

HIPS

Securing Hosts Using Cisco Security Agent

Version 2.0

Student Guide

CLS Production Services: 06.07.05

Copyright © 2005, Cisco Systems, Inc. All rights reserved.

Cisco Systems has more than 200 offices in the following countries and regions. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at www.cisco.com/go/offices.

Argentina • Australia • Austria • Belgium • Brazil • Bulgaria • Canada • Chile • China PRC • Colombia • Costa Rica
Croatia • Cyprus • Czech Republic • Denmark • Dubai, UAE • Finland • France • Germany • Greece
Hong Kong SAR • Hungary • India • Indonesia • Ireland • Israel • Italy • Japan • Korea • Luxembourg • Malaysia
Mexico • The Netherlands • New Zealand • Norway • Peru • Philippines • Poland • Portugal • Puerto Rico • Romania
Russia • Saudi Arabia • Scotland • Singapore • Slovakia • Slovenia • South Africa • Spain • Sweden • Switzerland
Taiwan • Thailand • Turkey • Ukraine • United Kingdom • United States • Venezuela • Vietnam • Zimbabwe



Copyright © 2005 Cisco Systems, Inc. All rights reserved. CCSP, the Cisco Square Bridge logo, Follow Me Browsing, 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 Access Registrar, Aironet, ASIST, BPX, Catalyst, CCDA, CCDP, CCIE, CCIP, CCNA, CCNP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Cisco Unity, Empowering the Internet Generation, Enterprise/Solver, EtherChannel, EtherFast, EtherSwitch, Fast Step, FormShare, GigaDrive, GigaStack, HomeLink, Internet Quotient, IOS, IP/TV, iQ Expertise, the iQ logo, iQ Net Readiness Scorecard, LightStream, Linksys, MeetingPlace, MGX, the Networkers logo, Networking Academy, Network Registrar, Packet, PIX, Post-Routing, Pre-Routing, ProConnect, RateMUX, ScriptShare, SlideCast, SMARTnet, StrataView Plus, 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 United States and certain other countries.

All other trademarks mentioned in this document or Website 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. (0501R)

DISCLAIMER WARRANTY: THIS CONTENT IS BEING PROVIDED "AS IS." CISCO MAKES AND YOU RECEIVE NO WARRANTIES IN CONNECTION WITH THE CONTENT PROVIDED HEREUNDER, EXPRESS, IMPLIED, STATUTORY OR IN ANY OTHER PROVISION OF THIS CONTENT OR COMMUNICATION BETWEEN CISCO AND YOU. CISCO SPECIFICALLY DISCLAIMS ALL IMPLIED WARRANTIES, INCLUDING WARRANTIES OF MERCHANTABILITY, NON-INFRINGEMENT AND FITNESS FOR A PARTICULAR PURPOSE, OR ARISING FROM A COURSE OF DEALING, USAGE OR TRADE PRACTICE. This learning product may contain early release content, and while Cisco believes it to be accurate, it falls subject to the disclaimer above.

Table of Contents

Volume 1

<i>Course Introduction</i>	1
Overview	1
Course Objectives	2
Lab Topology Overview	7
<i>Security Fundamentals</i>	1-1
Overview	1-1
Objectives	1-1
Need for Network Security	1-2
Network Security Policy	1-9
Primary Network Threats and Attacks	1-12
Reconnaissance Attacks and Mitigation	1-15
Access Attacks and Mitigation	1-22
Denial of Service Attacks and Mitigation	1-30
Worm, Virus, and Trojan Horse Attacks and Mitigation	1-35
Management Protocols and Functions	1-42
Summary	1-47
<i>Cisco Security Agent Overview</i>	2-1
Overview	2-1
Objectives	2-1
Defense in Depth	2-2
Cisco Security Agent Architecture	2-7
Anatomy of an Attack and Response	2-11
Key Features of Cisco Security Agent	2-14
Summary	2-17
<i>Cisco Security Agent Quick Start Installation</i>	3-1
Overview	3-1
Objectives	3-1
CSA MC System Requirements	3-2
CSA System Requirements	3-6
Installing the CSA MC	3-11
Configuring the CSA MC	3-13
Installing the CSA	3-21
Summary	3-23
<i>Cisco Security Agent Management Center Administration</i>	4-1
Overview	4-1
Objectives	4-1
Using Cisco Security Agent Management Center	4-2
Summary	4-15
<i>Configuring Groups and Managing Hosts</i>	5-1
Overview	5-1
Objectives	5-1
Configuring Groups	5-2
Building an Agent Kit	5-10
Managing Hosts	5-19
Deploying Scheduled Software Updates	5-26
Summary	5-30
<i>Building Policies</i>	6-1
Overview	6-1
Objectives	6-1

Developing a Security Policy	6-2
Policy Components	6-6
Building Policies and Rule Modules	6-10
Attaching Rule Modules to Policies	6-27
Summary	6-34
<i>Rule Basics</i>	<i>7-1</i>
Overview	7-1
Objectives	7-1
Basics of Rule Construction and Functionality	7-2
Rules Common to Windows and UNIX	7-13
Windows-Only Rules	7-38
UNIX-Only Rules	7-59
Summary	7-72
<i>System Correlation Rules</i>	<i>8-1</i>
Overview	8-1
Objectives	8-1
About System Correlation Rules	8-2
System API Control Rule	8-4
Network Shield Rule	8-8
Buffer Overflow Rule	8-14
E-Mail Worm Protection Rule Module	8-18
E-Mail Worm Event Correlation	8-19
Installation Applications Policy	8-20
Global Events	8-21
Correlation	8-21
Manage Dynamically Quarantined Files and IP Addresses	8-23
Summary	8-24
<i>Defining Application Classes</i>	<i>9-1</i>
Overview	9-1
Objectives	9-1
About Application Classes	9-2
Processes Created by Application Classes	9-2
Removing Processes from Application Classes	9-2
Shell Scripts and Application Classes	9-3
Preserving Application Process Classes	9-7
Configuring Static Application Classes	9-8
Dynamic Application Classes	9-12
Building Classes as Rule Consequences	9-12
Removing Processes from Classes	9-13
Create New Application Classes from Rule Pages	9-19
Summary	9-22
<i>Working with Variables</i>	<i>10-1</i>
Overview	10-1
Objectives	10-1
Variables	10-2
Display Only in Show All Mode Option	10-3
Data Sets	10-4
File Sets	10-7
Network Address Sets	10-11
Network Services Sets	10-13
Registry Sets	10-16
Included Registry Sets	10-16
COM Component Sets	10-20
Query Settings	10-23
Localized Language Version Support	10-26

Summary	10-27
<i>Using Cisco Security Agent Analysis</i>	<i>11-1</i>
Overview	11-1
Objectives	11-1
Application Deployment Investigation	11-3
Group Settings	11-4
Product Associations	11-6
Data Management	11-10
Application Deployment Reports	11-12
Viewing Reports	11-28
Exporting Reports	11-28
Application Behavior Investigation	11-30
Monitoring the Behavior Analysis	11-37
Start Behavior Analysis	11-37
Importing the Rule Module	11-38
Behavior Analysis Reports	11-39
Report Components	11-39
Working with Reports	11-44
Behavior Analysis Rule Modules	11-45
Reviewing the Rule Module	11-45
Variable and Application Class Creation	11-47
Summary	11-48
<i>Using Event Logs and Generating Reports</i>	<i>12-1</i>
Overview	12-1
Objectives	12-1
How Logging Works	12-2
The Event Log and Event Monitor	12-6
Start Date and End Date	12-7
Minimum and Maximum Severity Settings	12-7
Host	12-7
Events per Page	12-7
Filter Out Duplicates	12-7
Event Log Management	12-10
Event Sets	12-13
Configuring Alerts	12-17
Generating Reports	12-20
Summary	12-26

Course Introduction

Overview

This lesson includes the following topics:

- Course Objectives
- Course Agenda
- Participant Responsibilities
- General Administration
- Graphic Symbols
- Participant Introductions
- Cisco Security Career Certifications
- Lab Topology Overview

Course Objectives

This topic introduces the course and the course objectives.

Course Objectives

Cisco.com

Upon completion of this course, you will be able to perform the following tasks:

- Identify the platforms and infrastructure that support CSA and the CSA MC
- Describe the CSA architecture and the CSA MC
- Configure the way CSA protects a host system
- Install CSA with a default Agent kit
- Create host groups and build Agent kits

© 2005 Cisco Systems, Inc. All rights reserved. HIPS v2.0-0-3

Course Objectives (Cont.)

Cisco.com

- Define application classes and associate them with the appropriate security policies
- Use variables for granular control when creating rules
- Configure security policies and rules
- Configure System Correlation rules for CSA
- Identify which rules are for Windows, UNIX, and both platforms
- Perform data analysis and create policies with CSA Analysis
- Manage the Event Log and generate reports

© 2005 Cisco Systems, Inc. All rights reserved. HIPS v2.0-0-4

Course Agenda

Cisco.com

Day 1

- **Course Introduction**
- **Lesson 1: Security Fundamentals**
- **Lesson 2: Cisco Security Agent Overview**
- **Lunch**
- **Lesson 3: Cisco Security Agent Quick Start Installation**
- **Lesson 4: Cisco Security Agent Management Center Administration**
- **Lesson 5: Configuring Groups and Managing Hosts**

© 2005 Cisco Systems, Inc. All rights reserved.

HIPS v2.0—0-5

Course Agenda (Cont.)

Cisco.com

Day 2

- **Lesson 6: Building Policies**
- **Lesson 7: Rule Basics**
- **Lesson 8: System Correlation Rules**
- **Lunch**
- **Lesson 9: Defining Application Classes**
- **Lesson 10: Working with Variables**
- **Lesson 11: Cisco Security Agent Analysis**
- **Lesson 12: Using Event Logs and Generating Reports**

© 2005 Cisco Systems, Inc. All rights reserved.

HIPS v2.0—0-6

Participant Responsibilities

Cisco.com

Student responsibilities

- Complete prerequisites.
- Participate in lab exercises.
- Ask questions.
- Provide feedback.



© 2005 Cisco Systems, Inc. All rights reserved.

HIPS v2.0--0-7

General Administration

Cisco.com

Class-related

- Sign-in sheet
- Length and times
- Break and lunch room locations
- Attire

Facilities-related

- Participant materials
- Site emergency procedures
- Restrooms
- Telephones/faxes

© 2005 Cisco Systems, Inc. All rights reserved.

HIPS v2.0--0-8

Graphic Symbols

Cisco.com



IOS Router




PIX Firewall



VPN 3000



IDS Sensor



Catalyst 6500
w/ IDS Module



IOS Firewall



Network
Access Server



Policy Manager



CA
Server



PC



Laptop



Server
Web, FTP, etc.



Hub



Modem



Ethernet Link



VPN Tunnel



Network
Cloud

© 2005 Cisco Systems, Inc. All rights reserved.

HIPS v2.0—0.9

Participant Introductions

Cisco.com

- **Your name**
- **Your company**
- **Prerequisite skills**
- **Brief history**
- **Objective**



© 2005 Cisco Systems, Inc. All rights reserved.

HIPS v2.0—0.10


Cisco Security Career Certifications

Cisco.com

**Expand Your Professional Options —
and Advance Your Career**

Cisco Certified Security Professional (CCSP) Certification

Professional-level recognition in designing
and implementing Cisco security solutions



Required Exam	Recommended Training Through Cisco Learning Partners
642-501	Securing Cisco IOS Networks
642-511	Cisco Secure Virtual Private Networks
642-531	Cisco Secure Intrusion Detection System
642-521	Cisco Secure PIX Firewall Advanced
642-541	Cisco SAFE Implementation

www.cisco.com/go/training

© 2005 Cisco Systems, Inc. All rights reserved.

HIPS v2.0—0-11

Cisco Security Career Certifications

Cisco.com

**Enhance Your Cisco Certifications —
and Validate Your Areas of Expertise**

Cisco Firewall, VPN, and IDS Specialists

Cisco Firewall Specialist



Required Exam	Recommended Training Through Cisco Learning Partners
Prerequisite: Valid CCNA certification	
642-501	Securing Cisco IOS Networks
642-521	Cisco Secure PIX Firewall Advanced

Cisco VPN Specialist



Required Exam	Recommended Training Through Cisco Learning Partners
Prerequisite: Valid CCNA certification	
642-501	Securing Cisco IOS Networks
642-511	Cisco Secure Virtual Private Networks

Cisco IDS Specialist



Required Exam	Recommended Training Through Cisco Learning Partners
Prerequisite: Valid CCNA certification	
642-501	Securing Cisco IOS Networks
642-531	Cisco Secure Intrusion Detection System

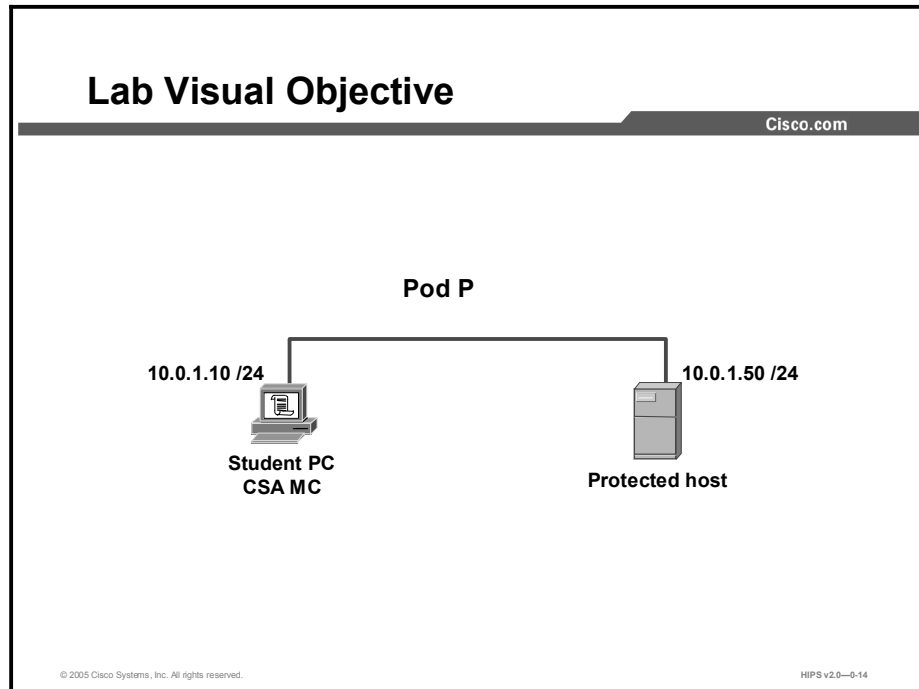
www.cisco.com/go/training

© 2005 Cisco Systems, Inc. All rights reserved.

HIPS v2.0—0-12

Lab Topology Overview

This topic explains the lab topology that is used in this course.



Each student will be assigned a server to run the Cisco Security Agent Management Console (CSA MC) and another server to be protected by Cisco Security Agent (CSA) software. In general, you will be creating security policies on the CSA MC and deploying them to the protected host.

Lesson 1

Security Fundamentals

Overview

This lesson describes security fundamentals. It includes the following topics:

- Objectives
- Need for Network Security
- Network Security Policy
- Primary Network Threats and Attacks
- Reconnaissance Attacks and Mitigation
- Access Attacks and Mitigation
- Denial of Service Attacks and Mitigation
- Worm, Virus, and Trojan Horse Attacks and Mitigation
- Management Protocols and Functions
- Summary

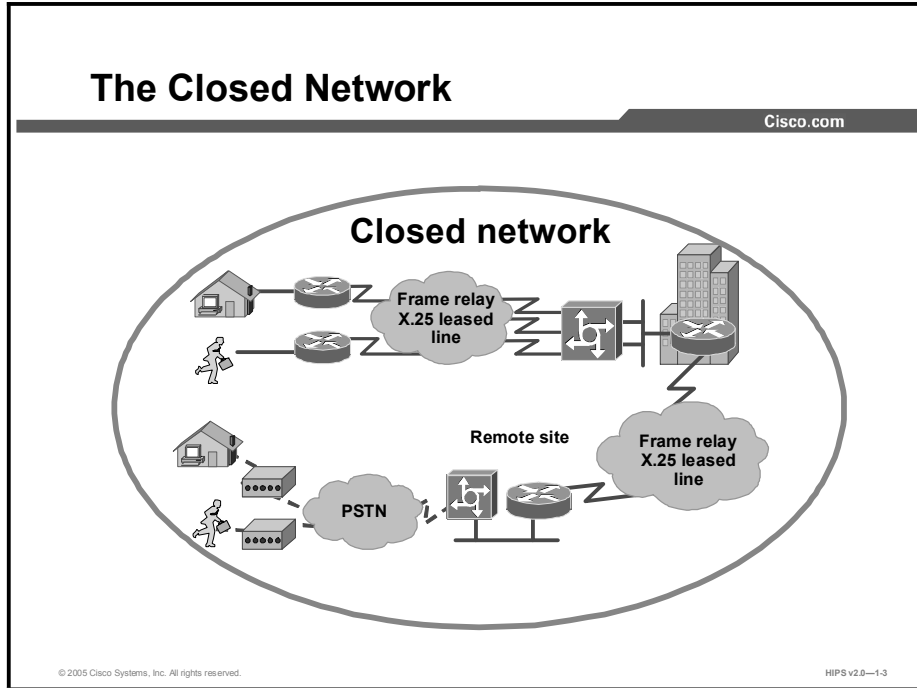
Objectives

Upon completion of this lesson, you will meet the following objectives:

- Describe the need for network security
- Identify the components of a complete security policy
- Explain security as an ongoing process
- Describe the four types of security threats
- Describe the four primary attack categories
- Describe the types of attacks associated with each primary attack category and their mitigation methods
- Describe the configuration management and management protocols and the recommendations for securing them

Need for Network Security

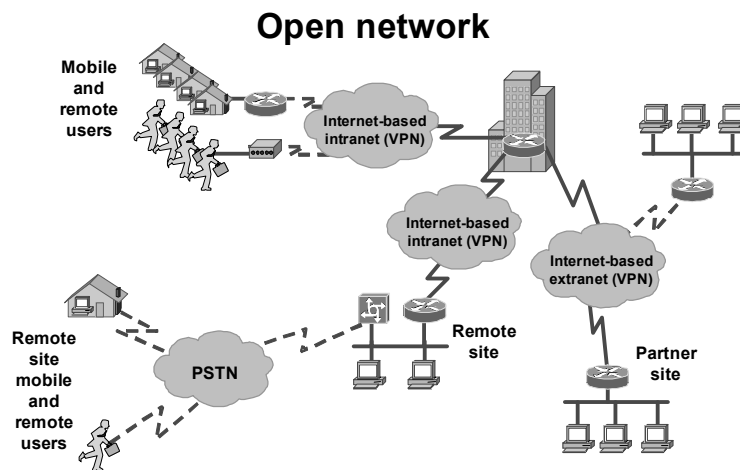
Over the past few years, Internet-enabled business, or e-business, has drastically improved efficiency and revenue growth of companies. E-business applications such as e-commerce, supply-chain management, and remote access enable companies to streamline processes, reduce operating costs, and increase customer satisfaction. Such applications require mission-critical networks that accommodate voice, video, and data traffic, and these networks must be scalable to support increasing numbers of users and the need for greater capacity and performance. As networks enable more and more applications and are available to more and more users, however, they become ever more vulnerable to a wider range of security threats. To combat those threats and ensure that e-business transactions are not compromised, security technology must play a major role in today's networks.



The closed network typically consists of a network designed and implemented in a corporate environment, and it provides connectivity only to known parties and sites without connecting to public networks. Networks were designed this way in the past and thought to be reasonably secure because there was no outside connectivity.

The Network Today

Cisco.com



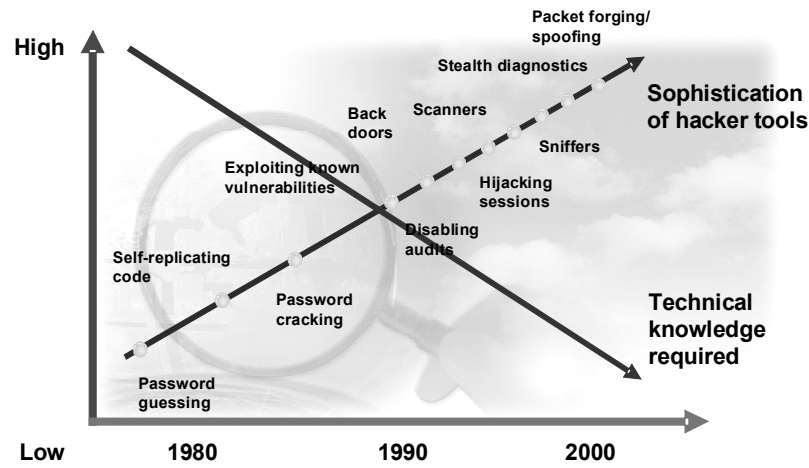
© 2005 Cisco Systems, Inc. All rights reserved.

HIPS v2.0-1.4

The networks of today are designed with availability to the Internet and public networks, which is a major requirement. Most of today's networks have several access points to other networks both public and private; therefore, securing these networks has become fundamentally important.

Threat Capabilities—More Dangerous and Easier to Use

Cisco.com



© 2005 Cisco Systems, Inc. All rights reserved.

HIPS v2.0—1.5

With the development of large open networks there has been a huge increase in security threats in the past 20 years. Not only have hackers discovered more vulnerabilities, but the tools used to hack a network have become simpler and the technical knowledge required has decreased. Downloadable applications that require little or no hacking knowledge to implement are available. There are also applications intended for troubleshooting a network that when used improperly can pose severe threats.

The Changing Role of Security

Cisco.com

As businesses become more open to supporting Internet-powered initiatives such as e-commerce, customer care, supply-chain management, and extranet collaboration, network security risks are also increasing.



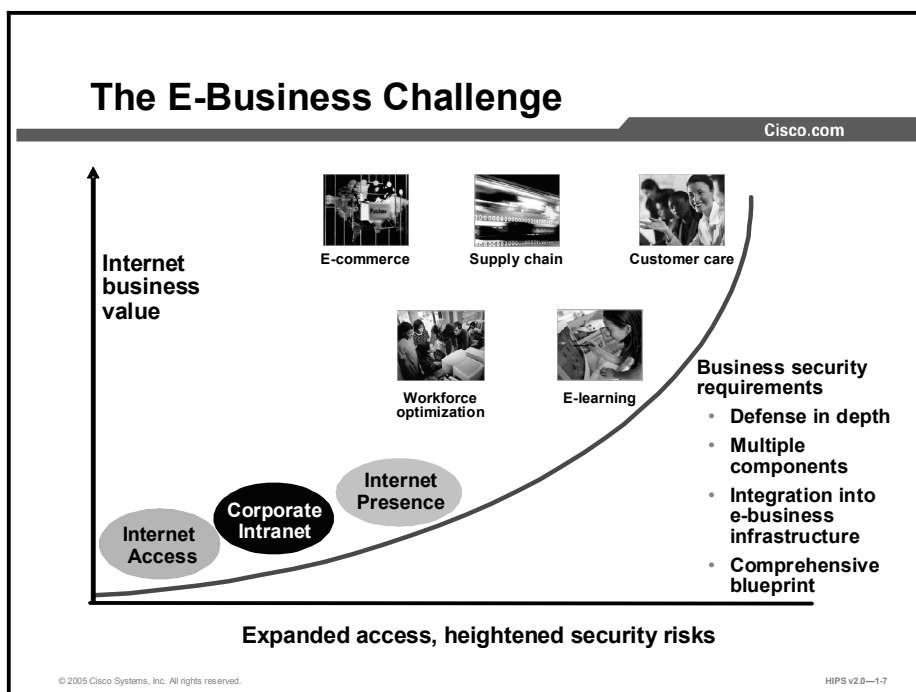
© 2005 Cisco Systems, Inc. All rights reserved.

HIPS v2.0-1-6

Security has moved to the forefront of network management and implementation. The survival of many businesses depends on allowing open access to network resources while ensuring that the data and resources are as secure as possible.

