, BUT THE MODIFIED SOFTWARE AND DOCUMENTATION MUST BE USED FOR THE SOLE PURPOSE OF IMPLEMENTING THE FASTCGI SPECIFICATION DEFINED BY OPEN MARKET OR DERIVATIVE SPECIFICATIONS PUBLICLY ENDORSED BY OPEN MARKET AND PROMULGATED BY AN OPEN STANDARDS ORGANIZATION AND FOR NO OTHER PURPOSE. IF MODIFICATIONS TO THIS SOFTWARE AND DOCUMENTATION HAVE NEW LICENSING TERMS, THE NEW TERMS MUST PROTECT OPEN MARKET'S PROPRIETARY RIGHTS IN THE SOFTWARE AND DOCUMENTATION TO THE SAME EXTENT AS THESE LICENSING TERMS AND MUST BE CLEARLY INDICATED ON THE FIRST PAGE OF EACH FILE WHERE THEY APPLY.

OPEN MARKET SHALL RETAIN ALL RIGHT, TITLE AND INTEREST IN AND TO THE SOFTWARE AND DOCUMENTATION, INCLUDING WITHOUT LIMITATION ALL PATENT, COPYRIGHT, TRADE SECRET AND OTHER PROPRIETARY RIGHTS.

OPEN MARKET MAKES NO EXPRESS OR IMPLIED WARRANTY WITH RESPECT TO THE SOFTWARE OR THE DOCUMENTATION, INCLUDING WITHOUT LIMITATION ANY WARRANTY OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. IN NO EVENT SHALL OPEN MARKET BE LIABLE TO YOU OR ANY THIRD PARTY FOR ANY DAMAGES ARISING FROM OR RELATING TO THIS SOFTWARE OR THE DOCUMENTATION, INCLUDING, WITHOUT LIMITATION, ANY INDIRECT, SPECIAL OR CONSEQUENTIAL DAMAGES OR SIMILAR DAMAGES, INCLUDING LOST PROFITS OR LOST DATA, EVEN IF OPEN MARKET HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. THE SOFTWARE AND DOCUMENTATION ARE PROVIDED "AS IS". OPEN MARKET HAS NO LIABILITY IN CONTRACT, TORT, NEGLIGENCE OR OTHERWISE ARISING OUT OF THIS SOFTWARE OR THE DOCUMENTATION.

CGIC License

Basic License

CGIC, copyright 1996, 1997, 1998, 1999, 2000, 2001, 2002 by Thomas Boutell and Boutell.Com, Inc.. Permission is granted to use CGIC in any application, commercial or noncommercial, at no cost. HOWEVER, this copyright paragraph must appear on a "credits" page accessible in the public online and offline documentation of the program. Modified versions of the CGIC library should not be distributed without the attachment of a clear statement regarding the author of the modifications, and this notice may in no case be removed. Modifications may also be submitted to the author for inclusion in the main CGIC distribution.

Mozilla 1.xx (libcurl)

The curl and libcurl are dual-licensed under the MPL and the MIT/X-derivative licenses. This software is licensed under an MIT/X-derivative license as shown here:

COPYRIGHT AND PERMISSION NOTICE

Copyright (c) 1996—2001, Daniel Stenberg, <daniel@haxx.se>.

All rights reserved.

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, provided that the above copyright notice(s) and this permission notice appear in all copies of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OF THIRD PARTY RIGHTS. IN NO EVENT SHALL THE COPYRIGHT HOLDER OR HOLDERS INCLUDED IN THIS NOTICE BE LIABLE FOR ANY CLAIM, OR ANY SPECIAL INDIRECT OR CONSEQUENTIAL DAMAGES, OR ANY DAMAGES WHATSOEVER RESULTING FROM LOSS OF USE, DATA OR PROFITS, WHETHER IN AN ACTION OF CONTRACT, NEGLIGENCE OR OTHER TORTIOUS ACTION, ARISING OUT OF OR IN CONNECTION WITH THE USE OR PERFORMANCE OF THIS SOFTWARE.

Except as contained in this notice, the name of a copyright holder shall not be used in advertising or otherwise to promote the sale, use or other dealings in this Software without prior written authorization of the copyright holder.