Security is becoming more important because of the following:

- **Requirements of e-business:** The importance of e-business and the need for private data to traverse public networks has increased the need for network security.
- **Requirements of communicating and doing business safely in potentially unsafe environments:** Today's business environment requires communication with many public networks and systems, which produces the need for as much security as possible.
- **Requirements of developing and implementing a corporate-wide security policy:** Establishing a security policy should be the first step in migrating a network to a secure infrastructure.



Security must be a fundamental component of any e-business strategy. As enterprise network managers open their networks to more users and applications, they also expose these networks to greater risk. The result has been an increase in business security requirements.

The Internet has radically shifted expectations of companies' abilities to build stronger relationships with customers, suppliers, partners, and employees. Driving companies to become more agile and competitive, e-business is giving birth to exciting new applications for e-commerce, supply-chain management, customer care, workforce optimization, and e-learning—applications that streamline and improve processes, speed up turnaround times, lower costs, and increase user satisfaction.

E-business requires mission-critical networks that accommodate ever-increasing constituencies and demands for greater capacity and performance. These networks also need to handle voice, video, and data traffic as networks converge into multiservice environments.

Legal and Governmental Policy Issues

Cisco.com

- **Many governments have formed cross-border task forces to deal with privacy issues.**
- **The outcome of international privacy efforts is expected to take several years to develop.**
- **National laws regarding privacy are expected to continue to evolve worldwide.**



© 2005 Cisco Systems, Inc. All rights reserved.

HIPS v2.0-1-8


As concerns about privacy increase, many governments have formed cross-border task forces to deal with privacy issues. International privacy efforts are expected to take several years to develop and even longer to implement globally. National laws regarding privacy are expected to continue to evolve worldwide.

Network Security Is a Continuous Process

Cisco.com

Network security is a continuous process built around a security policy:

- Step 1: Secure
- Step 2: Monitor
- Step 3: Test
- Step 4: Improve



© 2005 Cisco Systems, Inc. All rights reserved.

HIPS v2.0—1.0

After setting appropriate policies, a company or organization must methodically consider security as part of normal network operations. This process could be as simple as configuring routers so that they do not accept unauthorized addresses or services, or as complex as installing firewalls, intrusion detection systems (IDSs), centralized authentication servers, and encrypted Virtual Private Networks (VPNs). Network security is a continuing process:

- **Secure:** The following are methods used to secure a network:
 - Authentication
 - Encryption
 - Firewalls
 - Vulnerability patching
- **Monitor:** To ensure that a network remains secure, organizations need to monitor the state of security preparation. Network vulnerability scanners can proactively identify areas of weakness, and IDSs can monitor and respond to security events as they occur. Using security monitoring solutions, organizations can obtain unprecedented visibility into both the network data stream and the security posture of the network.
- **Test:** Testing security is as important as monitoring. Without testing the security solutions in place, existing or new attacks will go unnoticed. The hacker community is an ever-changing environment. You can perform this testing or outsource it to a third party, such as the Cisco Security Posture Assessment (SPA) group.
- **Improve:** Monitoring and testing provides the data necessary to improve network security. Administrators and engineers should use the information from the monitor and test phases to improve the security implementation as well as to adjust the security policy as vulnerabilities and risks are identified.

Network Security Policy

A security policy can be as simple as an acceptable use policy for network resources, or it can be several hundred pages in length and detail every element of connectivity and associated policies.

What Is a Security Policy?

Cisco.com

“A security policy is a formal statement of the rules by which people who are given access to an organization’s technology and information assets must abide.”

– RFC 2196, Site Security Handbook

© 2005 Cisco Systems, Inc. All rights reserved. HIPS v2.0—1-11

According to the *Site Security Handbook* (RFC 2196), “A security policy is a formal statement of the rules by which people who are given access to an organization’s technology and information assets must abide.” It further states, “A security policy is essentially a document summarizing how the corporation will use and protect its computing and network resources.”

Why Create a Security Policy?

Cisco.com

- **To create a baseline of your current security posture**
- **To set the framework for security implementation**
- **To define allowed and not-allowed behaviors**
- **To help determine necessary tools and procedures**
- **To communicate consensus and define roles**
- **To define how to handle security incidents**
- **To inform users of their responsibilities**
- **To define assets and the way to use them**
- **To state the ramifications of misuse**

© 2005 Cisco Systems, Inc. All rights reserved.

HIPS v2.0—1-12

Security policies provide many benefits and are worth the time and effort needed to develop them. Developing a security policy has the following benefits:

- Provides a process for auditing existing network security.
- Provides a general security framework for implementing network security.
- Defines which behavior is and is not allowed.
- Helps determine which tools and procedures are needed for the organization.
- Helps communicate consensus among a group of key decision makers and define responsibilities of users and administrators.
- Defines a process for handling network security incidents.
- Enables global security implementation and enforcement. Computer security is now an enterprise-wide issue, and computing sites are expected to conform to the network security policy.
- Creates a basis for legal action if necessary.

What Should the Security Policy Contain?

Cisco.com

- **Statement of authority and scope**
- **Acceptable use policy**
- **Identification and authentication policy**
- **Internet use policy**
- **Campus access policy**
- **Remote access policy**
- **Incident handling procedure**

© 2005 Cisco Systems, Inc. All rights reserved.

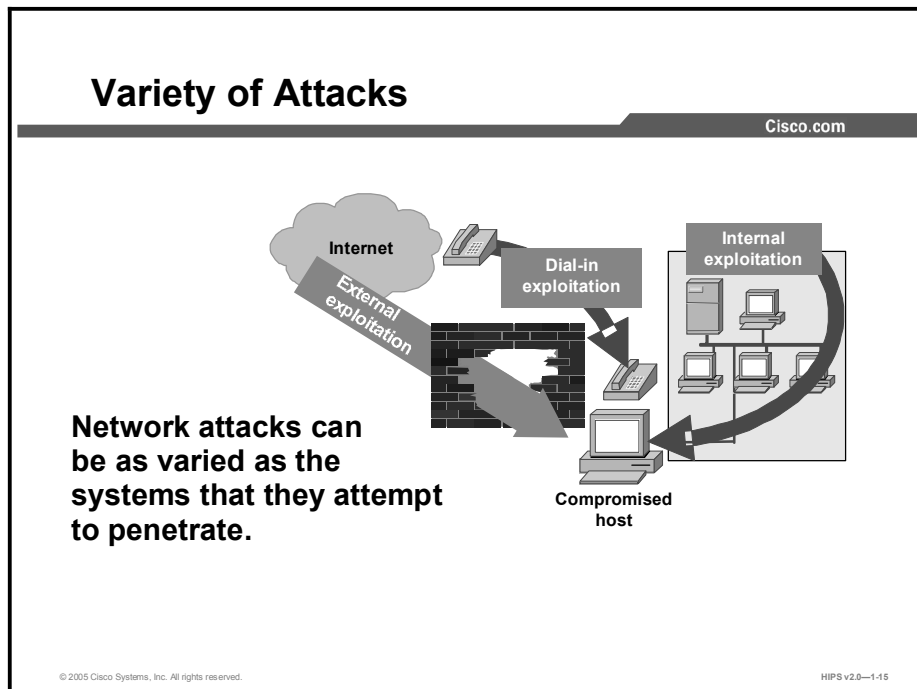
HIPS v2.0—1-13

The following are some key policy components:

- **Statement of authority and scope:** This topic specifies who sponsors the security policy and what areas the policy covers.
- **Acceptable use policy:** This topic specifies what the company will and will not allow regarding its information infrastructure.
- **Identification and authentication policy:** This topic specifies what technologies, equipment, or combination of the two the company will use to ensure that only authorized individuals have access to its data.
- **Internet access policy:** This topic specifies what the company considers ethical and proper use of its Internet access capabilities.
- **Campus access policy:** This topic specifies how on-campus users will use the company's data infrastructure.
- **Remote access policy:** This topic specifies how remote users will access the company's data infrastructure.
- **Incident handling procedure:** This topic specifies how the company will create an incident response team and the procedures it will use during and after an incident.

Primary Network Threats and Attacks

This topic provides an overview of primary network threats and attacks.



Without proper protection, any part of any network can be susceptible to attacks or unauthorized activity. Routers, switches, and hosts can all be violated by professional hackers, company competitors, or even internal employees. In fact, according to several studies, more than half of all network attacks are waged internally. The Computer Security Institute (CSI) in San Francisco estimates that between 60 and 80 percent of network misuse comes from inside the enterprises where the misuse has taken place. To determine the best ways to protect against attacks, information technology managers should understand the many types of attacks that can be instigated and the damage that these attacks can cause to e-business infrastructures.

Network Security Threats

Cisco.com

There are four general categories of security threats to the network:

- **Unstructured threats**
- **Structured threats**
- **External threats**
- **Internal threats**

© 2005 Cisco Systems, Inc. All rights reserved.

HIPS v2.0—1-16

There are four general threats to network security:

- **Unstructured threats:** These threats come primarily from random hackers using various common tools, such as malicious shell scripts, password crackers, credit card number generators, and dialer daemons. Although hackers in this category may have malicious intent, many are more interested in the intellectual challenge of cracking safeguards than in creating havoc.
- **Structured threats:** These threats are created by hackers who are highly motivated and technically competent. Typically, such hackers act alone or in small groups to understand, develop, and use sophisticated hacking techniques to penetrate unsuspecting businesses. These groups are often involved in the major fraud and theft cases reported to law enforcement agencies. Occasionally, such hackers are hired by organized crime, industry competitors, or state-sponsored intelligence collection organizations.
- **External threats:** These threats consist of structured and unstructured threats originating from an external source. These threats may have malicious and destructive intent, or they may simply be errors that generate a threat.
- **Internal threats:** These threats typically involve disgruntled former or current employees. Although internal threats may seem more ominous than threats from external sources, security measures are available for reducing vulnerabilities to internal threats and responding when attacks occur.

The Four Primary Attack Categories

Cisco.com

All of the following can be used to compromise your system:

- **Reconnaissance attacks**
- **Access attacks**
- **Denial of service attacks**
- **Worms, viruses, and Trojan horses**

© 2005 Cisco Systems, Inc. All rights reserved.

HIPS v2.0—1-17

There are four types of network attacks:

- **Reconnaissance attacks:** An intruder attempts to discover and map systems, services, and vulnerabilities.
- **Access attacks:** An intruder attacks networks or systems to retrieve data, gain access, or escalate access privileges.
- **Denial of service (DoS) attacks:** An intruder attacks your network in a way that damages or corrupts your computer system or denies you and others access to your networks, systems, or services.
- **Worms, viruses, and Trojan horses:** Malicious software is inserted onto a host in order to damage a system, corrupt a system, replicate itself, or deny services or access to networks, systems, or services.


Reconnaissance Attacks and Mitigation

This topic describes reconnaissance attacks and their mitigation.

Reconnaissance Attacks

Cisco.com

Reconnaissance refers to the overall act of learning information about a target network by using readily available information and applications.



© 2005 Cisco Systems, Inc. All rights reserved. HPS v2.0—1-19

Reconnaissance is the unauthorized discovery and mapping of systems, services, or vulnerabilities. It is also known as information gathering and, in most cases, precedes an actual access or DoS attack. The malicious intruder typically conducts a ping sweep of the target network first to determine which IP addresses respond. After this has been accomplished, the intruder determines which services or ports are active on the live IP addresses. From this information, the intruder queries the ports to determine the application type and version as well as the type and version of the operating system running on the target host.

Reconnaissance is somewhat analogous to a thief casing a neighborhood for vulnerable homes to break into, such as an unoccupied residence, a house with an easy-to-open door or window, and so on. In many cases the intruders go as far as “rattling the door handle,” not to go in immediately if the network is open, but to discover vulnerable services that they can exploit later when there is less likelihood that anyone is looking.

Reconnaissance attacks can consist of the following:

- Packet sniffers
- Port scans
- Ping sweeps
- Internet information queries

Packet Sniffers

Cisco.com



A packet sniffer is a software application that uses a network adapter card in promiscuous mode to capture all network packets. The following are packet sniffer features:

- **Packet sniffers exploit information passed in clear text. Protocols that pass information in the clear include the following:**
 - Telnet
 - FTP
 - SNMP
 - POP
 - HTTP
- **Packet sniffers must be on the same collision domain.**
- **Packet sniffers can be general purpose or can be designed specifically for attack.**

© 2005 Cisco Systems, Inc. All rights reserved.

HIPS v2.0—1-20

A packet sniffer is a software application that uses a network adapter card in promiscuous mode (a mode in which the network adapter card sends all packets received on the physical network wire to an application for processing) to capture all network packets that are sent across a LAN.

Several network applications distribute network packets in clear text; that is, the information sent across the network is not encrypted. Because the network packets are not encrypted, they can be processed and understood by any application that can pick them up off the network and process them.

A network protocol specifies how packets are identified and labeled, which enables a computer to determine whether a packet is intended for it. Because the specifications for network protocols, such as TCP/IP, are widely published, a third party can easily interpret the network packets and develop a packet sniffer. (The real threat today results from the numerous freeware and shareware packet sniffers that are available, which do not require the user to understand anything about the underlying protocols.)

Packet Sniffer Attack Mitigation

Cisco.com



The following techniques and tools can be used to mitigate sniffer attacks:

- **Authentication**—A first option for defense against packet sniffers is to use strong authentication, such as one-time passwords.
- **Switched infrastructure**—Deploy a switched infrastructure to counter the use of packet sniffers in your environment.
- **Antisniffer tools**—Use these tools to employ software and hardware designed to detect the use of sniffers on a network.
- **Cryptography**—The most effective method for countering packet sniffers does not prevent or detect packet sniffers, but rather renders them irrelevant.

© 2005 Cisco Systems, Inc. All rights reserved.

HIPS v2.0—1-21

The following techniques and tools can be used to mitigate packet sniffer attacks:

- **Authentication:** Using strong authentication is a first option for defense against packet sniffers. Strong authentication can be broadly defined as a method of authenticating users that cannot be easily circumvented. A common example of strong authentication is one-time passwords (OTPs).


An OTP is a type of two-factor authentication. Two-factor authentication involves using something that you have combined with something that you know. Automated teller machines (ATMs) use two-factor authentication. A customer needs both an ATM card and a personal identification number (PIN) to make transactions. With OTPs you need a PIN and your token card to authenticate to a device or software application. A token card is a hardware or software device that generates new, seemingly random passwords at specified intervals (usually 60 seconds). A user combines that password with a PIN to create a unique password that works for one instance of authentication. If a hacker learns that password by using a packet sniffer, the information is useless because the password has already expired. Note that this mitigation technique is effective only against a sniffer implementation that is designed to grab passwords. Sniffers deployed to learn sensitive information (such as e-mail messages) will still be effective.

- **Switched infrastructure:** This technique can be used to counter the use of packet sniffers in your network environment. For example, if an entire organization deploys switched Ethernet, hackers can gain access only to the traffic that flows on the specific port to which they connect. A switched infrastructure does not eliminate the threat of packet sniffers, but it can greatly reduce their effectiveness.
- **Antisniffer tools:** Software and hardware designed to detect the use of sniffers on a network can be employed. Such software and hardware does not completely eliminate the threat, but like many network security tools, they are part of the overall system. These so-called antisniffers detect changes in the response time of hosts to determine whether the hosts are processing more traffic than their own. One such network security software tool, which is available from Security Software Technologies, is called AntiSniff.

- **Cryptography:** Rendering packet sniffers irrelevant is the most effective method for countering packet sniffers, even more effective than preventing or detecting packet sniffers. If a communication channel is cryptographically secure, the only data that a packet sniffer will detect is cipher text (a seemingly random string of bits) and not the original message. The Cisco deployment of network-level cryptography is based on IPSec, which is a standard method for networking devices to communicate privately using IP. Other cryptographic protocols for network management include Secure Shell Protocol (SSH) and Secure Socket Layer (SSL).

Port Scans and Ping Sweeps

Cisco.com



These attacks can attempt to:

- **Identify all services on the network**
- **Identify all hosts and devices on the network**
- **Identify the operating systems on the network**
- **Identify vulnerabilities on the network**

© 2005 Cisco Systems, Inc. All rights reserved.

HIPS v2.0—1-22

Port scans and ping sweeps are typically applications built to run various tests against a host or device in order to identify vulnerable services. The information is gathered by examining IP addressing and port or banner data from both TCP and User Datagram Protocol (UDP) ports.

Port Scan and Ping Sweep Attack Mitigation

Cisco.com

- **Port scans and ping sweeps cannot be prevented entirely.**
- **IDSs at the network and host levels can usually notify an administrator when a reconnaissance attack such as a port scan or ping sweep is under way.**

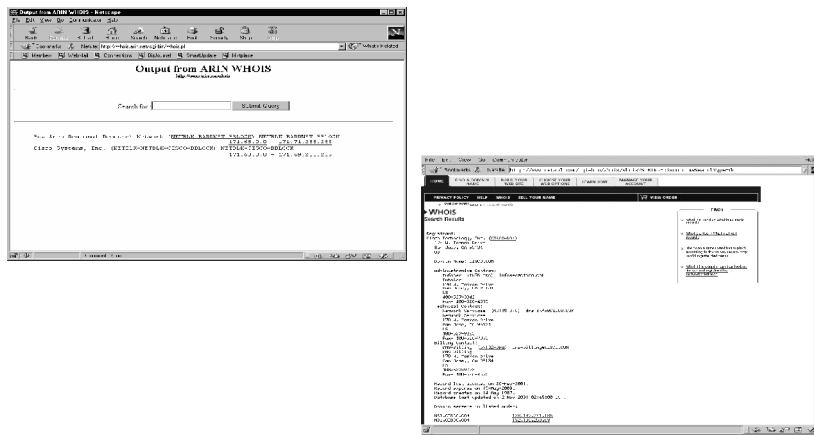
© 2005 Cisco Systems, Inc. All rights reserved.

HIPS v2.0—1-23

If ICMP echo and echo reply are turned off on edge routers, for example, ping sweeps can be stopped, but at the expense of network diagnostic data. However, port scans can be run easily without full ping sweeps; they simply take longer because they need to scan IP addresses that might not be live. IDSs at the network and host levels can usually notify an administrator when a reconnaissance attack is under way. This warning allows the administrator to better prepare for the coming attack or to notify the Internet Service Provider (ISP) of the system that is launching the reconnaissance probe.

Internet Information Queries

Cisco.com



Sample domain name query

© 2005 Cisco Systems, Inc. All rights reserved.

HIPS v2.0—1-24

The figure demonstrates how existing Internet tools can be used for network reconnaissance (for example, an IP address query or a Domain Name System [DNS] query).

DNS queries can reveal information such as who owns a particular domain and what addresses have been assigned to that domain. Ping sweeps of the addresses revealed by the DNS queries can present a picture of the live hosts in a particular environment. After such a list is generated, port scanning tools can cycle through all well-known ports to provide a complete list of all services running on the hosts that were discovered by the ping sweep. Finally, hackers can examine the characteristics of the applications that are running on the hosts. This step can lead to specific information that is useful when the hacker attempts to compromise that service.

IP address queries can reveal information such as who owns a particular IP address or range of addresses and what domain is associated with them.

Access Attacks and Mitigation


This topic describes specific access attacks and their mitigation.

Access Attacks

Cisco.com

In access attacks, intruders typically attack networks or systems to:

- Retrieve data
- Gain access
- Escalate their access privileges



© 2005 Cisco Systems, Inc. All rights reserved. HIPS v2.0—1-28

Access attacks exploit known vulnerabilities in authentication services, FTP services, and web services to gain entry to web accounts, confidential databases, and other sensitive information. Access attacks can consist of the following:


- Password attacks
- Trust exploitation
- Port redirection
- Man-in-the-middle attacks

Password Attacks

Cisco.com

Hackers can implement password attacks using several methods:

- Brute-force attacks
- Trojan horse programs
- IP spoofing
- Packet sniffers



© 2005 Cisco Systems, Inc. All rights reserved.

HIPS v2.0—1-27

Password attacks can be implemented using several methods, including brute-force attacks, Trojan horse programs, IP spoofing, and packet sniffers. Although packet sniffers and IP spoofing can yield user accounts and passwords, password attacks usually refer to repeated attempts to identify a user account, password, or both. These repeated attempts are called brute-force attacks.

Hackers often perform brute-force attacks using a program that runs across the network and attempts to log in to a shared resource, such as a server. When an attacker gains access to a resource, he or she has the same access rights as the user whose account has been compromised. If this account has sufficient privileges, the attacker can create a back door for future access, without concern for any status and password changes to the compromised user account.

Password Attack Mitigation

Cisco.com

The following are password attack mitigation techniques:

- **Do not allow users to use the same password on multiple systems.**
- **Disable accounts after a certain number of unsuccessful login attempts.**
- **Do not use plain text passwords. An OTP or a cryptographic password is recommended.**
- **Use “strong” passwords. Strong passwords are at least eight characters long and contain uppercase letters, lowercase letters, numbers, and special characters.**
- **Force periodic password changes.**

© 2005 Cisco Systems, Inc. All rights reserved.

HIPS v2.0—1-28

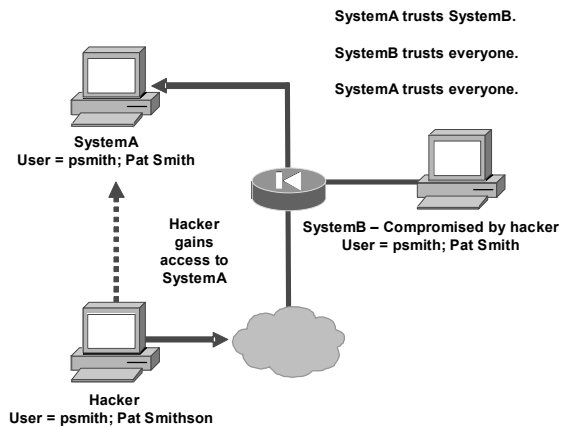
The following are password attack mitigation techniques:

- Do not allow users to have the same password on multiple systems. Most users will use the same password for each system they access, and personal system passwords will often be the same as well.
- Disable accounts after a specific number of unsuccessful logins. This practice helps to prevent continuous password attempts.
- Do not use plain-text passwords. Use of either an OTP or encrypted password is recommended.
- Use “strong” passwords. Many systems now provide strong password support and can restrict users to the use of strong passwords only. Strong passwords are at least eight characters long and contain uppercase letters, lowercase letters, numbers, and special characters.
- Force periodic password changes. Forcing users to periodically change their passwords can reduce the risk of password discovery.

Trust Exploitation

Cisco.com

- A hacker leverages existing trust relationships.
- Several trust models exist.
 - Windows
 - Domains
 - Active directory
 - Linux and UNIX
 - NFS
 - NIS+



© 2005 Cisco Systems, Inc. All rights reserved.

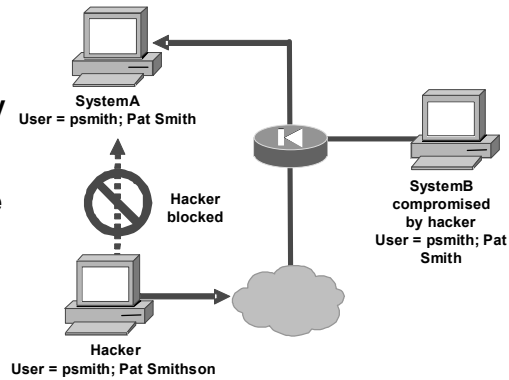
HIPS v2.0—1-29

Although it is not an attack in itself, trust exploitation refers to an individual's taking advantage of a trust relationship within a network. The classic example is a perimeter network connection from a corporation. These network segments often house DNS, Simple Mail Transfer Protocol (SMTP), and Hypertext Transfer Protocol (HTTP) servers. Because they all reside on the same segment, a compromise of one system can lead to the compromise of other systems if those other systems in turn trust systems attached to the same network. Another example is a system on the outside of a firewall that has a trust relationship with a system on the inside of a firewall. When the outside system is compromised, the attacker can leverage that trust relationship to attack the inside network.

Trust Exploitation Attack Mitigation

Cisco.com

- **Systems on the outside of a firewall should never be absolutely trusted by systems on the inside of a firewall.**
- **Such trust should be limited to specific protocols and should be validated by something other than an IP address where possible.**



© 2005 Cisco Systems, Inc. All rights reserved.

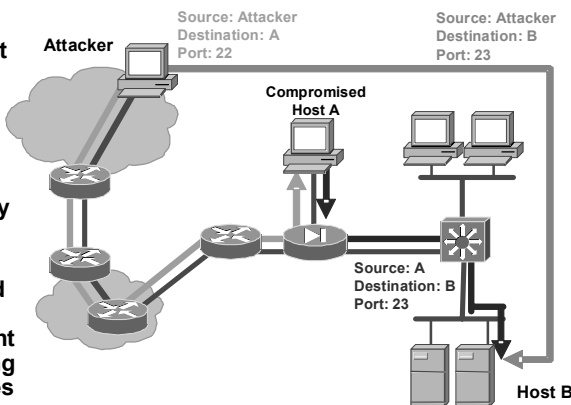
HIPS v2.0—1.30

You can mitigate attacks based on trust exploitation through tight constraints on trust levels within a network. Systems on the outside of a firewall should never be absolutely trusted by systems on the inside of a firewall. Such trust should be limited to specific protocols and should be authenticated by something other than an IP address where possible.

Port Redirection

Cisco.com

- Port redirection is a type of trust exploitation attack that uses a compromised host to pass traffic through a firewall that would otherwise be dropped.
- It is mitigated primarily through the use of proper trust models.
- Antivirus software and host-based IDS can help detect and prevent a hacker from installing port redirection utilities on the host.



© 2005 Cisco Systems, Inc. All rights reserved.

HIPS v2.0—1-31

Port redirection attacks are a type of trust exploitation attack that uses a compromised host to pass traffic through a firewall that would otherwise be dropped. Consider a firewall with three interfaces and a host on each interface. The host on the outside can reach the host on the public services segment (commonly referred to as a Demilitarized Zone [DMZ]), but not the host on the inside. The host on the public services segment can reach the host on both the outside and the inside. If hackers were able to compromise the public services segment host, they could install software to redirect traffic from the outside host directly to the inside host. Though neither communication violates the rules implemented in the firewall, the outside host has now achieved connectivity to the inside host through the port redirection process on the public services host. An example of an application that can provide this type of access is netcat.

Port redirection can be mitigated primarily through the use of proper trust models, which are network-specific (as mentioned earlier). Assuming a system is under attack, a host-based IDS can help detect a hacker and prevent installation of such utilities on a host.

Man-in-the-Middle Attacks

Cisco.com



- A man-in-the-middle attack requires that the hacker have access to network packets that come across a network.
- A man-in-the-middle attack is implemented using the following:
 - Network packet sniffers
 - Routing and transport protocols
- Possible uses for man-in-the-middle attacks include the following:
 - Theft of information
 - Hijacking of an ongoing session
 - Traffic analysis
 - DoS
 - Corruption of transmitted data
 - Introduction of new information into network sessions

© 2005 Cisco Systems, Inc. All rights reserved.


HIPS v2.0—1.32

A man-in-the-middle attack requires that the attacker have access to network packets that come across the network. Such attacks are often implemented using network packet sniffers and routing and transport protocols. The possible uses of such attacks are theft of information, hijacking of an ongoing session to gain access to your internal network resources, traffic analysis to derive information about your network and its users, denial of service, corruption of transmitted data, and introduction of new information into network sessions.

An example of someone who might commit a man-in-the-middle attack is a person working for your ISP who can gain access to all network packets transferred between your network and any other network.

Man-in-the-Middle Attack Mitigation

Cisco.com



Man-in-the-middle attacks can be effectively mitigated only through the use of cryptography (encryption).

© 2005 Cisco Systems, Inc. All rights reserved.

HIPS v2.0—1-33

Man-in-the-middle attack mitigation is achieved, as shown in the figure, by encrypting traffic in an IPsec tunnel, which allows the hacker to see only cipher text.


Denial of Service Attacks and Mitigation

This topic describes specific DoS attacks and their mitigation.

Denial of Service Attacks

Cisco.com

Denial of service attacks occur when an intruder attacks your network in a way that damages or corrupts your computer system or denies you and others access to your networks, systems, or services.



© 2005 Cisco Systems, Inc. All rights reserved. HIPS v2.0—1-35

Certainly the most publicized form of attack, DoS attacks are also among the most difficult to eliminate completely. Even within the hacker community, DoS attacks are regarded as trivial and considered bad form because they require so little effort to execute. Still, because of their ease of implementation and potentially significant damage, DoS attacks deserve special attention from security administrators. If you are interested in learning more about DoS attacks, researching the methods employed by some of the better-known attacks can be useful. DoS attacks can consist of the following:

- IP spoofing
- Distributed denial of service (DDoS)

IP Spoofing

Cisco.com

- **IP spoofing occurs when a hacker inside or outside a network impersonates the conversations of a trusted computer.**
- **Two general techniques are used during IP spoofing:**
 - **A hacker uses an IP address that is within the range of trusted IP addresses.**
 - **A hacker uses an authorized external IP address that is trusted.**
- **Uses for IP spoofing include the following:**
 - **IP spoofing is usually limited to the injection of malicious data or commands into an existing stream of data.**
 - **If a hacker changes the routing tables to point to the spoofed IP address, then the hacker can receive all the network packets that are addressed to the spoofed address and reply, just as any trusted user can.**

© 2005 Cisco Systems, Inc. All rights reserved.

HIPS v2.0—1-38

An IP spoofing attack occurs when an attacker outside your network pretends to be a trusted computer. The attacker has two ways of doing this: either by using an IP address that is within the range of IP addresses for your network or by using an authorized external IP address that you trust and to which you wish to provide access to specified resources on your network.

Normally, an IP spoofing attack is limited to the injection of data or commands into an existing stream of data passed between a client and server application or a peer-to-peer network connection. To enable bidirectional communication, the attacker must change all routing tables to point to the spoofed IP address. Another approach the attacker could take is simply not to worry about receiving any response from the applications. For example, if an attacker is attempting to get a system to mail him or her a sensitive file, application responses are unimportant.

If an attacker manages to change the routing tables to point to the spoofed IP address, he or she can receive all the network packets that are addressed to the spoofed address and reply just as any trusted user can. Like packet sniffers, IP spoofing use is not restricted to people who are external to the network.

IP spoofing can also provide access to user accounts and passwords, although this use is not common, and it can be used in other ways; for example, an attacker can emulate one of your internal users in ways that prove embarrassing for your organization. The attacker could send e-mail messages to business partners that appear to have originated from someone within your organization. Such attacks are easier when an attacker has a user account and password, but they are possible when simple spoofing attacks are combined with knowledge of messaging protocols.

IP Spoofing Attack Mitigation

Cisco.com

The threat of IP spoofing can be reduced, but not eliminated, through the following measures:

- **Access control**—The most common method for preventing IP spoofing is to properly configure access control.
- **RFC 2827 filtering**—Prevent any outbound traffic on your network that does not have a source address in your organization's own IP range.
- **Additional authentication require additional authentication that does not use IP-based authentication. Examples of this technique include the following:**
 - **Cryptographic (recommended)**
 - **Strong, two-factor, one-time passwords**

© 2005 Cisco Systems, Inc. All rights reserved.

HIPS v2.0—1.37

The threat of IP spoofing can be reduced, but not eliminated, through the following measures:

- **Access control:** The most common method for preventing IP spoofing is proper configuration of access control. To reduce the effectiveness of IP spoofing, configure access control to deny any traffic from the external network that has a source address that should reside on the internal network. Note that this helps prevent spoofing attacks only if the internal addresses are the only trusted addresses. If some external addresses are trusted, this method is not effective.
- **RFC 2827 filtering:** You can prevent users of your network from spoofing other networks (and be a good Internet citizen at the same time) by preventing any outbound traffic on your network that does not have a source address in your organization's own IP range.

This filtering denies any traffic that does not have the source address that is expected on a particular interface. For example, if an ISP is providing a connection to the IP address 15.1.1.0/24, the ISP could filter traffic so that only traffic sourced from address 15.1.1.0/24 can enter the ISP router from that interface. Note that unless all ISPs implement this type of filtering, its effectiveness is significantly reduced.

- **Additional authentication:** The most effective method for mitigating the threat of IP spoofing is the same as the most effective method for mitigating the threat of packet sniffers; namely, eliminating its effectiveness. IP spoofing can function correctly only when devices use authentication based on IP addresses; therefore, if you use additional authentication methods, IP spoofing attacks are irrelevant. Cryptographic authentication is the best form of additional authentication, but when that is not possible, strong two-factor authentication using OTPs can also be effective.

DoS and DDoS Attacks

Cisco.com

DoS attacks focus on making a service unavailable for normal use. They have the following characteristics:

- **Different from most other attacks because they are generally not targeted at gaining access to your network or the information on your network**
- **Require very little effort to execute**
- **Among the most difficult to eliminate completely**

© 2005 Cisco Systems, Inc. All rights reserved.

HIPS v2.0—1-38

DoS attacks are different from most other attacks because they are not targeted at gaining access to your network or the information on your network. These attacks focus on making a service unavailable for normal use, which is typically accomplished by exhausting some resource limitation on the network or within an operating system or application. These attacks require little effort to execute because they typically take advantage of protocol weaknesses or because the attacks are carried out using traffic that would normally be allowed into a network. DoS attacks are among the most difficult to eliminate completely because of the way they use protocol weaknesses and “native” traffic to attack a network.

DoS and DDoS Attack Mitigation

Cisco.com

The threat of DoS attacks can be reduced through the following three methods:

- **Antispoof features**—Proper configuration of antispoof features on routers and firewalls
- **Anti-DoS features**—Proper configuration of anti-DoS features on routers, firewalls, and intrusion detection systems
- **Traffic rate limiting**—Implementation of traffic rate limiting with the ISP of the network

© 2005 Cisco Systems, Inc. All rights reserved.

HIPS v2.0—1.39

When they involve specific network server applications, such as an HTTP server or an FTP server, DoS attacks focus on acquiring and keeping open all the available connections supported by that server, effectively locking valid users of the server or service. DoS attacks can also be implemented using common Internet protocols, such as TCP and ICMP. While most DoS attacks exploit a weakness in the overall architecture of the system being attacked rather than a software bug or security hole, some attacks compromise the performance of your network by flooding the network with undesired, and often useless, network packets and by providing false information about the status of network resources.

The threat of DoS attacks can be reduced through the following three methods:

- **Antispoof features:** Proper configuration of antispoof features on your routers and firewalls can reduce your risk. This configuration includes RFC 2827 filtering at a minimum. If hackers cannot mask their identities, they might not attack.
- **Anti-DoS features:** Proper configuration of anti-DoS features on routers, firewalls, and IDSs can help limit the effectiveness of an attack. These features often involve limits on the amount of half-open connections that a system allows at any given time.
- **Traffic rate limiting:** An organization can implement traffic rate limiting with its ISP. This type of filtering limits the amount of nonessential traffic that crosses network segments at a certain rate. A common example is to limit the amount of ICMP traffic allowed into a network because it is used only for diagnostic purposes. ICMP-based DDoS attacks are common.

Worm, Virus, and Trojan Horse Attacks and Mitigation


This topic describes worm, virus, and Trojan horse attacks and their mitigation.

Worm, Virus, and Trojan Horse Attacks

Cisco.com

The primary vulnerabilities for end user workstations are worm, virus, and Trojan horse attacks.

- A worm executes arbitrary code and installs copies of itself in the infected computer's memory, which infects other hosts.
- A virus is malicious software that is attached to another program to execute a particular unwanted function on a user's workstation.
- A Trojan horse is different from a virus only in that the entire application is written to look like something else, when in fact it is an attack tool.




© 2005 Cisco Systems, Inc. All rights reserved. HIPS v2.0—1-41

The primary vulnerabilities for end-user workstations are worm, virus, and Trojan horse attacks.

- A worm executes arbitrary code and installs copies of itself in the infected computer's memory, which infects other hosts.
- A virus is malicious software that is attached to another program to execute a particular unwanted function on a user's workstation.
- A Trojan horse is different from a virus only in that the entire application is written to look like something else, when in fact it is an attack tool.

Worm Attacks

Cisco.com



© 2005 Cisco Systems, Inc. All rights reserved.

HIPS v2.0—142

The anatomy of a worm attack is as follows:

- **Enabling vulnerability:** A worm installs itself using an exploit vector on a vulnerable system.
- **Propagation mechanism:** After gaining access to devices, a worm replicates and selects new targets.
- **Payload:** Once the device is infected with a worm, the attacker has access to the host—often as a privileged user. Attackers could use a local exploit to escalate their privilege level to administrator.

Typically, worms are self-contained programs that attack a system and try to exploit a vulnerability in the target. Upon successful exploitation of the vulnerability, the worm copies its program from the attacking host to the newly exploited system to begin the cycle again. A virus normally requires a vector to carry the virus code from one system to another. The vector can be a word-processing document, an e-mail message, or an executable program. The key element that distinguishes a computer worm from a computer virus is the requirement for human interaction to facilitate the spread of a virus.

Worm Attack Mitigation

Cisco.com

- **Containment**—Contain the spread of the worm inside your network and within your network. Compartmentalize parts of your network that have not been infected.
- **Inoculation**—Start patching all systems and, if possible, scanning for vulnerable systems.
- **Quarantine**—Track down each infected machine inside your network. Disconnect, remove, or block infected machines from the network.
- **Treatment**—Clean and patch each infected system. Some worms may require complete core system reinstallations to clean the system.

© 2005 Cisco Systems, Inc. All rights reserved.

HIPS v2.0—1-43

Worm attack mitigation requires diligence on the part of system and network administration staff. Coordination between system administration, network engineering, and security operations personnel is critical in responding effectively to a worm incident. The following are the recommended steps for worm attack mitigation:

- Containment
- Inoculation
- Quarantine
- Treatment

Typical incident response methodologies can be divided into six major categories. The following categories are based on the network service provider security (NSP-SEC) incident response methodology:

- **Preparation:** Acquire the resources to respond.
- **Identification:** Identify the worm.
- **Classification:** Classify the type of worm.
- **Traceback:** Trace the worm back to its origin.
- **Reaction:** Isolate and repair the affected systems.
- **Post mortem:** Document and analyze the process used for the future.

Virus and Trojan Horse Attacks

Cisco.com

- **Viruses are malicious software that is attached to another program to execute a particular unwanted function on a user's workstation. End-user workstations are the primary targets.**
- **A Trojan horse is different from a virus only in that the entire application is written to look like something else, when in fact it is an attack tool.**

© 2005 Cisco Systems, Inc. All rights reserved.

HIPS v2.0--1.44

The primary vulnerabilities for end-user workstations are viruses and Trojan horse attacks. Viruses are malicious software that is attached to another program to execute a particular unwanted function on a user's workstation. An example of a virus is a program that is attached to `command.com` (the primary interpreter for Windows systems) that deletes certain files and infects any other versions of `command.com` that it can find.

A Trojan horse is different from a virus only in that the entire application is written to look like something else, when in fact it is an attack tool. An example of a Trojan horse is a software application that runs a simple game on the user's workstation. While the user is occupied with the game, the Trojan horse mails a copy of itself to every user in the user's address book. The other users receive the game and then play it, thus spreading the Trojan horse.

Virus and Trojan Horse Attack Mitigation

Cisco.com

Virus and Trojan horse applications can be contained in the following ways:

- **Use antivirus software effectively.**
- **Keep up-to-date on the latest developments in attacks.**
- **Keep up-to-date on the latest antivirus software and application versions.**
- **Use intrusion protection effectively.**

© 2005 Cisco Systems, Inc. All rights reserved.

HIPS v2.0—145

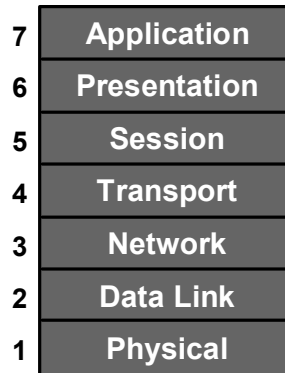
Virus and Trojan horse applications can be contained through the effective use of antivirus software and intrusion protection at the user level and potentially at the network level. Both methods can detect most viruses and many Trojan horse applications and prevent them from spreading in the network. Keeping up-to-date with the latest developments in these sorts of attacks can also lead to a more effective posture against them. As new virus or Trojan horse applications are released, enterprises need to keep up-to-date with the latest antivirus software and application versions.

Application-Layer Attacks

Cisco.com

Application-layer attacks have the following characteristics:

- Exploit well-known weaknesses, such as those in protocols, that are intrinsic to an application or system (for example, sendmail, HTTP, and FTP)
- Often use ports that are allowed through a firewall (for example, TCP port 80 used in an attack against a web server behind a firewall)
- Can never be completely eliminated, because new vulnerabilities are always being discovered



© 2005 Cisco Systems, Inc. All rights reserved.

HIPS v2.0--1-46

Application-layer attacks can be implemented using several different methods:

- One of the most common methods is exploiting well-known weaknesses in software commonly found on servers, such as sendmail, PostScript, and FTP. By exploiting these weaknesses, attackers can gain access to a computer with the permissions of the account running the application, which is usually a privileged, system-level account.
- Trojan horse attacks are implemented using programs that an attacker substitutes for common programs. These programs may provide all the functionality that the normal program provides, but also include other features that are known to the attacker, such as monitoring login attempts to capture user account and password information. These programs can capture sensitive information and distribute it back to the attacker. They can also modify application functionality, such as applying a blind carbon copy to all e-mail messages so that the attacker can read all of your organization's e-mail.

One of the oldest forms of application-layer attacks is a Trojan horse program that displays a screen, banner, or prompt that the user believes is the valid login sequence. The program then captures the information that the user enters and stores or e-mails it to the attacker. Next, the program either forwards the information to the normal login process (normally impossible on modern systems) or simply sends an expected error to the user (for example, Bad Username/Password Combination), exits, and starts the normal login sequence. The user, believing that he or she has incorrectly entered the password (a common mistake experienced by everyone), reenters the information and is allowed access.

- One of the newest forms of application-layer attacks exploits the openness of several new technologies: the HTML specification, web browser functionality, and HTTP. These attacks, which include Java applets and ActiveX controls, involve passing harmful programs across the network and loading them through a user's browser.

Application-Layer Attack Mitigation

Cisco.com

Measures you can take to reduce your risks include the following:

- **Read operating system and network log files, or have them analyzed by log analysis applications.**
- **Subscribe to mailing lists that publicize vulnerabilities.**
- **Keep your operating system and applications current with the latest patches.**
- **Use IDSs, which can scan for known attacks, monitor and log attacks, and in some cases, prevent attacks.**

© 2005 Cisco Systems, Inc. All rights reserved.

HIPS v2.0—1-47

The following are some measures you can take to reduce your risks for application-layer attacks:

- Read operating system and network log files or have them analyzed. Review all logs and take action accordingly.
- Subscribe to mailing lists that publicize vulnerabilities. Most application and operating system vulnerabilities are published on the web by various sources.
- Keep your operating system and applications current with the latest patches. Always test patches and fixes in a nonproduction environment. This practice prevents downtime and keeps errors from being generated unnecessarily.
- Use IDSs to scan for known attacks, monitor and log attacks, and in some cases, prevent attacks. The use of IDSs can be essential to identifying security threats and mitigating some of those threats. In most cases, they can be set to run automatically.

Management Protocols and Functions

The protocols that you use to manage your network can become a source of vulnerability. This topic examines common management protocols and how they can be exploited.

Configuration Management

Cisco.com

- **Configuration management protocols include SSH, SSL, and Telnet.**
- **Telnet issues include the following:**
 - **The data within a Telnet session is sent as clear text and may be intercepted by anyone with a packet sniffer located along the data path between the device and the management server.**
 - **The data may include sensitive information, such as the configuration of the device itself, passwords, and so on.**

© 2005 Cisco Systems, Inc. All rights reserved.HIPS v2.0—1-49

If the managed device does not support any of the recommended protocols, such as SSH and SSL, Telnet may be required (although this protocol is not highly recommended). The network administrator should recognize that the data within a Telnet session is sent as clear text and may be intercepted by anyone with a packet sniffer located along the data path between the managed device and the management server. The clear text may include important information, such as the configuration of the device itself, passwords, and other sensitive data.

Configuration Management Recommendations

Cisco.com

When possible, the following practices are advised:

- **Use IPSec, SSH, SSL, or any other encrypted and authenticated transport.**
- **ACLs should be configured to allow only management servers to connect to the device. All attempts from other IP addresses should be denied and logged.**
- **RFC 2827 filtering at the perimeter router should be used to mitigate the chance of an outside attacker spoofing the addresses of the management hosts.**

© 2005 Cisco Systems, Inc. All rights reserved.

HIPS v2.0—1-60

Regardless of whether SSH, SSL, or Telnet is used for remote access to the managed device, access control lists (ACLs) should be configured to allow only management servers to connect to the device. All attempts from other IP addresses should be denied and logged. RFC 2827 filtering at the ingress router should also be implemented to reduce the chance that an attacker from outside the network will spoof the addresses of the management hosts.

Management Protocols

Cisco.com

The following are management protocols that that can be compromised:

- **SNMP:** The community string information for simple authentication is sent in clear text.
- **Syslog:** Data is sent as clear text between the managed device and the management host.
- **TFTP:** Data is sent as clear text between the requesting host and the TFTP server.
- **NTP:** Many NTP servers on the Internet do not require any authentication of peers.

© 2005 Cisco Systems, Inc. All rights reserved.

HIPS v2.0—1-51

Simple Network Management Protocol (SNMP) is a network management protocol that can be used to retrieve information from a network device (commonly referred to as read-only access) or to remotely configure parameters on the device (commonly referred to as read-write access). SNMP uses passwords, called community strings, within each message as a very simple form of security. Unfortunately, most implementations of SNMP on networking devices today send the community string in clear text along with the message. Therefore, anyone with a packet sniffer located along the data path between the device and the management server can intercept SNMP messages and compromise the community string.

Syslog, which is information generated by a device that has been configured for logging, is sent as clear text between the managed device and the management host. Syslog has no packet-level integrity checking to ensure that the packet contents have not been altered in transit. An attacker may alter syslog data in order to confuse a network administrator during an attack.

Trivial File Transfer Protocol (TFTP) is used for transferring configuration or system files across the network. TFTP uses UDP for the data stream between the requesting host and the TFTP server.

As with other management protocols that send data in clear text, the network administrator should recognize that the data within a TFTP session might be intercepted by anyone with a packet sniffer located along the data path between the device and the management server. Where possible, TFTP traffic should be encrypted within an IPSec tunnel in order to reduce the chance of its being intercepted.

Network Time Protocol (NTP) is used to synchronize the clocks of various devices across a network. Synchronization of the clocks within a network is critical for digital certificates and for correct interpretation of events within syslog data.

A secure method of providing clocking for the network is for network administrators to implement their own master clocks for private networks synchronized to Coordinated Universal Time (UTC) via satellite or radio. For network administrators who do not wish to implement

their own master clocks because of cost or other reasons, clock sources are available for synchronization via the Internet.

An attacker could attempt a DoS attack on a network by sending bogus NTP data across the Internet to change the clocks on network devices in such a manner that digital certificates are considered invalid. Further, an attacker could attempt to confuse a network administrator during an attack by disrupting the clocks on network devices. This scenario would make it difficult for the network administrator to determine the order of syslog events on multiple devices.

Management Protocol Recommendations

Cisco.com

- **SNMP recommendations:**
 - Configure SNMP with only read-only community strings.
 - Set up access control on the device you wish to manage.
 - Use SNMP Version 3 or above.
- **Logging recommendations:**
 - Encrypt syslog traffic within an IPsec tunnel.
 - Implement RFC 2827 filtering.
 - Set up access control on the firewall.
- **TFTP recommendations:**
 - Encrypt TFTP traffic within an IPsec tunnel.
- **NTP recommendations:**
 - Implement your own master clock.
 - Use NTP Version 3 or above.
 - Set up access control that specifies which network devices are allowed to synchronize with other network devices.

© 2005 Cisco Systems, Inc. All rights reserved.

HIPS v2.0—1-52

The following are SNMP recommendations:

- Configure SNMP with only read-only community strings.
- Set up access control on the device you wish to manage via SNMP to allow access by only the appropriate management hosts.
- Use SNMP Version 3 or above.