Symbols

- @CD [2-24](#)
- @dynamic [2-23, 4-44, 7-9](#)
- @fixed [2-20](#)
- @floppy [2-24](#)
- @local [4-52](#)
- @network [2-25](#)
- @regpath [2-23](#)
- @removable [2-24](#)
- @system [2-22](#)

A

ActiveX

- crystal report viewer [9-8](#)
- preventing download [5-7](#)

Address set syntax [2-26](#)

Add to application class [4-6, 6-17](#)

Administrator

- by operating system [2-5](#)
- role-based administration [2-5](#)

Agent

- architecture [1-6](#)
- Clear saved responses [A-11](#)
- kits [3-8](#)
- optional reboot after install [A-6](#)
- scripted silent installs and uninstalls [3-19](#)
- user interface [A-7](#)

Agent (UNIX)

- check status [A-19](#)
- commands [A-18](#)
- csasctl utility [A-18](#)
- installing [A-16](#)
- polling [A-18](#)
- re-enable logging [A-18](#)
- software update [A-19](#)
- uninstalling [A-18](#)

Agent kits

- automatic reboot [3-11](#)
- kit creation combinations [3-15](#)
- network shim install [3-11](#)
- preconfigured sample [3-3, 3-9](#)
- quiet install [3-10](#)
- reboot vs. no reboot [3-17](#)
- status [3-14](#)
- UNIX installation [A-16](#)

- Agent polling [11-9](#)
 - Agent registration [3-18](#)
 - Registration control [3-19](#)
 - Agent service
 - stop and start [A-13](#)
 - Agent service control [4-24](#)
 - suspend agent security [4-26](#)
 - Alert types
 - configuring [8-27](#)
 - log file, generate [8-34](#)
 - Analysis job
 - analysis process overview [11-3](#)
 - configuration [11-6](#)
 - End job at time [11-9](#)
 - import policy [11-12](#)
 - job start job at time [11-9](#)
 - Job status field [11-6](#)
 - monitor event log [11-10](#)
 - overview [11-3](#)
 - progression [11-10](#)
 - start analysis [11-11](#)
 - Analysis job wizard [8-17](#)
 - Analysis process overview [11-3](#)
 - Analysis Reports [11-18](#)
 - Application-builder rule [6-15](#)
 - Application Classes [6-2, 6-6](#)
 - application-builder rule [6-15](#)
 - application class management [6-22](#)
 - configure downloaded content [6-9](#)
 - creation from rule page [6-20](#)
 - double-click to view [6-21](#)
 - dynamic [6-11](#)
 - enable/disable for product [6-22](#)
 - managing application classes [6-7, 6-22](#)
 - Network Applications [6-4](#)
 - Processes created by Network Applications [6-4](#)
 - Processes created by Servers (TCP and UDP) [6-4](#)
 - Processes executing downloaded content [6-4, 6-9](#)
 - Processes with elevated privileges [6-5](#)
 - Remote clients [6-5](#)
 - remove process [6-8, 6-14](#)
 - Server (TCP based) [6-5](#)
 - Server (UDP based) [6-5](#)
 - shell scripts [6-3](#)
 - static [6-6](#)
 - System Process [6-5](#)
- Application classes to analyze [11-7](#)
 - Application class management [6-7, 6-22](#)
 - Application control [4-28](#)
 - Are you there?, VPN client [13-1](#)
 - Audit trail [2-6](#)
 - Automatic reboot on agents [3-11](#)

Available button

agent [A-8](#)

Available software updates [3-26](#)

B

Backing up configurations [10-3](#)

differential backup [10-4](#)

full backup [10-4](#)

BootExecute [7-20](#)

broadcast messages [4-53](#)

Browser requirements [2-2](#)

Buffer overflow attack

preventing [5-7](#)

Buffer Overflow rule [5-16](#)

replicate feature [5-9](#)

Built-in agent self-protection [4-26](#)

Bulk transfer hosts [3-25](#)

C

CD ROM drives [2-24](#)

Check Point

integration [13-2](#)

Third Party directory [13-3](#)

Check Point OPSEC integration [13-2](#)

Cisco.com, accessing [xv](#)

Cisco Security Monitor

integration [13-2](#)

Cisco VPN client support [13-1](#)

Clear Pending Alerts [8-28](#)

Cloak system [5-12](#)

Clone configurations [2-17, 11-5](#)

COM component

access control rule [4-55](#)

extract utility [10-9](#)

COM Component Sets [7-21](#)

extract utility [10-9](#)

Compare configurations [2-18, 4-19](#)

Compare policies [4-19](#)

configuration manager [B-3](#)

Configuration shortcuts [2-15](#)

Configuration view [2-15](#)

Connection Rate Limit [4-33](#)

Contact information

agent [3-22](#)

Content matching, file sets UNIX [7-9](#)

Copy rules [4-18](#)

Correlating events [5-18](#)

CSA_uninstall.bat [3-19](#)

csactl Utility (UNIX agent) [A-18](#)

CSA MC

logging in [2-4](#)

D

Data access control [4-35](#)

Database restoring log [10-6](#)

- Data filter installation, required for data access control rule [10-10](#)
 - Data Sets [7-3](#)
 - Delete configurations [2-17](#)
 - Deployment Overview [1-4](#)
 - Detailed descriptions [4-2](#)
 - Details link [8-5](#)
 - Directory path protection (UNIX only) [2-24](#)
 - Disable application class [6-22](#)
 - Disabled rules [11-16](#)
 - Disable policy rule enforcement [8-17, 11-8](#)
 - Disable rule [4-17](#)
 - Disk space
 - shrink your database files [10-7](#)
 - Distribute agent kits [3-12](#)
 - documentation
 - feedback, submitting electronically [xvi](#)
 - obtaining [xv](#)
 - CD-ROM [xv](#)
 - Cisco.com [xv](#)
 - ordering [xvi](#)
 - other Cisco publications and information [xix](#)
-
- E**
- Enable rule [4-17](#)
 - End job at time [11-9](#)
 - Ephemeral ports
 - using [2-27](#)
 - ephemeral ports [7-15](#)
 - Event code [4-69](#)
 - Event Correlation [5-18](#)
 - Event ID [4-69](#)
 - Event log [8-2](#)
 - change filter [8-2](#)
 - Details link [8-5](#)
 - event log generation time [8-2](#)
 - filter text [8-2](#)
 - Find Similar [8-5](#)
 - minimum and maximum severity [8-3](#)
 - Rule number link [8-5](#)
 - start and end dates [8-2](#)
 - Event log analysis job messages [11-10](#)
 - Event log generation time [8-2](#)
 - Event logging exception wizard [8-15](#)
 - Event Log Management [8-6](#)
 - Event Management Wizard [8-9](#)
 - Event Monitor [8-5](#)
 - Events [8-2](#)
 - managing [8-6](#)
 - minimum severity [8-3](#)
 - third party access [8-24](#)
 - Events by Group reports [9-4](#)
 - Events by Severity reports [9-2](#)
 - Event Sets [8-20](#)
 - Purge events [8-23](#)
 - View button [8-23](#)
 - Exception rule wizard [8-11](#)
 - Explain rules link [4-23](#)

Export configurations [10-12](#)
Export reports [9-10](#)
extract_com [10-9](#)

F

File access control [4-40](#)
file interceptor [B-5](#)
File monitor [4-45](#)
File Sets [7-6](#)
File set syntax [2-19](#)
File version control [4-59](#)
Find database items [2-9](#)
Find Similar [8-5](#)
Floppy drives [2-20](#)
floppy drives [2-24](#)
Free up disk space [10-7](#)

G

Generating Rules [4-97](#)
 Details link [4-97](#)
 fails due to polling interval [4-97](#)
 pending changes [2-17](#)
Global event correlation [5-18](#)
Global Event Management
 Correlation [5-18](#)
global event manager [B-2](#)

Groups
 about [3-2](#)
 attaching policies to [4-92](#)
 configure [3-3](#)
 No user interaction [3-6, A-7](#)
 Polling interval [3-6](#)
 preconfigured sample [3-3](#)
 Test Mode [3-5](#)
 Verbose Logging Mode [3-5](#)