When possible, the following practices are advised:

- Encrypt syslog traffic within an IPsec tunnel.
- Implement RFC 2827 filtering at the perimeter router when allowing syslog access from devices on the outside of a firewall.
- Implement ACLs on the firewall in order to allow syslog data from only the managed devices themselves to reach the management hosts.
- Encrypt TFTP traffic when possible within an IPsec tunnel in order to reduce the chance of its being intercepted.

The following are NTP recommendations:

- Implement your own master clock for private network synchronization.
- Use NTP Version 3 or above because these versions support a cryptographic authentication mechanism between peers.
- Use ACLs that specify which network devices are allowed to synchronize with other network devices.

Summary

This topic summarizes what you learned in this lesson.

Summary

Cisco.com

- The need for network security has increased as networks have become more complex and interconnected.
- The following are the components of a complete security policy:
 - Statement of authority and scope
 - Acceptable use policy
 - Identification and authentication policy
 - Internet use policy
 - Campus access policy
 - Remote access policy
 - Incident handling procedure
- The Security Wheel details the view that security is an ongoing process.
- The Security Wheel comprises four phases: secure, monitor, test, and improve.

© 2005 Cisco Systems, Inc. All rights reserved. HPS v2.0—1-54

Summary (Cont.)

Cisco.com

- The following are the four types of security threats:
 - Structured
 - Unstructured
 - Internal
 - External
- The following are the four primary attack categories:
 - Reconnaissance attacks
 - Access attacks
 - Denial of service attacks
 - Worms, viruses, and Trojan horses
- Configuration management and management protocols are an important part of securing a network.

© 2005 Cisco Systems, Inc. All rights reserved. HPS v2.0—1-55

Lesson 2

Cisco Security Agent Overview

Overview

This lesson describes security fundamentals. It includes the following topics:

- Objectives
- Defense in Depth
- Cisco Security Agent Architecture
- Anatomy of an Attack and Response
- Key Features of Cisco Security Agent
- Summary

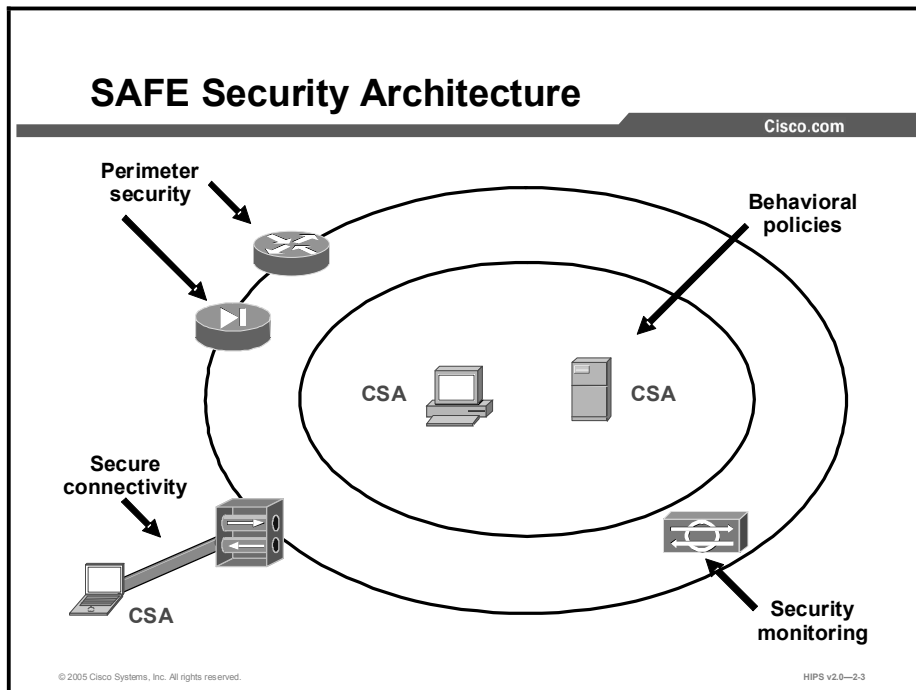
Objectives

Upon completion of this lesson, you will meet the following objectives:

- Explain the defense-in-depth concept of protecting a network
- Discuss Cisco Security Agent architecture
- Understand the life cycle of an attack
- Explain how Cisco Security Agent protects against attacks

Defense in Depth

This topic describes implementation of the SAFE architecture's defense-in-depth objective. In addition, the Cisco Security Agent (CSA) is introduced.



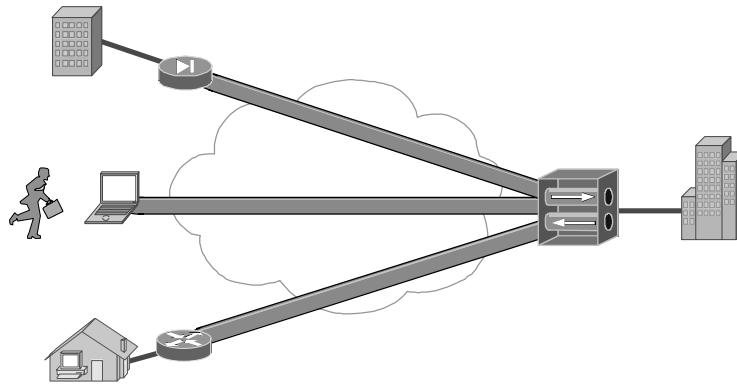
The Cisco SAFE architecture mandates that attacks that succeed in penetrating the first line of defense, or that originate from inside the network, must be accurately detected and quickly contained to minimize their effect on the rest of the network. In today's environment, new threats occur more frequently and are often disseminated faster than patches or new signatures can be made available. Effective network protection cannot depend solely upon a hardened perimeter or alerts and reactive measures. Cisco provides a variety of mechanisms to achieve this layered approach to network security.

Each component of a network defense in depth makes its own security contribution:

- **Cisco Virtual Private Network (Cisco VPN):** Provides secure connectivity across public networks.
- **Router:** Filters out spoofed IP addresses. Exploits using falsified Layer 3 information are blocked.
- **Cisco Security Appliance:** Controls network traffic by stateful sessions. Security is enforced with knowledge of connection information at Layers 3 and 4.
- **Cisco Intrusion Prevention System (Cisco IPS):** Monitors intranetwork traffic for suspicious signatures. Packets are screened up to Layer 7 for matches to the signature database.
- **Cisco Security Agent (CSA):** Filters system resource calls to the kernel by applications. CSA uses behavioral policies to enforce security at Layers 3 through 7 in network traffic and applies the behavioral requirements to internal system resources calls.

Virtual Private Networks

Cisco.com



Cisco VPNs: Provide secure, reliable, and authenticated connectivity over a shared public network.

© 2005 Cisco Systems, Inc. All rights reserved.


HIPS v2.0—2.4

Cisco VPNs deliver a first layer of defense with secure, reliable, and authenticated connectivity over a shared public network for mobile users and for home and branch offices. A VPN insures the security of connections through three services:

- Confidentiality is provided by encryption algorithms, such as Data Encryption Standard (DES) and Triple-Data Encryption Standard (3DES).
- Authentication of peers can be done by Internet Key Exchange (IKE), and data can be authenticated by Authentication Header (AH) or Encapsulating Security Payload (ESP).
- Integrity of the data is maintained by hash algorithms, such as Message Digest 5 (MD5) and Secure Hash Algorithm (SHA).

At the Perimeter

Cisco.com



- **Cisco security appliances: Perform stateful packet filtering and block SYN DoS attacks.**
- **Cisco routers: Prevent IP spoofing.**

© 2005 Cisco Systems, Inc. All rights reserved.

HIPS v2.0--2.5

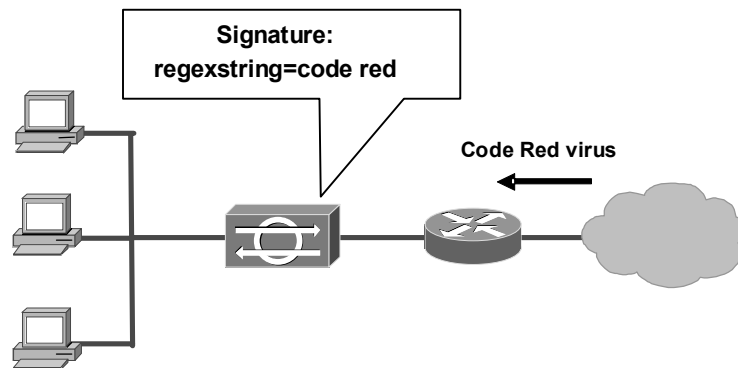
The security appliance efficiently guards the perimeter of the network using the Adaptive Security Algorithm (ASA). By combining a session flow table with stateful knowledge of how certain applications and protocols function, the security appliance creates temporary openings in the firewall only as needed for connections that were initiated from the higher-security interfaces. The security appliance also monitors the number of embryonic TCP sessions started to the network hosts in order to prevent SYN denial of service (DoS) attacks.

Cisco routers at the network's edge increase network security by blocking traffic from spoofed IP addresses. A Cisco router with Cisco IOS Firewall software can also function as a stateful firewall, although not with the very high efficiency levels of the security appliance.

Firewalls are key building blocks in a network defense in depth. They do not, however, provide protection from attacks from within the network or from exploits that enter the network disguised as normal traffic. Two examples of this are worms hidden in e-mail messages or attacks inside innocent-seeming traffic to the corporate servers that are exposed to public networks.

Monitoring the Intranetwork

Cisco.com



Cisco IPS matches traffic to signatures of known exploits.

© 2005 Cisco Systems, Inc. All rights reserved.

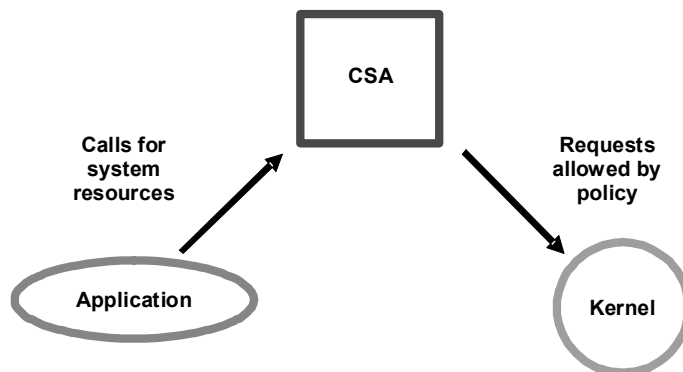
HIPS v2.0—2-6

Cisco Intrusion Prevention Systems (IPSs) provide a second line of defense by matching traffic in the intranetwork against signatures of known exploits. When suspicious traffic is detected, the IPS can report it to a log, send an alert, or take reactive steps to block any malicious activity. An IPS monitors traffic inside the network, so it can provide alerts for attacks launched internally. Packets are examined at higher levels by an IPS, detecting harmful payloads even inside traffic whose session state matches that of innocent packets.

The effectiveness of an IPS depends on the update status of its signature database. A brand-new exploit can slip past an IPS that does not yet have the signature of that exploit. Another issue is that signatures can be too broad, which can generate many false positive alerts.

Host-Based Intrusion Protection System

Cisco.com



CSA compares application calls for system resources to the security policy.

© 2005 Cisco Systems, Inc. All rights reserved.

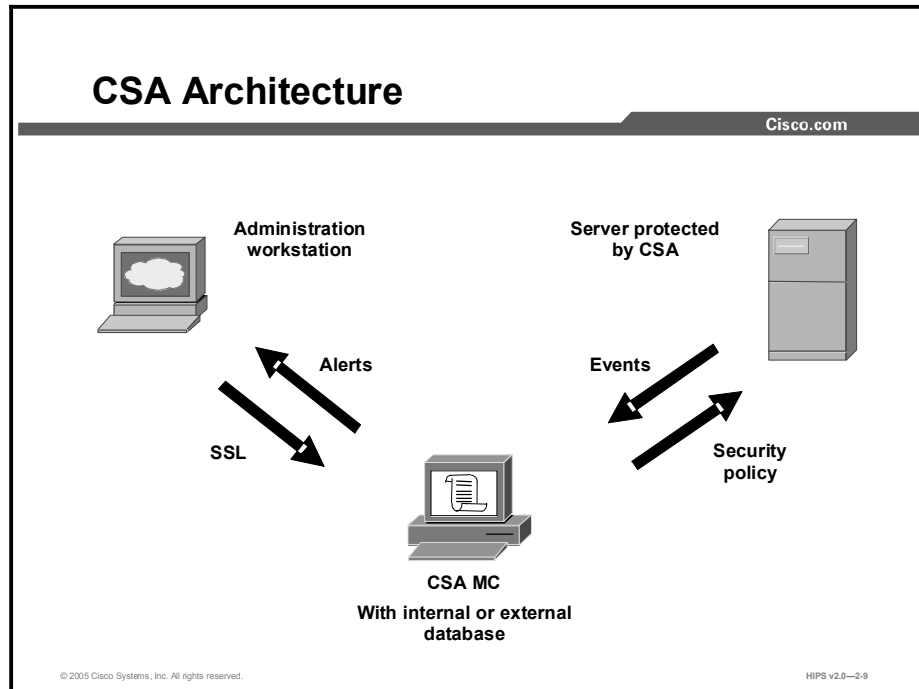
HIPS v2.0--2.7

CSA is a host intrusion prevention system (HIPS) that provides a third layer of depth to the network defense by applying security policy to system behavior at the host level. With this different approach to security, CSA stops attacks that are missed at other levels of network security. CSA has the following advantages:

- CSA proactively blocks intrusive attacks by comparing all requests for system resources to the behaviors allowed by the security policy.
- CSA is not dependent upon signatures or updates to recognize attacks; in other words, it provides Day Zero protection from previously unknown attacks.
- CSA creates significantly fewer false positive alerts than any IDS; therefore, less administrative time is needed.

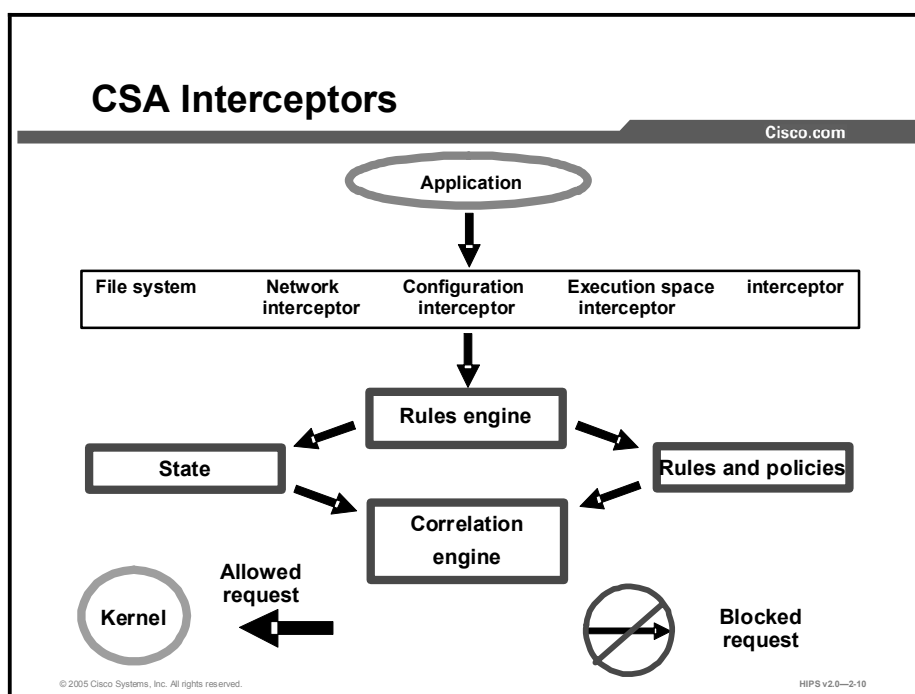
Cisco Security Agent Architecture

This topic describes the CSA architecture.



The CSA architecture model consists of three elements:

- **Cisco Security Agent Management Center (CSA MC):** The administrator divides network hosts into groups by function and security requirements and then configures security policies for those groups. The CSA MC maintains a log of security violations and sends alerts through e-mail or a pager.
- **CSA:** Software installed in the host systems continually monitors local system activity and analyzes the operations of that system. CSA takes proactive action to block attempted malicious activity. CSA also polls the CSA MC at configurable intervals for policy updates.
- **Administration workstation:** Any workstation can be connected securely to the CSA MC using a web interface with Secure Socket Layer (SSL) enabled.



When an application needs access to system resources, it makes an operating system call to the kernel. CSA intercepts these operating system calls and compares them to the cached security policy. If the request does not violate policy, it is passed to the kernel for execution.

If the request does violate policy, CSA takes the following actions:

- The request is blocked; it is not passed to the kernel.
- An appropriate error message is passed back to the application.
- An alert is generated and sent to the CSA MC.

CSA correlates this particular operating system call with others made by that application or process and correlates these events to detect malicious activity.

CSA provides protection through deployment of four interceptors:

- **File system interceptor:** All file read or write requests are intercepted and allowed or denied based on the security policy.
- **Network interceptor:** Network driver interface specification (NDIS) changes are controlled and network connections are cleared through the security policy by port/IP address pairs. The number of network connections allowed within a specified time can also be limited to prevent DoS attacks.
- **Configuration interceptor:** Read/write requests to the registry on Windows or to rc files on UNIX are intercepted. Because modification of operating system configuration is highly unusual, it is tightly controlled by CSA.
- **Execution space interceptor:** This interceptor deals with maintaining the integrity of each application's dynamic run-time environment. It detects and blocks requests to write to memory that are not owned by the requesting application. Attempts by one application to inject code, such as a shared library or dynamic link library (DLL), into another application are also detected and blocked. Buffer overflow attacks are detected by this interceptor as well. The result is that not only is the integrity of dynamic resources, such as the file

system and configuration, preserved, but the integrity of highly dynamic resources such as memory and network I/O is also preserved.

CSA Interceptors (Cont.)

Cisco.com

Security Application	Network Interceptor	File System Interceptor	Configuration Interceptor	Execution Space Interceptor
Distributed firewall	X	—	—	—
Host intrusion detection	X	—	—	X
Application sandbox	—	X	X	X
Network worm prevention	X	—	—	X
File integrity monitor	—	X	X	—

© 2005 Cisco Systems, Inc. All rights reserved.

HIPS v2.0—2-11

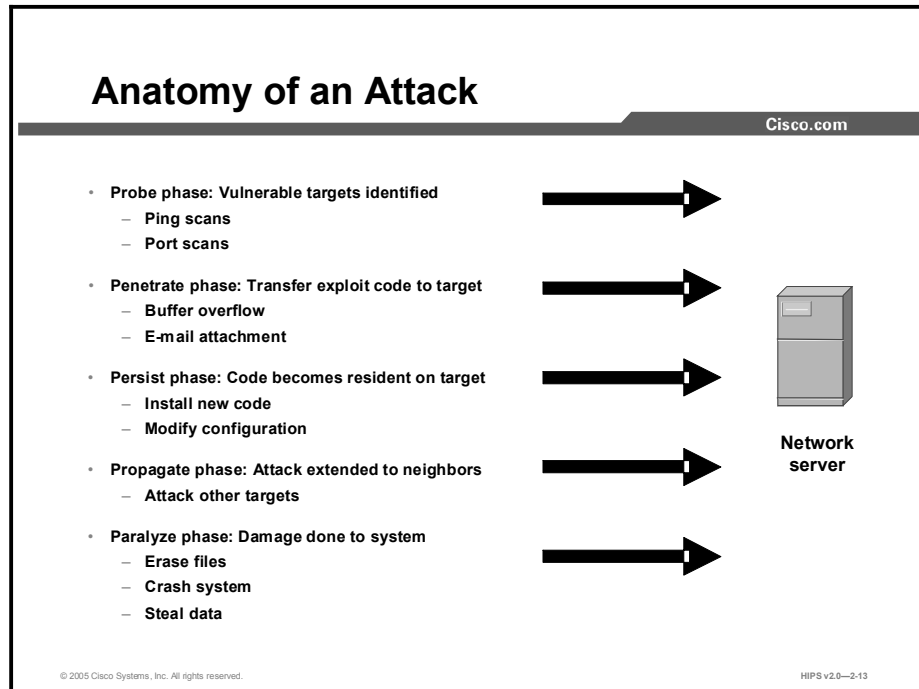
By intercepting communication between applications and the underlying system, CSA combines the functionality of the following traditional security approaches:

- **Distributed firewall:** The network interceptor performs the duties of a host firewall.
- **Host intrusion detection system (HIDS):** The network interceptor teams with the execution space interceptor to provide the alerting capability of an HIDS with the proactive enforcement of a security policy.
- **Application sandbox:** An application sandbox is an execution space in which suspect programs can be run with less than normal access to system resources. A combination of the file system, configuration, and the execution space interceptors provides this security service.
- **Network worm prevention:** The network and execution space interceptors provide Day Zero worm prevention without a need for updates.
- **File integrity monitor:** The file system and configuration interceptors act as a file integrity monitor.

The default policies preconfigured on CSA implement all of these security features. Customers can easily create or change policies, but the default policies provide all of these protections at once.

Anatomy of an Attack and Response

This topic describes the progression of an attack and the CSA response.



Malicious attacks come in thousands of varieties, and new attacks are being devised constantly to exploit newly discovered vulnerabilities, but their basic goals have remained nearly constant over time.

An analysis of the logical progression of an attack helps illustrate how almost all exploits are intended to gain control of core mechanisms in the target system:

- **Probe phase:** Vulnerable targets are identified in this phase. The goal of this phase is to find computers that can be subverted. ICMP ping scans are used to map networks, and application port scans identify operating systems and vulnerable software. Passwords can be obtained through social engineering, a dictionary attack, a brute-force attack, or network sniffing.
- **Penetrate phase:** Exploit code is transferred to the vulnerable target in this phase. The goal of this phase is to get the target to execute the exploit code via an attack vector like a buffer overflow, ActiveX or Common Gateway Interface (CGI) vulnerabilities, or an e-mail virus.
- **Persist phase:** Once an exploit has been successfully launched into memory, the exploit code tries to persist on the target system. The goal of this phase is to ensure that the attacker's code will be running and available to the attacker even if the system reboots. The exploit code achieves this goal by modifying system files, making registry changes, installing new code, and so on.
- **Propagate phase:** Upon establishing a beachhead in the organization, the attacker attempts to extend the attack to other targets. This phase looks for vulnerable neighboring machines. Propagation vectors would include e-mailing copies of the attack to other systems, uploading files to other systems using file shares or FTP services, active web connections, and file transfers via Internet Relay Chat (IRC).

- **Paralyze phase:** This is the phase in which actual damage is done to the system. Files can be erased, systems can be crashed, information can be stolen, and distributed DoS attacks can be launched.

There are significant differences between the attack mechanisms used at the probe and penetrate phases and those used at the persist phase.

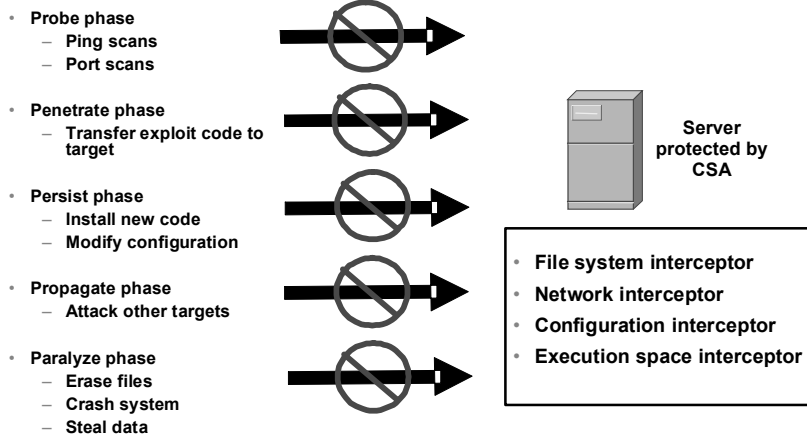
The first two stages mutate constantly, with new vulnerabilities being discovered and custom exploits crafted almost daily. Combating attacks at the probe and penetrate phases requires constant updating of malicious IDS signatures and firewall defenses as these attacks evolve. Attacks at these early phases also lend themselves to evasion techniques like Unicode encoding of web strings or overlapping packet fragments. The mutability of attacks at the penetrate stage requires a significant amount of interpretation; many false alarms are generated, requiring time-consuming review by a security administrator.