H

Help
 online [2-11](#)
help [xvi](#)
 Cisco.com [xvii](#)
 TAC [xvii](#)
 Escalation Center [xix](#)
 website [xviii](#)
Host groups [3-2](#)
Hosts
 about [3-3](#)
 bulk transfer [3-25](#)
 changing groups [3-23](#)
 unprotected hosts [2-10](#)
 view status [3-20](#)

Host Status

- Active [3-20](#)
- Last Poll [3-20](#)
- Latest Software [3-20](#)
- Protected [3-20](#)
- Test Mode [3-20](#)

HTTPS [A-4](#)

I

ID

- rule [4-17](#)

IDS [4-95](#)

Import configurations [10-12](#)

- renaming duplicate items [10-17](#)

Import policy [11-5, 11-12](#)

Install

- agent [A-2](#)

Insufficient disk space event [10-7](#)

interceptors [B-4](#)

- file [B-5](#)
- network application [B-5](#)
- network traffic [B-5](#)
- registry [B-5](#)

Internet Explorer

- version requirements [2-2](#)

Intrinsic security [1-1](#)

IPlanet server data filter [4-36](#)

IP security checks [5-10](#)

IPV6 addresses [2-26, 7-12](#)

J

Job status field [11-6](#)

K

Kernel Protection [4-64](#)

Keystroke trapping

- preventing [5-6](#)
-

L

Lifecycle of an attack [1-2](#)

List view [2-15](#)

local event manager [B-5](#)

Logging [8-8](#)

- suppression of log messages [8-8](#)

Logging agent [11-3](#)

Log size [11-9](#)

M

Maintenance

- Available Software Updates [3-26](#)

- Event Log Management [8-6](#)

Make kit button [3-12](#)

- Manage application classes [6-7, 6-22](#)
- Mandatory policies
 - about [4-9](#)
- Mandatory policy for all groups [4-15](#)
- Manual data filter installation [10-10](#)
- Media type syntax [2-24](#)
- Menu bar [2-7, 2-17](#)
- Merge configurations [2-18, 4-22](#)
- Merge policies [4-22](#)
- minimum severity [8-3](#)
- Modify agent configuration
 - self-protection [4-26](#)
- Monitor
 - Event Log [8-2](#)
 - Event Monitor [8-5](#)
 - Event Sets [8-20](#)
 - Status Summary [2-13](#)
- Monitor analysis job [11-10](#)
- Monitoring Access [4-9](#)
- multicast packet signals [4-53, 8-8](#)
- Navigation shortcuts
 - insert link [2-15](#)
 - new application class link [2-15](#)
 - new item creation [2-15](#)
 - right click menu access [2-16](#)
 - show reference list [2-15](#)
 - variables [2-15](#)
- netForensics
 - integration [13-2](#)
- Netscape
 - version support [2-2](#)
- net start command [10-2, A-13](#)
- net stop command [10-2, A-13](#)
- Network access control [4-49](#)
- Network Address Sets [7-11](#)
- network application interceptor [B-5](#)
- Network Applications [6-4](#)
- Network interface control [4-81](#)
- Network service ephemeral ports [2-27](#)
- Network Services [7-14](#)
- Network service syntax [2-26](#)

N

- Naming convention
 - Analysis job policy [11-14](#)

Network shield 5-9cloak system (prevent port scans) **5-12**dangerous ICMP messages **5-13**detect port scans **5-11**ICMP covert channels **5-13**invalid IP addresses **5-10**invalid IP headers **5-10**invalid transport headers **5-11**IP source routing **5-10**prevent SYN floods **5-13**randomize TCP sequence numbers **5-12**replicate feature **5-9**restrict network connectivity **5-14**suspicious packets **5-13**trace route **5-10****Network shim 3-11**optional **A-4**network traffic interceptor **B-5**Network worm protection **5-3**New button **2-17, 11-5**No user interaction **3-6, A-7**NT Event log **4-67****O**Operating system changes, agent **A-4****P**packet sniffers **4-78, 4-81**

Password stealing

preventing **5-6**Peripherals **2-20, 2-24**

Policies

about mandatory policies **4-9**adding rules **4-16**agent service control **4-24**application control **4-28**applying **4-92**buffer overflow **5-16**combining **4-6**com component access control **4-55**compare, copy, merge **4-19**configuring **4-13**copying rules between policies **4-18**file access control **4-33, 4-35, 4-40, 4-45, 4-55, 4-70**file monitor **4-45**file version control **4-59**generate rules **4-97**mandatory policies **4-15**network access control **4-49**network interface control **4-81**network shield **5-9**NT event log **4-67**port scan detection **5-11**preparing a security policy **1-7**