In contrast, attack mechanisms at the persist phase and later phases have been remarkably stable over time. Attackers can commit a limited number of malicious activities, and they all involve making a system call to the kernel to access system resources. The malicious code can attempt to modify the operating system, modify files, create or alter network connections, or violate the memory space of active processes. The list of potential attacks on system resources has remained stable. These attacks use different vectors to access the target systems, but the actions performed in the persist phase are very similar.

Because consistently identifying attacks at the early phases of a newly developed exploit can be nearly impossible, CSA focuses on providing proactive security by controlling access to system resources. This approach avoids the race to update defenses in order to keep up with the latest exploit and protects hosts even on Day Zero of a new attack. For example, the Nimda and Slammer worms did millions of dollars in damage to enterprises on the first day of their appearance, before updates were even available, but CSA stopped these attacks without any updates by identifying their behavior as malicious.

CSA Attack Response

Cisco.com



© 2005 Cisco Systems, Inc. All rights reserved.

HIPS v2.0—2-14

When an application attempts to write to a file, make registry changes, or access system resources in any way, it must make an operating system call to the kernel. CSA provides complete enforcement of your security policy by policing these requests from applications to the kernel.

CSA intercepts operating system calls and compares them with a cached policy that is centrally defined on the CSA MC. If the request does not violate policy, it is passed to the kernel for execution, but if the request does violate policy, it is blocked. An alert is then generated by the host's CSA and sent to the CSA MC.

By controlling behavior at the operating system call level, CSA blocks attacks at the persist, propagate, and paralyze phases without the constant updates that are required at the probe and penetrate phases.

Key Features of Cisco Security Agent

This topic describes the key features of CSA.

CSA Features

Cisco.com

- **Real-time protection decisions**
- **Defense-in-depth approach**
 - **Intercepts communication between applications and the kernel**
 - **Protects system from attacks at all phases**
- **Ease of deployment**
 - **Deploys with default policies in 30 minutes**
 - **Allows easy configuration of custom policies**
- **Broad platform support**
 - **Windows or UNIX (Solaris and Linux)**
 - **Servers and desktops**

© 2005 Cisco Systems, Inc. All rights reserved. HIPS v2.0—2.16

The following are the key features of CSA:

- **Real-time protection decisions:** CSA provides protection in real time rather than spotting attacks after they have happened.
- **Defense-in-depth approach:** More than a network perimeter defense or an attempt to detect attacks inside the network, CSA controls access to host system resources for complete protection.
 - Intercepts communication between applications and the kernel
 - Protects system from attacks at all phases
 - Network
 - File system
 - Configuration
 - Execution space
- **Ease of deployment**
 - Deploys with default policies in 30 minutes
 - Allows easy configuration of custom policies
- **Broad platform support**
 - Windows or UNIX (Solaris and Linux)
 - Servers and desktops

CSA Features (Cont.)

Cisco.com

- **Real-time correlation at Agent and enterprise-wide**
- **Ease of administration**
 - No need for constant review of logs
 - No updates: Day Zero ready
 - Manage from any web browser
- **Centralized event management**
 - E-mail, pager, SNMP alerts controlled at CSA MC
 - Logging and report-generating capability
- **Enforce and Detect rule organization**
- **Internationalization and localization for Windows agents**
- **Integrated with Cisco Trust Agent**

© 2005 Cisco Systems, Inc. All rights reserved.

HIPS v2.0—2-17

- **Real-time correlation at Agent and enterprise-wide:** Reduces false positives and allows adaptability to new threats enterprise-wide.
 - A network scan over multiple systems within a configured time period logs network events.
 - Worm events on multiple systems cause all systems to quarantine the contaminated files.
 - NT Event Logs and virus scanner logs can be correlated across the enterprise.
- **Ease of administration**
 - Less need for constant review of logs: proactive defense approach minimizes requirements for administrator involvement.
 - No updates: Day Zero ready.
 - Secure management from any web browser.
- **Centralized event management**
 - E-mail, pager, Simple Network Management Protocol (SNMP) alerts controlled at the CSA MC.
 - Logging and report-generating capability.
- **Enforce and Detect rule organization:** Combined rule lists are now organized as a combination of Enforce and Detect rules. Enforce rules are primarily access control rules that allow, deny, or terminate an action. Detect rules are monitoring, logging, and tagging rules. In rule display lists, enforce rules are shown at the top of the list and detect rules are shown at the bottom. These rule types work together to monitor actions, build application classes, and protect systems.
- **Internationalization and localization for Windows Agents:** The Cisco Security Agent now accepts and displays query text characters appropriately for the selected language type. It

also displays events in non-ASCII characters so that internationalization of events is possible.

- Integrated with Cisco Trust Agent: The Cisco Security Agent is a supported configuration for the Cisco Trust Agent feature.

Summary

This topic summarizes the information that you learned in this lesson.

Summary

Cisco.com

- **Defense in depth can be achieved by using a layered deployment of security mechanisms.**
- **CSA consists of the CSA MC, the Agent installed on each host, and any web-based workstation for administration.**
- **CSA intercepts operating system calls and clears them against a behavioral security policy.**
- **Because CSA protects at the system resource level, the ever-changing form of attack is irrelevant and evasion techniques are ineffective.**

© 2005 Cisco Systems, Inc. All rights reserved.

HIPS v2.0--2-18

Lesson 3

Cisco Security Agent Quick Start Installation

Overview

This lesson includes the following topics:

- Objectives
- CSA MC System Requirements
- CSA System Requirements
- Installing the CSA MC
- Configuring the CSA MC
- Installing the CSA
- Summary
- Lab Exercise

Objectives

Upon completion of this lesson, you will meet the following objectives:

- Identify the CSA MC and CSA system requirements
- Identify the administration workstation requirements
- Install the CSA MC
- Configure the CSA MC
- Install the CSA

CSA MC System Requirements

This topic identifies the requirements for installing the Cisco Security Agent Management Center (CSA MC).

CSA MC Installation Requirements

Cisco.com

- **Processor: 1 GHz or faster**
- **Memory: 1 GB minimum**
- **Virtual memory: 2 GB minimum**
- **Disk space: 9 GB minimum available**
- **File system: NTFS**
- **Operating system**
 - **Windows 2000 Server (Service Pack 4)**
 - **Windows Advanced Server (Service Pack 4)**
- **Modem: Hayes compatible, if pager alerts are required**

© 2005 Cisco Systems, Inc. All rights reserved.HIPS v2.0--3.3

The following are the recommended platform specifications for the CSA MC:

- **Processor:** 1 GHz or faster
- **Memory:** 1 GB minimum
- **Virtual memory:** 2 GB minimum
- **Disk space:** 9 GB minimum available

Note The actual amount of disk space required depends on the number of CiscoWorks Common Services client applications that you are installing and the number of devices that you are managing with the client applications.

- **File system:** New Technology File System (NTFS)
- **Operating system:** You have a choice of operating system:
 - Windows 2000 Server (Service Pack 4)
 - Windows Advanced Server (Service Pack 4)

Note Terminal services must be turned off on Windows Advanced Server.

- **Modem:** Hayes compatible, if pager alerts are required

Note The CSA MC is a component of the CiscoWorks Virtual Private Network/Security Management Solution.

CSA MC Database Installation

Cisco.com

- **CSA MC has the option to install a local database using the Microsoft SQL Server Desktop Engine (MSDE), or you may use Microsoft SQL Server 2000.**
- **The other option is to use an external database.**

© 2005 Cisco Systems, Inc. All rights reserved.

HIPS v2.0--3-4

CSA MC has two options for database installation. The first is to install a local database using the Microsoft SQL Server Desktop Engine (MSDE), or you may use Microsoft SQL Server 2000. The second option is to use an external database installed on another server.

CSA MC Installation Recommendations

Cisco.com

- **Place the system in a physically secure location with limited access.**
- **Install CSA MC required software only.**
- **Use a static IP address.**
- **Use HTTPs for communication.**

© 2005 Cisco Systems, Inc. All rights reserved.

HIPS v2.0—3-5

Installation recommendations for the CSA MC are as follows:

- The system on which you are installing the CSA MC software should be placed in a physically secure location with restricted access.
- Only software required by the CSA MC itself should be installed on the CSA MC system.
- The CSA MC system must have a static IP address or a fixed Dynamic Host Configuration Protocol (DHCP) address.
- CSA systems must be able to communicate with the CSA MC over Hypertext Transfer Protocol secure (HTTPS).

CSA System Requirements

This topic identifies the requirements for installing Cisco Security Agent (CSA).

CSA Installation Requirements (Windows)

Cisco.com

- **Processor: 200 MHz or faster**
- **Memory: 128 MB minimum**
- **Disk space: 15 MB minimum available**
- **File system: NTFS**
- **Operating system**
 - **Windows 2000 Professional, Server, or Advanced Server (Service Pack 0–4)**
 - **Windows 2003**
 - **Windows XP Professional (Service Pack 0, 1, or 2)**
 - **Windows NT Workstation, Server, or Enterprise Server (Service Pack 5 or higher)**
- **Network: Ethernet or dial-up**

© 2005 Cisco Systems, Inc. All rights reserved. HIPS v2.0–3-7

The following are the recommended Windows platform specifications for CSA:

- **Processor:** 200 MHz or faster

Note Uniprocessor and dual-processor systems are supported.

- **Memory:** 128 MB minimum

Note CSA uses approximately 20 MB of memory.

- **Disk space:** 15 MB minimum available
- **File system:** NTFS
- **Operating system:** The choices are as follows:
 - Windows 2000 Professional, Server, or Advanced Server (Service Pack 0–4)
 - Windows 2003
 - Windows XP Professional (Service Pack 0, 1, or 2)
 - Windows NT Workstation, Server, or Enterprise Server (Service Pack 5 or higher)

Note Terminal services are not supported on Windows NT.

- **Network:** Ethernet or dial-up

Note A maximum of 64 IP addresses are supported on a system.

For Agents and browsers to successfully communicate with the CSA MC, the CSA MC system name must be resolvable through Domain Name System (DNS) or Windows Internet Name Service (WINS).

CSA Installation Requirements (Solaris)

Cisco.com

- **Processor:** UltraSPARC, 400 MHz or faster
- **Memory:** 256 MB minimum
- **Disk space:** 15 MB minimum available
- **Operating system:** Solaris 8, 64 bit 12/02 Edition or higher
- **Network:** Ethernet

© 2005 Cisco Systems, Inc. All rights reserved.

HIPS v2.0--3-8

The following are the recommended Solaris platform requirements for CSA:

- **Processor:** UltraSPARC 400 MHz or faster
- **Memory:** 256 MB minimum
- **Disk space:** 15 MB minimum available
- **Operating system:** Solaris 8, 64 bit 12/02 Edition or higher

Note Before you install CSA, the system must have the SUNWlibCxx library installed. It can be installed from the first Solaris 8 CD in the /Solaris_8/Product directory using the **pkgadd -d SUNWlibCxx** command.

- **Network:** Ethernet

Note A maximum of 64 IP addresses are supported on a system.

Caution If a new type of Ethernet interface is added to a UNIX system running CSA, the system must be rebooted twice for the Agent to detect it and apply rules to it accordingly.

CSA Installation Requirements (Linux)

Cisco.com

- **Processor:** Intel Pentium 500 MHz or faster
- **Memory:** 256 MB minimum
- **Disk space:** 15 MB minimum available
- **Operating system:** RedHat Enterprise Linux 3.0 WS, ES, and AS
- **Network:** Ethernet

© 2005 Cisco Systems, Inc. All rights reserved.

HIPS v2.0—3-9

The following are the recommended Linux platform requirements for CSA:

- **Processor:** Intel Pentium 500 MHz or faster
- **Memory:** 256 MB minimum
- **Disk space:** 15 MB minimum available
- **Operating system:** RedHat Enterprise Linux 3.0 WS, ES, and AS
- **Network:** Ethernet

Note A maximum of 64 IP addresses are supported on a system.

Administrator Workstation Requirements

Cisco.com

The following web browsers are supported:

- **Netscape**
 - **Version 7.1 or higher.**
 - **Cookies must be enabled.**
 - **JavaScript must be enabled.**
- **Internet Explorer**
 - **Version 6.0 or higher.**
 - **Cookies must be enabled.**
 - **JavaScript must be enabled.**

© 2005 Cisco Systems, Inc. All rights reserved.

HIPS v2.0—3-10

The following are the recommended browsers for the administrator workstation:

- **Netscape**
 - Version 7.1 or higher.
 - Cookies must be enabled.
 - JavaScript must be enabled.
- **Internet Explorer**
 - Version 6.0 or higher.
 - Cookies must be enabled.
 - JavaScript must be enabled.

Note To access the CSA MC GUI from CiscoWorks, you must have Secure Socket Layer (SSL) enabled in CiscoWorks.

Installing the CSA MC

This topic describes the installation of the CSA MC.



When choosing to use a local database, the installation checks whether Microsoft SQL Server is installed. CSA MC uses SQL Server for its configuration database. If this software is not detected, you are prompted to install it.

Caution The setup program installs the Microsoft SQL Server Desktop Engine (MSDE). If the CSA MC installation detects any other database type attached to an existing installation of MSDE, the installation will abort. This database configuration is not supported by Cisco.

Step 1 Click **Yes** to install MSDE.

Step 2 Proceed through the Microsoft SQL Server installation. The first installation screen prompts you to accept the default SQL Server install directory path. The default is selected by searching the system disk for a location that provides the most space for the database. You can select a different path if you choose.


Note When the Microsoft SQL Server installation finishes, you must begin the CSA MC installation again. The system will require a reboot before you can restart the CSA MC installation.

Step 3 Begin the CSA MC installation again. This time, after you click **Install**, the installation detects the Microsoft SQL Server software and proceeds by displaying the introduction screen.

Step 4 Click **Next** to continue.

CSA MC Installation (Cont.)

Cisco.com



- **Click Yes to install the license file.**
- **Browse to the license file location.**

© 2005 Cisco Systems, Inc. All rights reserved.

HIPS v2.0—3-14

Step 5 You are reminded that you must obtain a license key. If you already have a license key file on the system to which you are installing the CSA MC, you can copy it to the installation directory at this time by clicking the Yes button and browsing to it on the system. You can also click No and copy it any time after the installation.

Note If you copy a valid license key to the CSA MC during the installation, after the system reboots, all downloaded and installed Agent kits immediately operate with full functionality. You do not have to log in and generate rules to have this occur. If you copy your license after the installation completes, downloaded and installed Agents operate in test mode until you generate rules.

Once all the files are copied, the installation performs some preliminary system setup tasks.

Note When the CSA MC installation is complete, an Agent installation automatically begins. It is recommended that an Agent protect the CSA MC system, and this process is done automatically for you. (You may uninstall the Agent separately if you choose, but this configuration is not recommended.)

Step 6 The installer program announces a reboot in 2 minutes and then reboots the system at the end of that time to complete the installation.

Note Installation of the CSA MC and CSA produces log files: CSCOp\CSA MC\log\Management Center for Cisco Security Agents\InstallInfo.txt and Cisco\CSAgent\log\Cisco Security Agent\InstallInfo.txt. These text files can help with troubleshooting if there are problems during installation.

Configuring the CSA MC

This topic discusses the configuration of the CSA MC.

CSA MC Configuration

Cisco.com

Basic configuration steps for the CSA MC:

- **Log in to CiscoWorks.**
- **Verify SSL on CiscoWorks.**
- **Select a default group.**
- **Obtain the Agent kit URL for the group.**

© 2005 Cisco Systems, Inc. All rights reserved.HIPS v2.0—3-16

The following are the basic configuration steps for the CSA MC:

- **Log in to CiscoWorks:** Configuring the CSA MC requires a CiscoWorks administrator login.
- **Verify SSL on CiscoWorks:** CiscoWorks is required to have SSL enabled for communication with the CSA MC. SSL is enabled automatically during the installation of the CSA MC. You should not ever disable SSL in CiscoWorks after installing the CSA MC.
- **Select a default group:** 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 with similar security requirements.
- **Obtain the Agent kit URL for the group:** The user or administrator of the host can use the Agent kit URL to register with the CSA MC and install the CSA software.

Note CSA default Agent kits, groups, policies, and configuration variables are designed to provide a high level of security coverage for desktops and servers. These default Agent kits, groups, policies, and configuration variables cannot anticipate all possible local security policy requirements specified by your organization's management, nor can they anticipate all local combinations of application usage patterns. It is recommended that you deploy CSA using the default configurations and then monitor and tune it for your environment.

CSA MC Configuration

Cisco.com

Basic configuration steps for the CSA MC:

- **Log in to CiscoWorks.**
- **Verify SSL on CiscoWorks.**
- **Select a default group.**
- **Obtain the Agent kit URL for the group.**

© 2005 Cisco Systems, Inc. All rights reserved.

HIPS v2.0-3-16

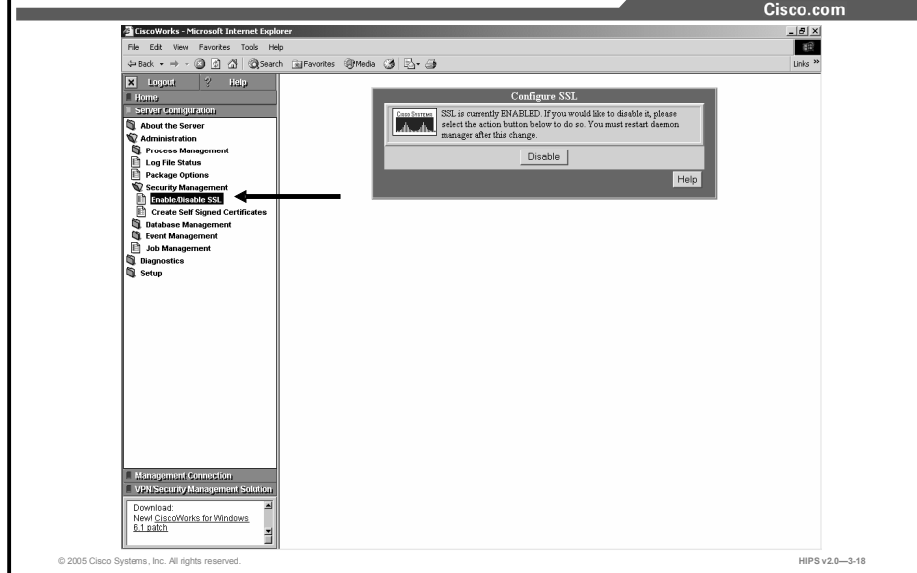
When the installation is complete and the system has rebooted, the CSA MC interface is available on the local system hosting the CiscoWorks software by choosing **Start > Programs > CiscoWorks > CiscoWorks** to open the CiscoWorks GUI. Next, log in to CiscoWorks.

Administrators can have different levels of CSA MC database access privileges. The initial administrator created by the CiscoWorks installation automatically has configuration privileges.

The following are CSA MC administrator roles:

- **Configure:** This role provides full read and write access to the CSA MC database.
- **Deploy:** This role 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:** This role provides administrators with read access to the entire CSA MC database. Administrators can also create reports, alerts, and event sets.

Initiating Secure Communications



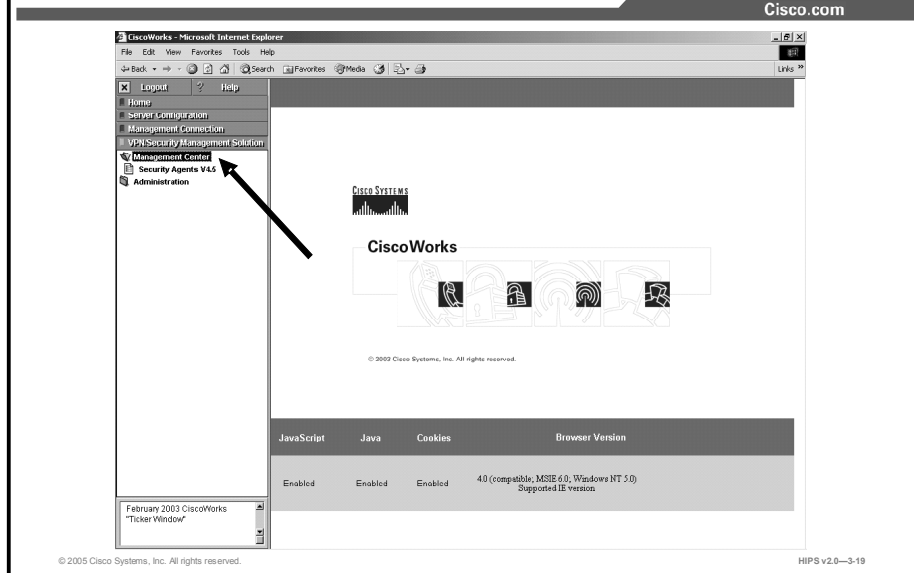
The CSA MC uses SSL to secure all communications to the CSA MC user interface, locally and remotely. All configuration data travels over secure channels irrespective of the location of the CSA MC host system.

During installation, the CSA MC generates private and public keys to be used for secure communications between any system accessing the CSA MC user interface and the CSA MC itself. To access the CSA MC user interface from CiscoWorks, you must have SSL enabled in CiscoWorks for the connection to be allowed.

Caution SSL is enabled during the installation of the CSA MC. Do not disable SSL under CiscoWorks, or the CiscoWorks management console can become inaccessible.

Note When your browser connects to the server, it receives the server's certificate. You are then prompted to accept this certificate. It is recommended that you import it into your local certificate database so that you are not prompted to accept the certificate each time that you log in.

Accessing the CSA MC Interface



To access the CSA MC interface on the system running CiscoWorks, choose **VPN/Security Management Solution > Management Center > Security Agents** (as shown in the figure).

To access the CSA MC from a remote system, launch a browser on the remote host and enter the following URL: **https://(ciscoworks system hostname):1741**. Next, log in to CiscoWorks and click **VPN/Security Management Solution > Management Center > Security Agents**.

Selecting a Default Group

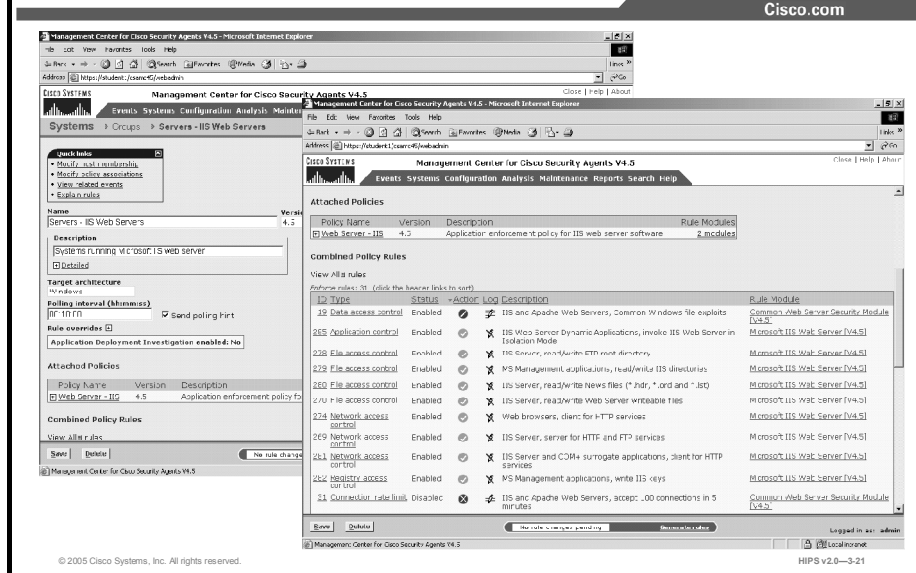
The screenshot displays the Management Center for Cisco Security Agents V4.5. The interface shows a list of system groups under the 'Data SYSTEMS' tab. The groups are organized into categories like Servers, Desktops, and Systems. The 'Servers - IIS Web Servers' group is highlighted, and an arrow points to it. The interface also shows navigation tabs for Events, Systems, Configuration, Analysis, Maintenance, Reports, Search, and Help. The status bar at the bottom indicates 'No rules changes pending' and 'Logged in as: admin'.

Group Name	Version	Description	Platform	Host Count
Servers - All Types	4.5	Default group for systems that install the Server agent kit	Solaris	
Servers - Apache Web Servers	4.5	Apache web server systems	Solaris	
Servers - Externally Deployed	4.5	Default group for servers deployed on public networks	Solaris	
Servers - iPlanet Web Servers	4.5	iPlanet web server systems	Solaris	
Systems - IDS Mode	4.5	Servers running in intrusion detection mode with no preventive capabilities	Solaris	
Systems - Mission Critical	4.5	Systems that need to be monitored at a higher priority	Solaris	
Systems - Restricted Networking	4.5	Systems which are under network lockdown	Solaris	
Systems - Test Mode	4.5	Systems operating in test mode	Solaris	
All Windows		Auto-enrollment group for Windows hosts	Windows	1 host
Desktops - All Types	4.5	Default group for systems that install the Desktop agent kit	Windows	
Desktops - Remote	4.5	Systems that may operate from a remote network	Windows	
Servers - All Types	4.5	Default group for systems that install the Server agent kit	Windows	
Servers - Apache Web Servers	4.5	Systems running the Apache web server	Windows	
Servers - DHCP and DNS Servers	4.5	Systems running DHCP and DNS servers	Windows	
Servers - Externally Deployed	4.5	Default group for servers deployed on public networks	Windows	
Servers - IIS Web Servers	4.5	Systems running Microsoft IIS web server	Windows	
Servers - SQL Server 2000	4.5	Systems running Microsoft SQL Server 2000 database server	Windows	
Systems - IDS Mode	4.5	Systems running in intrusion detection mode with no preventive capabilities	Windows	
Systems - Mission Critical	4.5	Systems that need to be monitored at a higher priority	Windows	
Systems - Restricted Networking	4.5	Systems which are under network lockdown	Windows	
Systems - Test Mode	4.5	Systems operating in test mode	Windows	
VMS CiscoWorks Systems	4.5	Systems running the CiscoWorks VMS product bundle	Windows	1 host

Host groups reduce the administrative burden of managing a large number of Agents. Grouping hosts together lets you apply the same policy to hosts with similar security requirements. A group is the only element required to build Agent kits. When hosts register with the CSA MC, they are automatically put into their assigned group or groups. Once hosts are registered you can edit their grouping at any time.

In the Quick Start configuration example used in this lesson, you will use the Servers – IIS Web Servers for Windows group. The IIS Web Servers group requires no additional configuration, but the preconfigured policies can be examined by choosing **Systems > Groups** and clicking the **Servers – IIS Web Servers** link in the lower box, which holds the Windows default groups.

Selecting a Default Group (Cont.)



The Systems > Groups > Servers – IIS Web Servers window displays deployment configuration options and the policies attached to this group.

Caution It is recommended that you allow the installation program to install the preconfigured CSA MC Agent kit on the CSA MC system. It provides the appropriate security policies for protecting the CSA MC.

Sending Agent Kit URL to Host

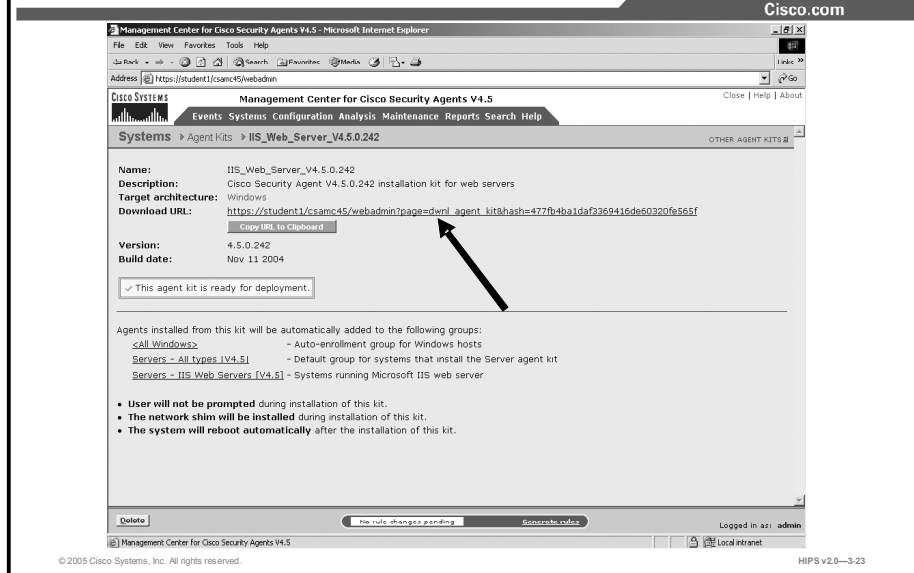
The screenshot shows the Cisco Management Center for Cisco Security Agents V4.5 interface. The page title is "Management Center for Cisco Security Agents V4.5" and the breadcrumb navigation is "Systems > Agent Kits". The table below lists the available agent kits:

Name	Status	Description	Architecture
Apache_Web_Server_V4.5.0.242	Ready	Cisco Security Agent V4.5.0.242 installation kit for web servers	Windows
CiscoWorks_VMS_V4.5.0.242	Ready	Cisco Security Agent V4.5.0.242 installation kit for systems running the Management Center for Cisco Security Agents	Windows
Desktop_V4.5.0.242	Ready	Cisco Security Agent V4.5.0.242 installation kit for desktops	Windows
External_Apache_Web_Server_V4.5.0.242	Ready	Cisco Security Agent V4.5.0.242 installation kit for web servers	Windows
External_IIS_Web_Server_V4.5.0.242	Ready	Cisco Security Agent V4.5.0.242 installation kit for web servers	Windows
IIS_Mode_Desktop_V4.5.0.242	Ready	Cisco Security Agent V4.5.0.242 installation kit for desktops running in test mode	Windows
IIS_Mode_Server_V4.5.0.242	Ready	Cisco Security Agent V4.5.0.242 installation kit for servers running in detection mode	Windows
IIS_Web_Server_V4.5.0.242	Ready	Cisco Security Agent V4.5.0.242 installation kit for web servers	Windows
Remote_Desktop_V4.5.0.242	Ready	Cisco Security Agent V4.5.0.242 installation kit for remote laptops	Windows
Server_V4.5.0.242	Ready	Cisco Security Agent V4.5.0.242 installation kit for servers	Windows
SQL_Web_Server_V4.5.0.242	Ready	Cisco Security Agent V4.5.0.242 installation kit for web servers	Windows

An arrow points to the "IIS_Web_Server_V4.5.0.242" link in the table. The interface also shows a "New" button, a "Delete" button, and a "Generate rules" button. The status bar indicates "No rule changes pending" and "Logged in as: admin".

You can obtain the Agent kit URL for the Web Servers group by choosing **Systems > Agent Kits** and then clicking the **IIS_Web_Server_V4.5.0.242** link in the lower (for Windows) Agent Kits box.

Sending Agent Kit URL to Host (Cont.)

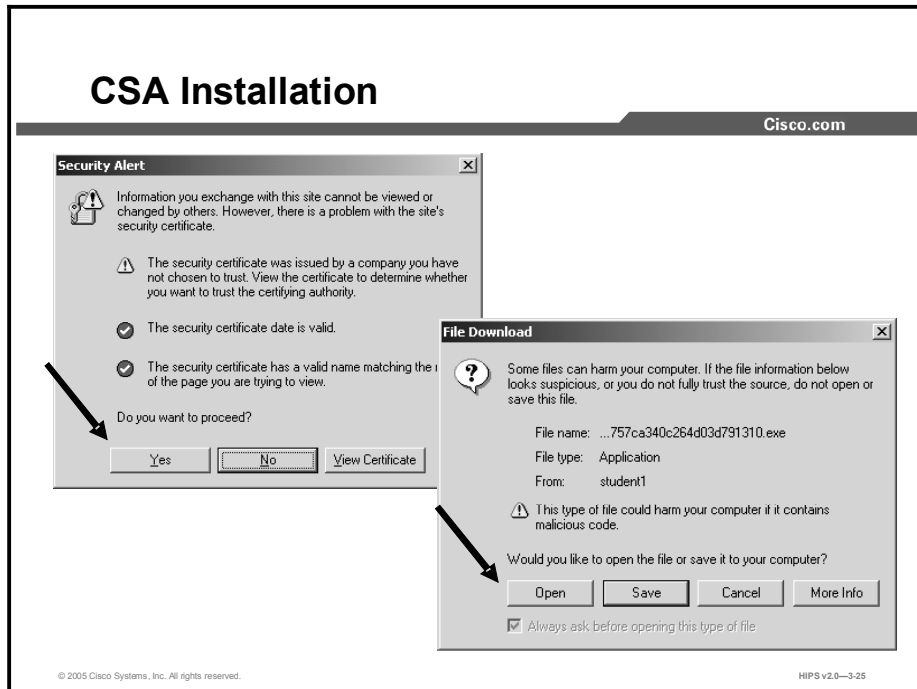


You can distribute this URL, via e-mail for example, to the host systems that the kit is designated for. They access the URL to download and then install the kit.

Note If you type the URL rather than cutting and pasting it, remember that the spaces that appear between the characters in the URL are actually underscore characters.

Installing the CSA

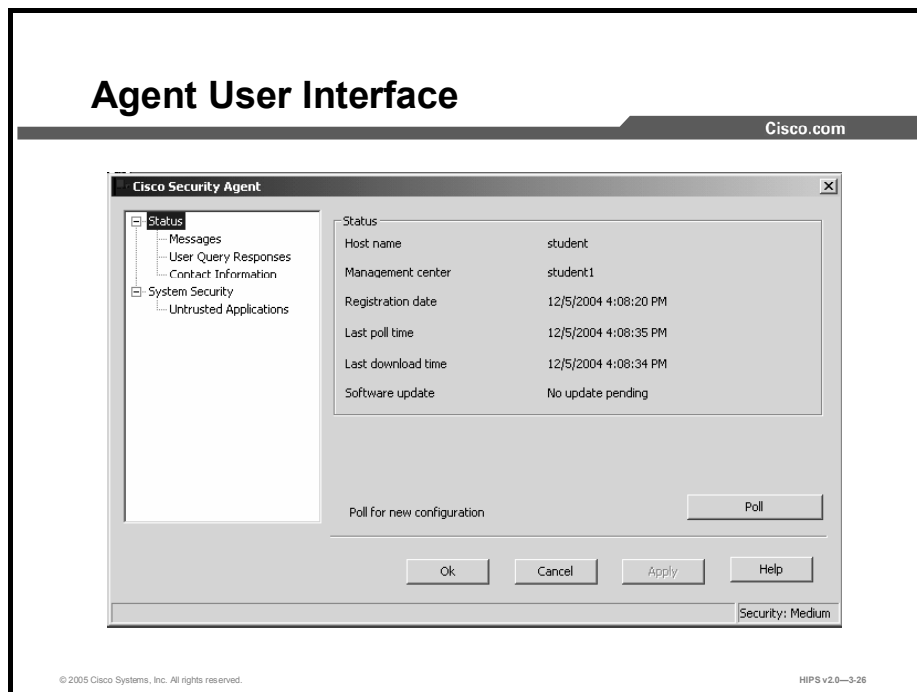
This topic discusses the deployment of the CSA to a host.



You must have local administrator privileges to install CSA on a host. To begin installation, enter the Agent kit URL in your browser, or click **Start > Run** and enter the URL on the run line. A succession of alert messages may open. Click **Yes** and **Open** to proceed with the installation.

Once you successfully download and install Agents, the system will inform you that it will reboot in 2 minutes. When the system restarts, the Agent service starts immediately, and the flag icon appears in the system tray. At this time, the Agent automatically and transparently registers with the CSA MC. The Agent is now ready to receive rules and begin protecting the host.

Agent User Interface



To open the Agent user interface, end users can double-click the flag icon in their system trays. The user interface opens on the desktop. Most fields are read-only status displays.

You can view successfully registered hosts by choosing **Systems > Hosts** from the menu bar on the CSA MC.

Summary

This topic summarizes what you have learned in this lesson.

Summary

Cisco.com

- **Organizing hosts into groups on the CSA MC makes administration of security policies for the hosts easier.**
- **CSA MC administrators can be given different levels of database access:**
 - **Configure: Full access**
 - **Deploy: Most management access**
 - **Monitor: Read-only access**
- **The CSA MC installs with preconfigured groups for many network host desktops and servers.**
- **CSA default groups can be deployed quickly with information available under Agent kits.**
- **CSA requires a post-installation reboot and begins protecting the system immediately afterward.**

© 2005 Cisco Systems, Inc. All rights reserved.

HIPS v2.0—3-27

Lesson 4

Cisco Security Agent Management Center Administration

Overview

This lesson includes the following topics:

- Objectives
- Using Cisco Security Agent Management Center
- Summary

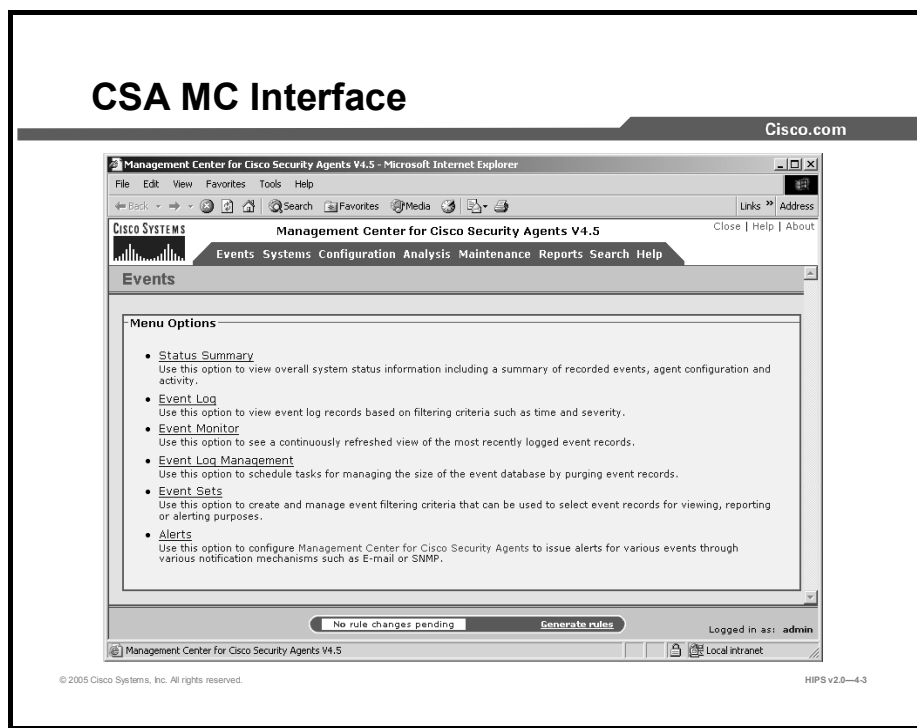
Objectives

Upon completion of this lesson, you will be able to perform the following tasks:

- Describe various components of the menu bar in the CSA MC interface
- Describe the options available in each menu
- Configure the CSA MC using the menu bar.
- Create, save, and delete data

Using Cisco Security Agent Management Center

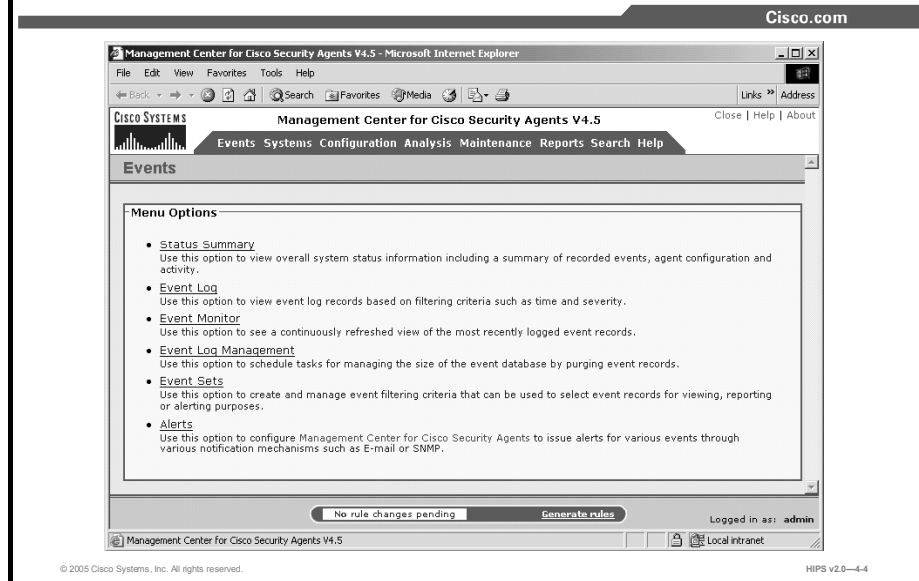
This topic describes the various components you should understand in order to configure Cisco Security Agents (CSAs) using Cisco Security Agent Management Center (CSA MC).



The CSA MC supports editing of the database by multiple administrators. It also provides role-based administration, allowing some administrators to edit configurations, while others can only monitor status. 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.

The menu bar at the top of the CSA MC window provides links to all configuration windows and list views. Arrows indicate that you can choose subcategories from the top-level items. The subcategories appear when you move the mouse over the main item. You may also click the item from the main menu to view the options available for each drop-down menu.

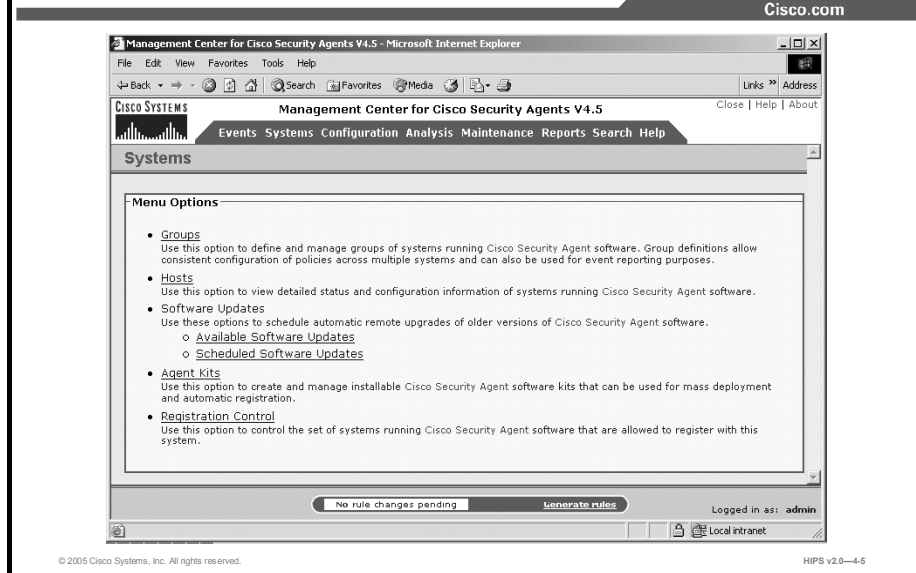
CSA MC Events Menu



The Events menu allows the administrator to view events and configure event settings. The options available are as follows:

- **Status Summary:** Use this option to view overall system status information, including a summary of recorded events, agent configuration, and activity.
- **Event Log:** Use this option to view event log records based on filtering criteria such as time and severity.
- **Event Monitor:** Use this option to see a continuously refreshed view of the most recently logged event records.
- **Event Log Management:** Use this option to schedule tasks for managing the size of the event database by purging event records.
- **Event Sets:** Use this option to create and manage event filtering criteria that can be used to select event records for viewing, reporting, or alerting purposes.
- **Alerts:** Use this option to configure Management Center for Cisco Security Agents to issue alerts for various events through various mechanisms such as e-mail or Simple Network Management Protocol (SNMP).

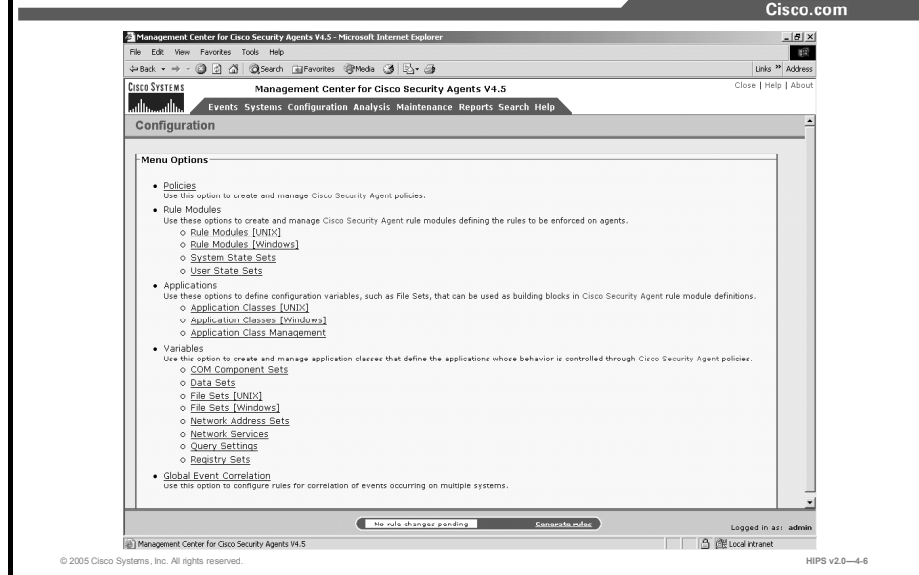
CSA MC Systems Menu



The Systems menu allows the administrator to view and configure groups, hosts, software updates, Agent kits, and registration control. The options available are as follows:

- **Groups:** Use this option to define and manage groups of systems running CSA software. Group definitions allow consistent configuration of policies across multiple systems and can also be used for event reporting purposes.
- **Hosts:** Use this option to view detailed status and configuration information of systems running CSA software.
- **Software Updates:** Use these options to schedule automatic remote upgrades of older versions of CSA software.
 - Available software updates
 - Scheduled software updates
- **Agent Kits:** Use this option to create and manage installable CSA software kits that can be used for mass deployment and automatic registration.
- **Registration Control:** Use this option to control the set of systems that are running CSA software and are allowed to register with this system.

CSA MC Configuration Menu



The Configuration menu allows the administrator to view and configure policies, rule modules, applications, variables, and global event correlation. The options available are as follows:

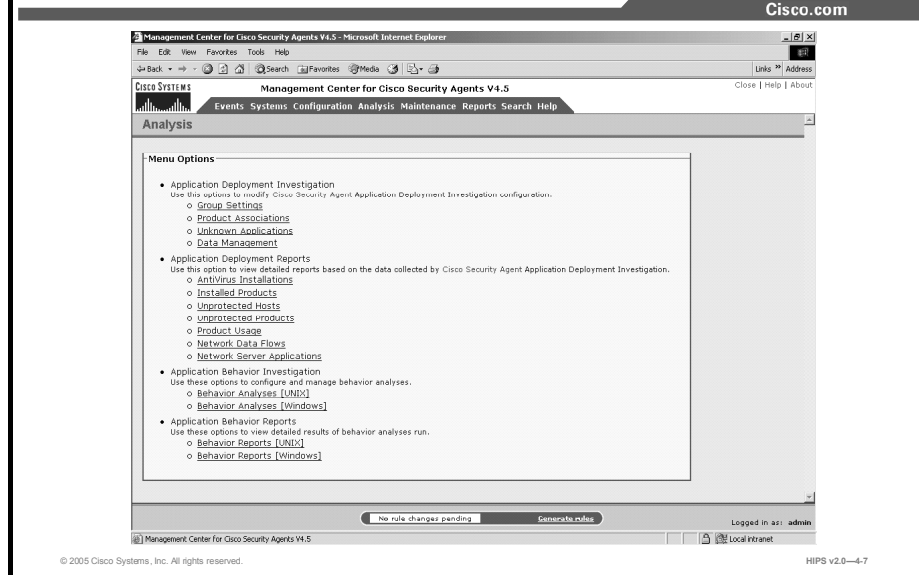
- **Policies:** Use this option to create and manage CSA policies.
- **Rule Modules:** Use these options to create and manage CSA rule modules that define the rules to be enforced on Agents.
 - Rule Modules (UNIX)
 - Rule Modules (Windows)
 - System State Sets
 - User State Sets
- **Applications:** Use these options to define configuration variables, such as file sets, that can be used as building blocks in CSA rule module definitions.
 - Application Classes (UNIX)
 - Application Classes (Windows)
 - Application Class Management
- **Variables:** Use this option to create and manage application classes that define the applications whose behavior is controlled through CSA policies.
 - COM Component Sets
 - Data Sets
 - File Sets (UNIX)
 - File Sets (Windows)
 - Network Address Sets
 - Network Services

— Query Settings

— Registry Sets

- **Global Event Correlation:** Use this option to configure rules for correlation of events that occur on multiple systems.