- registry access control [4-70](#)
- resource access control [4-84](#)
- rootkit /kernel protection [4-87](#)
- service restart [4-74](#)
- sniffer and protocol detection [4-77](#)
- SYN flood protection [5-13](#)
- Syslog control [4-88](#)
- Policy creation guidelines [12-1](#)
- Policy enforcement [8-17, 11-8](#)
- polling [1-9](#)
- Polling interval [3-6](#)
 - fast polling [A-11](#)
- Port scan detection [5-11](#)
- Processes created by Network Applications [6-4](#)
- Processes created by Servers (TCP and UDP) [6-4](#)
- Processes executing downloaded content [6-9](#)
- Processes with elevated privileges [6-5](#)
- Profiler
 - Analysis process [11-3](#)
 - Analysis Reports [11-18](#)
 - Configure Analysis Job [11-6](#)
 - Overview [11-1](#)
 - Policy creation methodology [11-15](#)
 - Policy enforcement [11-8](#)
 - policy naming convention [11-14](#)
- promiscuous mode [4-83](#)
- Purge events [8-23](#)

Q

- Quarantined files [2-23, 4-44, 5-19, 7-9](#)
 - quarantined file events [2-23, 5-21, 7-9](#)
- Query User [4-10, 4-38, 4-42, 4-50, 4-57, 4-72](#)
 - cached responses [4-12](#)
 - Text used to query user [4-38, 4-42](#)
- Quiet install [3-10](#)

R

- Reboot operations [7-20](#)
- Reboot optional
 - agent [A-5, A-6](#)
- Reboot vs. no reboot for agents [3-17](#)
- Registration control [3-19](#)
- Registry access control [4-70](#)
- registry interceptor [B-5](#)
- Registry sets
 - BootExecute [7-20](#)
 - Reboot operations [7-20](#)
 - Run keys [7-20](#)
 - Shell commands [7-20](#)
- Remote access [2-4](#)
- Remote clients [6-5](#)
- Removable media [2-24](#)
- Remove process from application class [6-8, 6-14](#)
- Renaming duplicate configuration items [10-17](#)
- Replace items, search [2-10](#)

replicate feature (rule option) [5-9](#)

report generator [B-3](#)

Reports

Events by Group [9-4](#)

Events by Severity [9-2](#)

exporting [9-10](#)

generating [9-2](#)

Group Detail [9-7](#)

Host Detail [9-5](#)

Policy Detail [9-7](#)

printing [9-9](#)

refresh [9-10](#)

Report viewer [9-8](#)

Requirements

agent [A-2](#)

Resource access control [4-84](#)

Restoring configurations [10-6](#)

Resume agent security [A-13](#)

Right-click menu shortcuts [2-18](#)

Role-based administration [2-5](#)

Rootkit / kernel protection [4-87](#)

Rule ID [4-17](#)

Rules

about [4-5](#)

agent service control [4-24](#)

allow vs. deny [4-6](#)

application control [4-28](#)

buffer overflow [5-16](#)

com component access control [4-55](#)

compare, copy, merge [4-19](#)

connection rate limit [4-33](#)

copying rules between policies [4-18](#)

data access control [4-35](#)

enable/disable [4-17](#)

Events column [4-17](#)

explain link [4-23](#)

file access control [4-40](#)

file monitor [4-45](#)

file version control [4-59](#)

ID column [4-17](#)

kernel protection [4-64](#)

Logging options [4-8](#)

manipulating precedence [4-8](#)

network access control [4-49](#)

network interface control [4-81](#)

network shield [5-9](#)

network worm protection [5-3](#)

NT event log [4-67](#)

port scan detection [5-11](#)

precedence, ordering [4-6](#)

registry access control [4-70](#)

- resource access control [4-84](#)
- rootkit / kernel protection [4-87](#)
- service restart [4-74](#)
- show enabled rules only checkbox [4-17](#)
- sniffer and protocol detection [4-77](#)
- SYN flood protection [5-13](#)
- syslog control [4-88](#)
- Trojan detection [5-4](#)
- View All rules [4-17](#)
- view change history [4-22](#)