CSA MC Analysis Menu

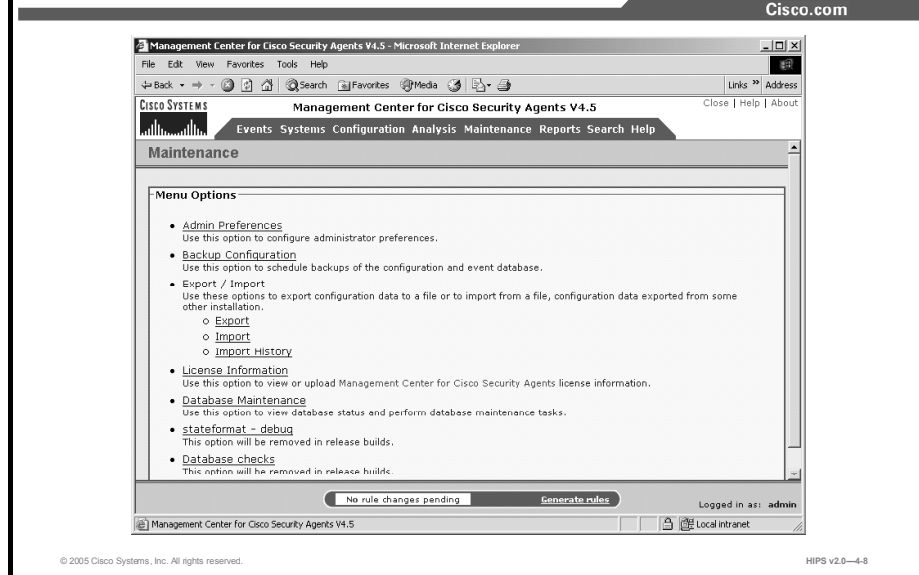


The Analysis menu allows the administrator to view and configure application deployment investigation, application behavior investigation, and reports for both. The options available are as follows:

- **Application Deployment Investigation:** Use these options to modify CSA application deployment investigation configuration.
 - Group Settings
 - Product Associations
 - Unknown Applications
 - Data Management
- **Application Deployment Reports:** Use this option to view detailed reports based on the data collected by CSA application deployment investigation.
 - Antivirus Installations
 - Installed Products
 - Unprotected Hosts
 - Unprotected Products
 - Product Usage
 - Network Data Flows
 - Network Server Applications
- **Application Behavior Investigation:** Use these options to configure and manage behavior analysis.
 - Behavior Analysis (UNIX)
 - Behavior Analysis (Windows)

- **Application Behavior Reports:** Use these options to view detailed results of behavior analyses that you have run.
 - Behavior Reports (UNIX)
 - Behavior Reports (Windows)

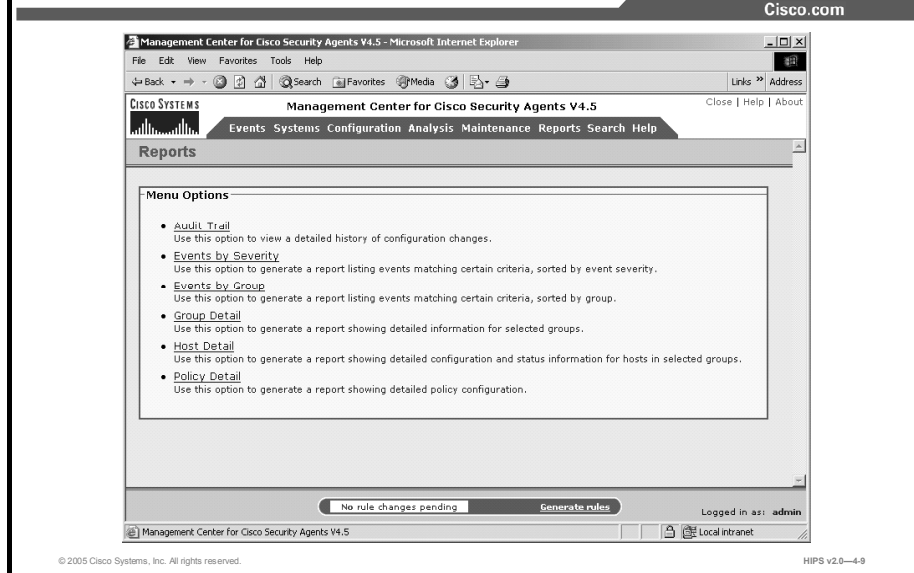
CSA MC Maintenance Menu



The Maintenance menu allows the administrator to view and configure administrative preferences, backup configuration, export and import configuration data, license information, and database maintenance. The options available are as follows:

- **Admin Preferences:** Use this option to configure administrator preferences.
- **Backup Configuration:** Use this option to schedule backups of the configuration and event database.
- **Export/Import:** Use these options to export configuration data to a file or, when configuration data has been exported from some other installation, to import it from a file.
 - Export
 - Import
 - Import History
- **License Information:** Use this option to view or upload CSA MC license information.
- **Database Maintenance:** Use this option to view database status and perform database maintenance tasks.

CSA MC Reports Menu



The Reports menu allows the administrator to view and configure reports according to specified categories. The options available are as follows:

- **Audit Trail:** Use this option to view a detailed history of configuration changes.
- **Events by Severity:** Use this option to generate a report that lists events matching certain criteria, sorted by event severity.
- **Events by Group:** Use this option to generate a report that lists events matching certain criteria, sorted by group.
- **Group Detail:** Use this option to generate a report that shows detailed information for selected groups.
- **Host Detail:** Use this option to generate a report that shows detailed configuration and status information for hosts in selected groups.
- **Policy Detail:** Use this option to generate a report that shows detailed policy configuration.

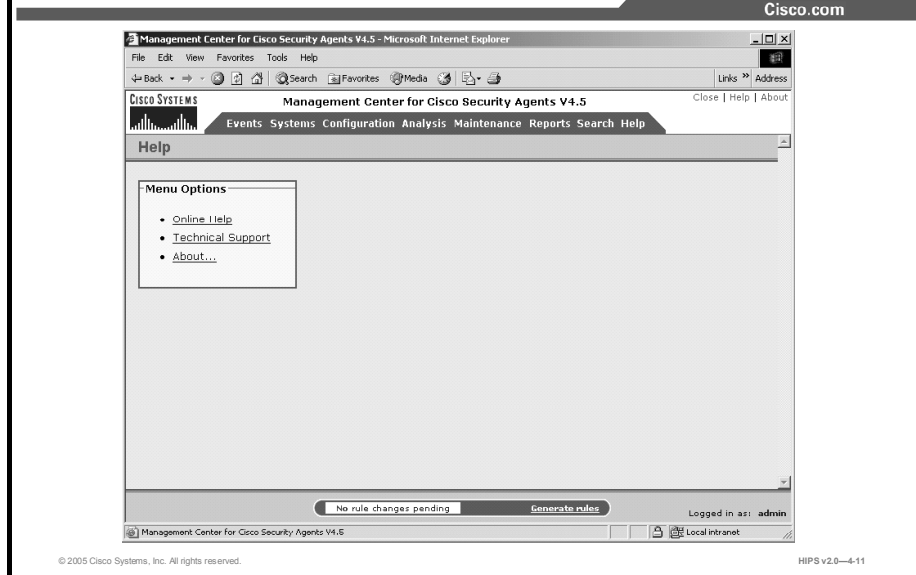
CSA MC Search Menu



The Search menu allows the administrator to search the database according to specified criteria. The options available are as follows:

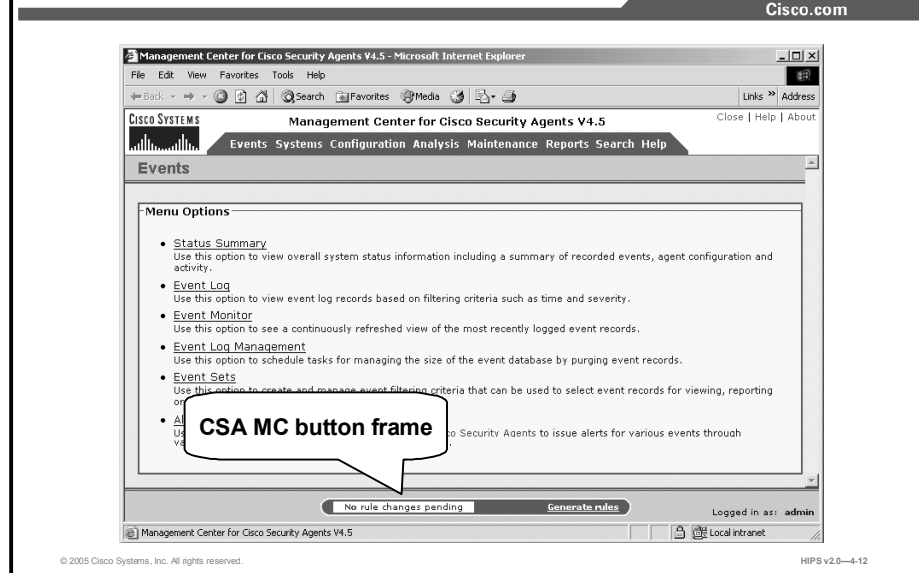
- **Hosts:** Use this option to search host records. Additional search criteria such as host names, group membership, polling activity, software versions, etc. may be specified.
- **Groups:** Use this option to search group records.
- **Policies:** Use this option to search policies. Additional search criteria that may be specified include group associations.
- **Rule Modules:** Use this option to search rule modules. Additional search criteria that may be specified include policy associations, rules, etc.
- **Rules:** Use this option to search policy rules. Additional search criteria such as rule IDs, associated actions, rule state, etc. may be specified.
- **Variables:** Use this option to search various types of configuration variables.
- **Application Classes:** Use this option to search application classes.
- **All:** Use this option to search the entire configuration database for all types of items. Note that this does not include the event database.

CSA MC Help Menu



The Help menu allows the administrator to use online tools and technical support information on the CSA MC as well as to get information about the product version and other topics.

CSA MC Button Frame: Creating, Saving, and Deleting Data



All CSA MC action items appear in a frame at the bottom of the CSA MC window. The buttons in this frame change in accordance with the actions available for the window that you are 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 CSA systems, you must click this link in the button frame first to view all pending database changes and then to generate them.

Note Most list view windows in the CSA MC contain New, Clone, and Delete buttons. (Clone is not present in all list view windows because you can clone only 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 window for that item.
- **Clone:** Use the Clone button in conjunction with the check boxes beside each list view item. To clone a particular configuration, check its check box 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 check boxes beside each list view item. To delete a configuration, check its check box (you can check several at once) and click the **Delete** button. All checked items are deleted. To quickly check all check boxes, check the top check box 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 after you have finished in order to save your configuration in the CSA MC database. If you do not click Save before moving to another window in the 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.

- **Compare:** Policies, variables, and application classes provide a Compare button in their list views. When you check the check boxes next to two items (you cannot compare more than two configurations at a time) and click the Compare button, the CSA MC displays the configurations side by side and highlights the differences in red. After you have examined how the configurations compare, you can choose to merge them.

The purpose of the Compare tool is to assist you after you have imported configurations or upgraded the 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.

Note Right-clicking your mouse on a CSA MC window displays a shortcut menu for performing the tasks provided by buttons on that window and for additional configuration tasks that are not as easy to access from your current window.

Summary

This topic summarizes what you have learned in this lesson.

Summary

Cisco.com

- Administration of the CSA MC is role-based.
- The CSA MC provides a menu bar for configuration tasks.
- The following configuration options are available on the menu bar:
 - Events
 - System
 - Configuration
 - Analysis
 - Maintenance
 - Reports
 - Search
 - Help
- Creating, saving, and deleting data are done using the CSA MC button frame.

© 2005 Cisco Systems, Inc. All rights reserved. HIPS v2.0—4-13

Lesson 5

Configuring Groups and Managing Hosts

Overview

This lesson explains the configuration of hosts in the Cisco Security Agent Management Center (CSA MC) and the utilization of groups to make security policy deployment easier and more effective. This lesson includes the following topics:

- Objectives
- Configuring groups
- Building an Agent kit
- Managing hosts
- Deploying scheduled software updates
- Summary
- Lab exercise

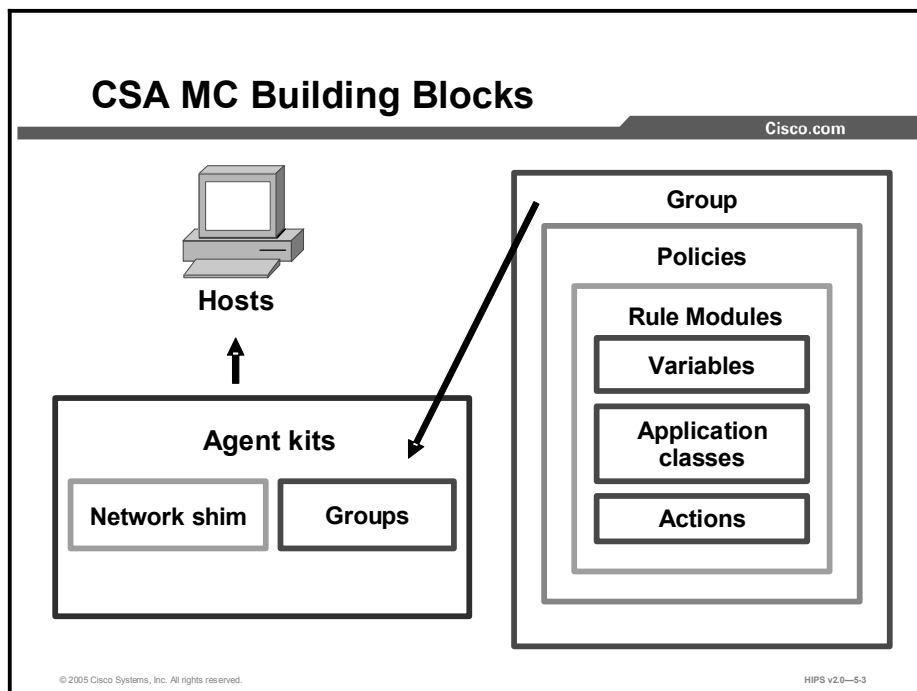
Objectives

Upon completion of this lesson, you will be able to perform the following tasks:

- Create groups to ease host management and security policy deployment
- Build Agent kits for the newly created groups
- View host status and modify host configuration
- Distribute software updates to hosts

Configuring Groups

This topic discusses the configuration of groups in the CSA MC.

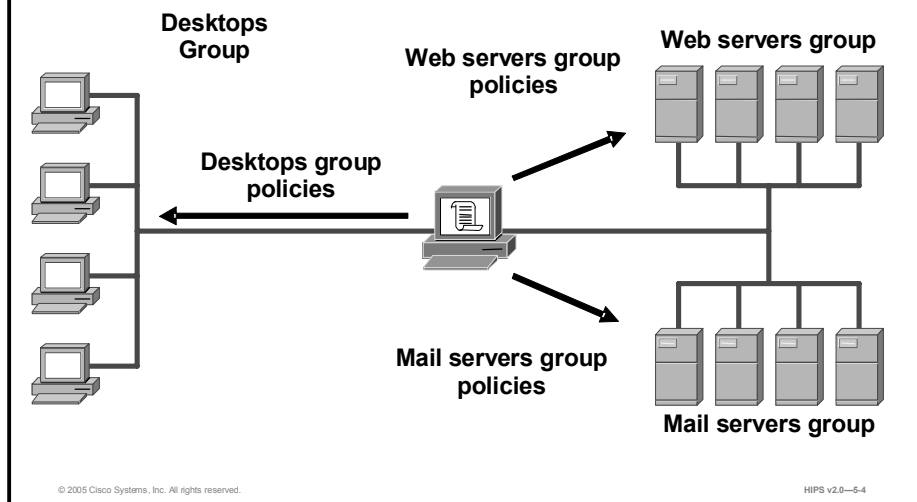


The figure illustrates the various components that you use to form the groups contained in Agent kits:

- Variables, application classes, and actions combine to create rules.
- Rule modules contain rules, variables, application classes, and actions and combine to form policies.
- Policies contain rules and can be applied to a group or multiple groups.
- Groups contain associations with policies; they accept hosts as members.
- Agent kits contain groups and (optionally) the network shim. Agent kits are deployed to hosts; they install the CSA software and all of the policies and rules that have been built into them.

Configuring Groups

Cisco.com



System hosts across your network, including mobile systems in the field, must download Cisco Security Agent (CSA) software and register with the CSA MC to receive the security policies configured for them. When you are ready to apply policies to the hosts that are running Agents, having those hosts placed into common groups streamlines the process of assigning policies to several hosts at once. Using groups reduces the administrative burden of managing a large number of Agents.

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.

The CSA MC ships with several preconfigured groups for you to use. If the included groups do not suit your needs, use the instructions in this lesson to configure new groups or to edit existing ones.

Advantages to Forming Hosts into Groups

Cisco.com

Groups allow you to

- **Apply the same set of policies across multiple host systems.**
- **Apply alerts and event set parameters based on group configurations.**
- **Use test mode to try out policies on groups of hosts before you actively enforce those policies.**
- **Run reports based on group settings.**

© 2005 Cisco Systems, Inc. All rights reserved.

HIPS v2.0--5-5

Grouping individual host systems together provides the following advantages:

- Grouping lets you apply the same set of policies consistently across multiple host systems. Rather than configuring a security policy on each host, you deploy a common policy to any number of hosts grouped by administrator-selected criteria.
- Grouping eases deployment of alerts by applying alerts to many hosts at once. The use of groups sharpens the filtering granularity of event sets, thus improving analysis of network events.
- By using groups, you can use test mode to try out policies on many hosts before you enforce those policies in production.
- Using groups enhances reporting capabilities.

Grouping Criteria

Cisco.com

- **System function**
- **Business groups**
- **Geographical or topological location**
- **Importance to your enterprise**

© 2005 Cisco Systems, Inc. All rights reserved.

HIPS v2.0—5-6

Hosts can be grouped together based on many different criteria:

- **System function:** You can create a security policy that corresponds specifically to the needs of your web servers, for example, and distribute it to that group.
- **Business group:** You can distribute policies based on the needs of each business group, such as finance, operations, or marketing.
- **Geographical or topological location:** You can group hosts based on their subnet, office, or data center location, for reporting purposes.
- **Importance to your enterprise:** You can 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 hosts to a group or remove them at any time. However, the policy configuration of a host that is moved to another group will not take effect until you generate your rule programs and distribute them.

Mandatory Group Enrollment

Cisco.com

CSA MC provides three auto-enrollment architectural groups that are mandatory for all hosts of a given OS:

- **Windows**
- **Solaris**
- **Linux**

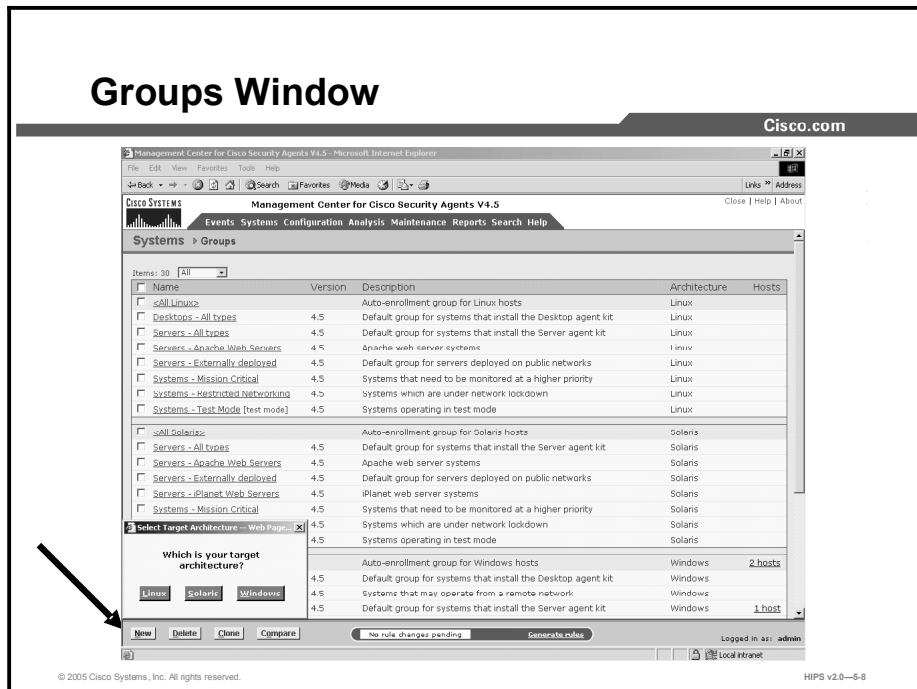
© 2005 Cisco Systems, Inc. All rights reserved.

HIPS v2.0--5-7

CSA MC provides three auto-enrollment architectural groups (Windows, Solaris, and Linux) that are mandatory for all hosts of a given architecture. For example, all Windows hosts are enrolled in the <All Windows> group when they register with the CSA MC. This is in addition to any other groups assigned by the administrator. Hosts cannot be removed from these mandatory groups.

By providing group auto-enrollment for hosts, any policies attached to these hosts also become mandatory by association. An administrator could use these mandatory groups to apply policies that prevent some critical service from being inadvertently banned. For example, you could attach policies to prevent the Dynamic Host Configuration Protocol (DHCP) from being disabled by an overly restrictive rule.

Groups Window



When hosts across your network download and install Agent kits, they automatically and transparently register with the CSA MC. Hosts inherit membership to the groups that are associated with the Agent kit that they install.

Complete the following steps to configure a group:

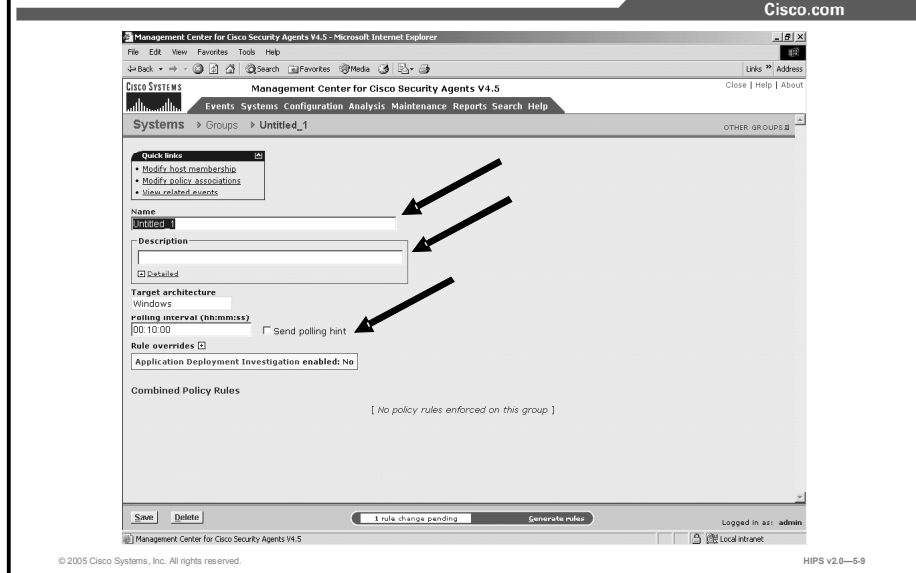
Step 1 Move the mouse over **Systems** in the menu bar and select **Groups** from the drop-down list that appears. The list of existing groups is displayed.

Note Management Center for Cisco Security Agents ships with several preconfigured 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 designated "All" as the operating system type for your administrator session, you are prompted to select whether this is a Windows, Solaris, or Linux group. You cannot combine hosts of differing OS architectures in the same group.