Run keys [7-20](#)

S

- Sample policies [8-20](#)
 - network [11-16](#)
- Save configurations [2-17](#)
- Save job [11-5](#)
- Scheduled database backups [10-4](#)
- Scheduled software updates [3-28](#)
- Scripted agent installs and uninstalls [3-19](#)
- Scripts, writing rules for [6-3](#)
- Search
 - how to use it [2-9](#)
 - replace [2-10](#)
- Security Monitor
 - integration [13-2](#)
- Security policy [4-2](#)
 - preparing [1-7](#)
- Self-protection rules
 - agent [4-26](#)
- Server (TCP based) [6-5](#)
- Server (UDP based) [6-5](#)
- service, agent start/stop [10-2, A-14](#)
- Service restart [4-74](#)
- Shell commands [7-20](#)
- Shell scripts, writing rules for [6-3](#)
- show enabled rules only checkbox [4-17](#)
- Show reference list [2-15](#)
- Silent agent install and uninstall [3-19](#)
- Sniffer and protocol detection rule [4-77](#)
- Software updates [3-26](#)
 - agents [A-14](#)
 - Installing updates on agents [A-14](#)
 - Scheduled software updates [3-28](#)
- Solaris agent install directory [A-17](#)
- Solaris requirements
 - agent [A-3](#)
- SQL Server database [B-3](#)
- SSL [1-8](#)
- Start agent service [10-2, A-14](#)
- Start analysis [11-5](#)
- Start job at time [11-9](#)
- Status
 - Event Log [8-2](#)
- Status, agent kits [3-14](#)

- Status Summary [2-13](#)
 - colored chart [2-13](#)
 - Stop after application invocations [11-9](#)
 - Stop agent service [4-26, 10-2, A-14](#)
 - Stop and start agent security [A-13](#)
 - Stop logging button [11-8](#)
 - Suspend agent security [A-13](#)
 - all users [4-26](#)
 - Symbolic link protection
 - file access control [2-25, 4-44, 4-86](#)
 - general [4-85](#)
 - SYN Flood Protection [5-13](#)
 - Syntax [2-18](#)
 - Syslog control [4-88](#)
 - System components [B-1](#)
 - System Process [6-5](#)
 - System Startup Security checks [5-14](#)
-
- T**
- TAC (Technical Assistance Center) [xvii](#)
 - Escalation Center [xix](#)
 - website [xviii](#)
 - Take precedence over other rules [4-8, 4-9](#)
 - technical support [xvi](#)
 - Cisco.com [xvii](#)
 - TAC [xvii](#)
 - Escalation Center [xix](#)
 - website [xviii](#)
 - Terminal services [A-2](#)
 - Test Mode [3-5, 4-95](#)
 - Text used to query user [4-38, 4-42](#)
 - Third party
 - access to events [8-24](#)
 - Third Party directory [13-3](#)
 - Tool tips [A-12](#)
 - Trace route, prevent [5-10](#)
 - Transport security checks [5-11](#)
 - Trojan Detection [5-4](#)
 - replicate feature [5-9](#)
-
- U**
- UNIX administrator [2-5](#)
 - UNIX agent install directory [A-17](#)
 - Unprotected hosts [2-10](#)
 - Utilities
 - csactl (UNIX agent) [A-18](#)
 - extract_com [10-9](#)
 - net stop/start [10-2](#)

V

Variables

COM Component Sets [7-21](#)

Data Sets [7-3](#)

Event Sets [8-20](#)

File Sets [7-6](#)

Network Address Sets [7-11](#)

Network Services [7-14](#)

Registry Sets [7-17](#)

using [7-2](#)

Verbose logging mode [3-5, 8-9](#)

View All rules

filter rule display [4-17](#)

View change history [7-10](#)

View change history, rules [4-22](#)

Vulnerable applications [12-9](#)

W

Waving system tray flag [A-10](#)

Windows administrator [2-5](#)

Windows requirements

agent [A-2](#)

Wizard

analysis job [8-15, 8-17](#)

events [8-9](#)

exception rule [8-11](#)