Groups Configuration Window



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 is 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.

Step 4 Optionally, you can change the default polling interval to any value between 10 seconds and 24 hours (formatted as hh:mm:ss). This controls how often Agents in this group poll into CSA MC for policy updates. Decreasing 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 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 5 Optionally, enable the Send Polling Hint capability. Normally, if you make changes to a policy, schedule a software update, or make any other change to a host's configuration, the host does not receive that change until it next polls into the MC. But if you have checked the Send Polling Hint checkbox, certain changes that occur on the MC will cause a "non-reliable" signed User Datagram Protocol (UDP) message to be sent to the appropriate hosts. This message tells hosts to poll into the MC earlier than their next scheduled polling interval. The UDP message would be sent if a policy change occurred, if a global correlation caused a file to be added to

the global quarantine list, and if you selected to retrieve status information from a particular host.

Step 6 Optionally, enable one or more Rule Overrides for the group. 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.

Step 7 Optionally, enable Verbose Logging Mode to change the event log timer to log all reoccurring events rather than suppressing duplicates.

Step 8 Optionally, enable Log All Deny Actions to turn on logging for all deny rules that are running on hosts within the group regardless of the individual rule settings for the policy attached to the group. You may wish to use this feature to turn on all deny logging for diagnostic purposes.

Step 9 Optionally, you can select the Filter User Info from Events checkbox for this group. Due to privacy issues, you may not want this username information displayed in events or in the additional information screen available from the event Details link.

Step 10 Optionally, for Windows groups, you can enable Application Deployment and Analysis. This analysis functionality works with CSA MC and the Agent, serving as a data collection tool for administrators deploying policies across systems and networks. If this feature is enabled, you can access analysis reports from a link on this page.

Step 11 When all required information is entered, click the **Save** button to enter and save your group in the CSA MC database.

Building an Agent Kit

In this section you will learn about building an Agent kit.

Agent Kits

Cisco.com

- **CSA MC allows the creation of custom Cisco Security Agent installation kits to greatly reduce the administrative burden required to deploy Agents on new systems.**
- **Upon creation, new Agent kits may be associated with one or more groups.**

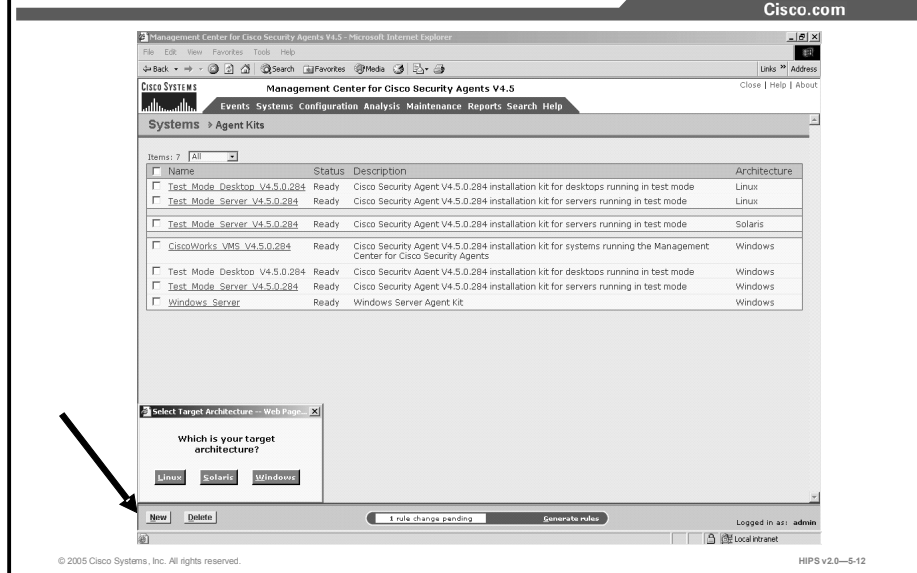
© 2005 Cisco Systems, Inc. All rights reserved. HIPS v2.0—5-11

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. When you create the Cisco Security Agent kit, you have the option of associating it with one or more groups. The particular Agent kit that 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 are associated with the installed kit.

Note CSA MC ships with preconfigured Agent kits for desktops, servers, and many other needs. 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.)

Building an Agent Kit

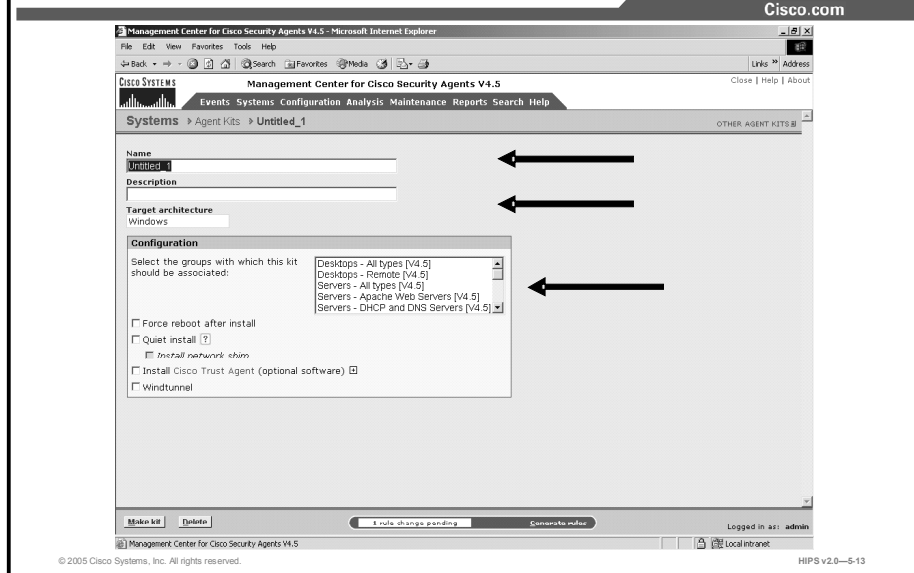


To configure an Agent kit, complete the following steps:

- Step 1** Move the mouse over **Systems** 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 kit.

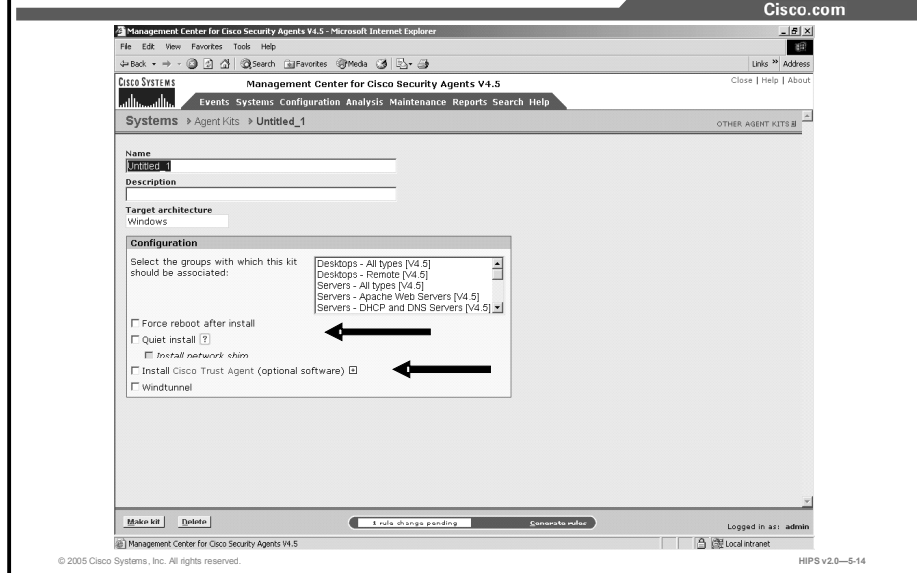
Note If you have designated "All" as the operating system type for your administrator session, you are prompted to select whether this is a Windows, Linux, or Solaris kit. You cannot select a Solaris group for an Agent kit that you have configured for Windows systems.

Building an Agent Kit (Cont.)



- Step 3** In the Agent kit configuration view, enter a unique name for this kit. (You cannot use spaces in Agent kit names.) Generally, it is 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.
- 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 Control key as you select each item. To deselect a single item, hold down the Control 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.

Building an Agent Kit (Cont.)



Step 6 Select whether or not to have Agents install "quietly" on end-user systems (Windows and Linux only). Quiet Install requires users to download the self-extracting executable as does the nonquiet install. The difference is that in a Quiet Install, no prompts appear and the user is not required to enter any information or select any options. A nonquiet install prompts the user for installation options, such as enabling the network shim, in addition to the reboot prompt.

Step 7 For Windows kits, if you select Quiet Install, you can also select whether or not the network shim is installed 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 Virtual Private Network (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 the network shim, you would create kits as nonquiet installations. (Do not select the Quiet Install checkbox.) This way, users are prompted to enable the network shim during the Agent installation.

Note Not enabling the network shim does not mean that network access control rules will not work. It only means that the system hardening features (configured in the Network Shield rule page) are not enabled.

Step 8 If you select Quiet Install, you can also select whether the system is automatically rebooted once the installation 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. Windows NT systems must be rebooted after an Agent installation.

Step 9 You can optionally install the Cisco Trust Agent (CTA) with the Cisco Security Agent. The fields found here also allow you to specify CTA initialization settings.

Step 10 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. You may distribute this URL, via e-mail for example, to the host systems that 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>/csamc45/kits>. If you are pointing users to the kits URL and you have multiple Agent kits listed, be sure to tell users which kits to download.

Note The Registration Control feature also applies to the <ciscoworks system name>/csamc45/kits URL. If the Registration Control feature prevents your IP address from registering, it also prevents you from viewing this kit's URL.

Note You must regenerate your rule program after Agent kits are created.

Note If you installed Management Center for Cisco Security Agents to the default directory, all Agent kits are placed in the %Program Files%\CSCOp\CSAMC45\bin\webserver\htdocs\deploy_kits directory.

Agent Kit Status

Cisco.com

When you create an Agent kit it is given one of three status levels:

- **Ready**
- **Needs rule generation**
- **Incomplete**

© 2005 Cisco Systems, Inc. All rights reserved.

HIPS v2.0—5-15

When you create an Agent kit, it is given one of three status levels based on how far into the configuration you have 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 of the necessary parameters for this Agent kit. You must complete the configuration and then generate rules before the kit can be downloaded.

Agent Reboot vs. No Reboot

Cisco.com

When the Agent kit has completed installation and is not rebooted, some CSA functionality of Windows, Solaris, and Linux hosts is not available.

© 2005 Cisco Systems, Inc. All rights reserved.

HIPS v2.0—5-16

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 System API page for Windows) is only enforced for new processes.
- COM component access control rules are only enforced for new processes.
- Data access control rules are not applied until the web server service is restarted.

■ Solaris and Linux Agents

- 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).

Scripted Agent Installs and Uninstalls

Cisco.com

You can use scripts to complete the following functions on Windows Cisco Security Agent kits on end-user systems:

- **Scripted install**
- **Scripted uninstall**

© 2005 Cisco Systems, Inc. All rights reserved.

HIPS v2.0—5-17

You can use scripts to silently install and uninstall Windows Cisco Security Agents on end-user systems. (Linux and Solaris do not support scripted Agent installs and uninstalls.)

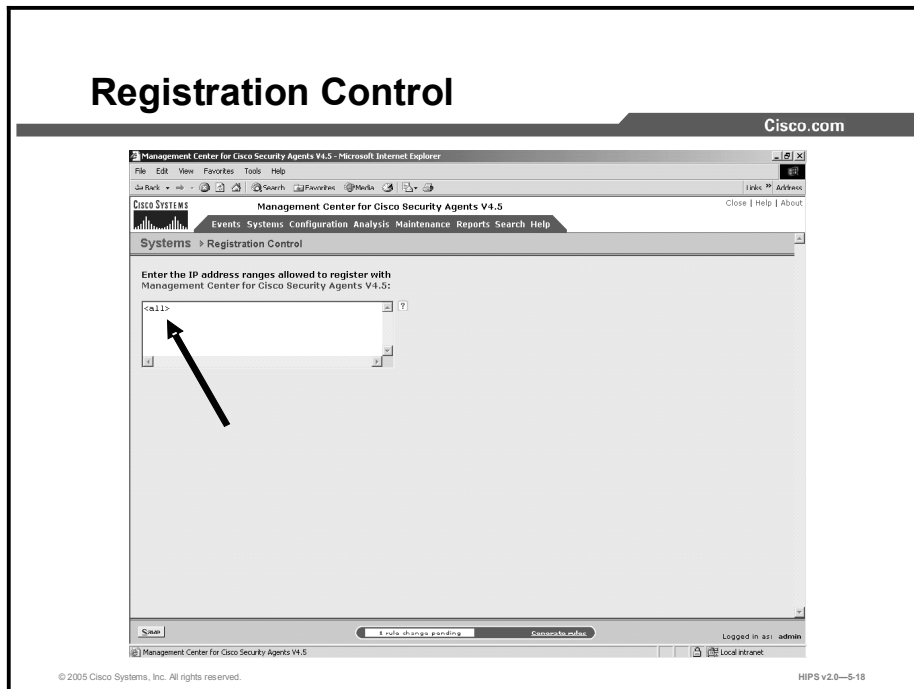
- **Scripted install:** The Agent kit is a self-extracting executable placed in the following directory on the server: %Program Files%\CSCOp\CSAMC45\bin\webserver\htdocs\deploy_kits. (Retrieve the kit from this directory or download it from the server.) You can then use a script to copy and silently install Agent kits on systems.

Note 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



The Registration Control feature, which is accessible from the Maintenance item in the menu bar, prevents unauthorized hosts from downloading Agent kits and receiving rules. On the Registration Control page, you enter a range of addresses, which restricts Agent hosts attempting to successfully register with CSA MC to those with addresses listed here.

Note Any user who is logged in to CSA MC can download an Agent kit.

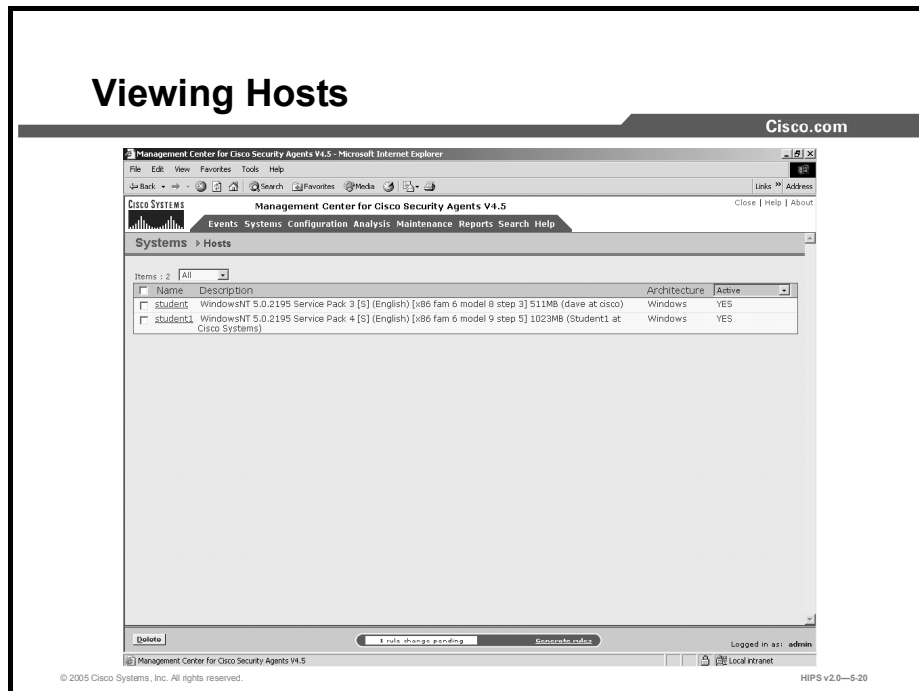
The default entry here is <all> (0.0.0.0-255.255.255.255), which applies no address registration restrictions. An example entry of restricted registration addresses is as follows. (Only those addresses within the range listed can register. This range is inclusive.)

192.168.10.0-192.168.10.255

172.16.20.0-172.16.20.255

Managing Hosts

In this section you will learn about managing hosts.



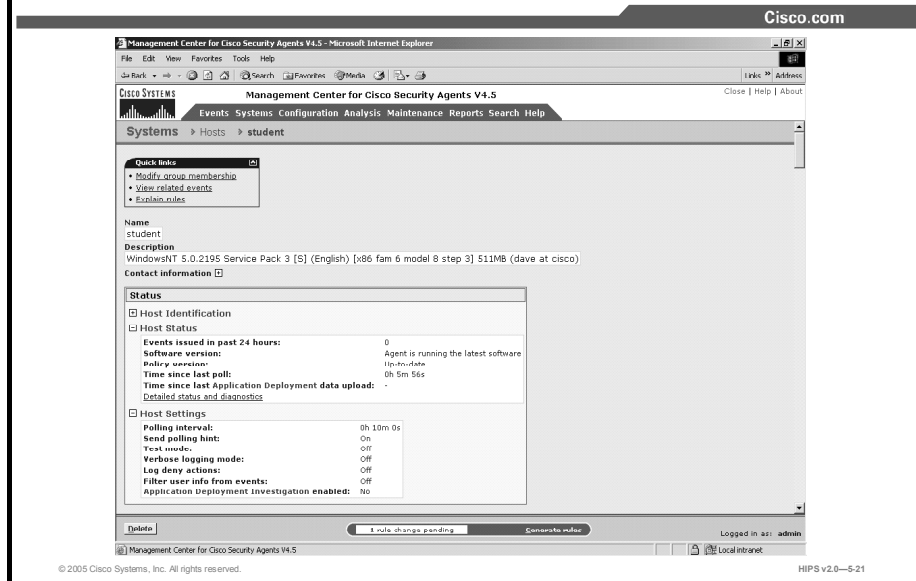
You can see which hosts have successfully registered with the CSA MC by choosing Systems > Hosts. You can view the current status overview of hosts by moving your mouse over Events in the menu bar and clicking Status Summary in the drop-down list. You can also use the Search<Hosts page to locate hosts based on some of the following information:

- **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 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.

Host Detail



Click on the host link itself for detailed host information. In the Host Detail page, these additional options and information are available:

- Click the Modify Group Membership link in the Host Detail page 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.

Host Detail (Cont.)

The screenshot displays the Cisco Management Center for Cisco Security Agents V4.5 interface. The main content area is titled "Host Detail" and is divided into several sections:

- Status**
 - Host Identification:** Product information: Cisco Security Agent Version 4.5.0.284; Last known IP address: 10.0.1.50 (lastpoll); Host ID: 111; UID: (9CF86C79-E284-4F5D-9B65-DBFC6749F7FD); Registration time: 12/02/2004 9:18:05 AM; Operating system: Windows 2000 [WindowsNT 5.0.2195 SP 3 Server;English]; Cisco Trust Agent installed: No.
 - Host Status:** Events issued in past 24 hours: 0; Software version: Agent is running the latest software; Policy version: Up-to-date; Time since last poll: 0h 2m 23s; Time since last Application Deployment data upload: [Link]; Detailed status and diagnostics: [Link].
 - Host Settings:** Polling interval: 0h 10m 0s; Send polling hint: On; Test mode: Off; Verbose logging mode: Off; Log deny actions: Off; Filter user info from events: Off; Application Deployment Investigation enabled: No.
- Group Membership and Policy Inheritance:** A table listing group and policy information.

Group Name	Version	Description	Policies
<input type="checkbox"/> All Windows		Auto-enrollment group for Windows hosts	4 policies
<input type="checkbox"/> AD Test Policy	4.5	Experimental policy module (caveat emptor)	1 module

At the bottom of the interface, there is a status bar showing "1 rule change pending" and "Connected to server". The footer contains the text "© 2005 Cisco Systems, Inc. All rights reserved." and "HIPS v2.0-6-22".

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. Click on a host to view more detailed information on that host system.

- **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; it does not have to be unique. CSA MC assigns each registering host a unique number by which the database identifies it.
- **Contact Information:** Click this link to view the contact information that the user provided to the Agent. (The available fields for the user are first name, last name, e-mail, telephone, and location.)

The following options are available in the Host Identification pane:

- **Product Information:** This is the Cisco Security Agent version for this particular machine.
- **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 five IP addresses can be listed.)
- **Host ID:** CSA MC assigns each registering host a unique number by which the database identifies it.
- **UID:** This is a globally unique ID for your Agent that 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.
- **Registration Time:** This is the time that the Agent registered with CSA MC.
- **Operating System:** This is the operating system installed on this particular machine.
- **Cisco Trust Agent Installed:** This displays whether optional CTA software is installed on the system. If CTA software is installed, this field also displays the current CTA posture status.

The following options are available in the Host Status pane:

- **Events Issued in Past 24 Hours:** This is the number of events (rule triggers) that have occurred on the host system in the given time frame.
- **Software Version:** This is the version of Cisco Security Agent software that the system is running. If a software update is 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.
- **Policy Version:** This field reads "Up-to-date" or "Not up-to-date," indicating whether the Agent has the latest policy configuration from CSA MC.
- **Time Since Last Poll:** This is the interval since the host system's last polling request.
- **Time Since Last Application Deployment Data Upload:** If application deployment data collection is enabled on the end-user system, this indicates the time of the most recent upload of analysis logging data.
- **Detailed Status and Diagnostics:** Click this link to view status information for the host in question. The window that is opened by this link uploads information from the Agent. You can use this information to diagnose Agent issues, to view the current states and policies running on the Agent system, and to reset the system to factory default settings.

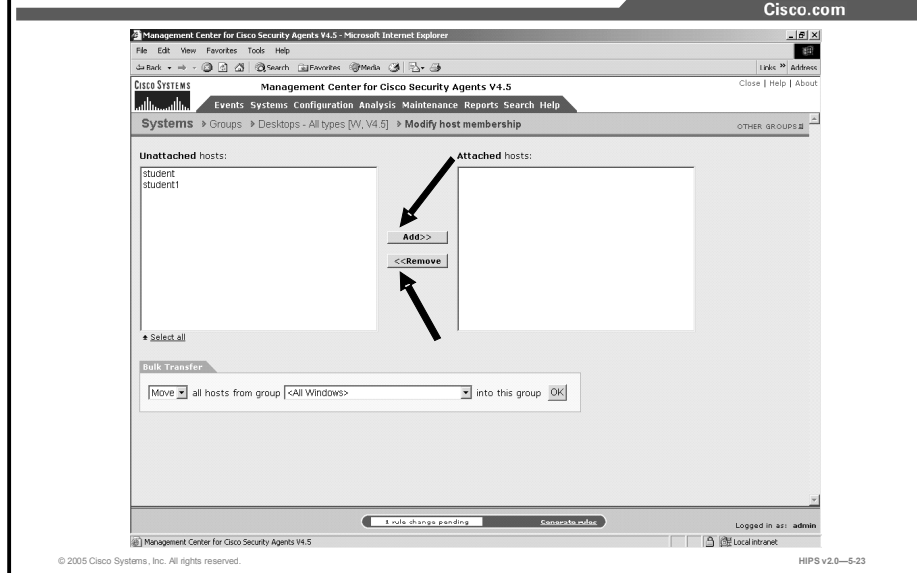
The following options are available in the Host Settings pane:

- **Polling Interval (seconds):** The value shown here indicates the time interval in which this system polls in to the management server. This feature is configurable through the Groups page.
- **Send Polling Hint:** This field indicates whether the polling hint capability is turned on for the group in which this host is a member.
- **Test Mode:** If this host is part of a group operating in test mode, that information is displayed here.
- **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.
- **Log Deny Actions:** This field indicates whether the "log all deny actions" capability is turned on for the group in which this host is a member.
- **Filter User Info from Events:** this field indicates whether the "filter user from events" capability is turned on for the group in which this host is a member.
- **Application Deployment Investigation Enabled:** This appears if the application deployment data collection capability, available from the Analysis menu bar item, is enabled on the end-user system. If this feature is enabled, you can access analysis reports from a link on this page.

Optionally, you can enter contact information such as username, location, e-mail, 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.

Adding Hosts to a Group



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 have 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 effect until you generate your rule programs and distribute them.

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

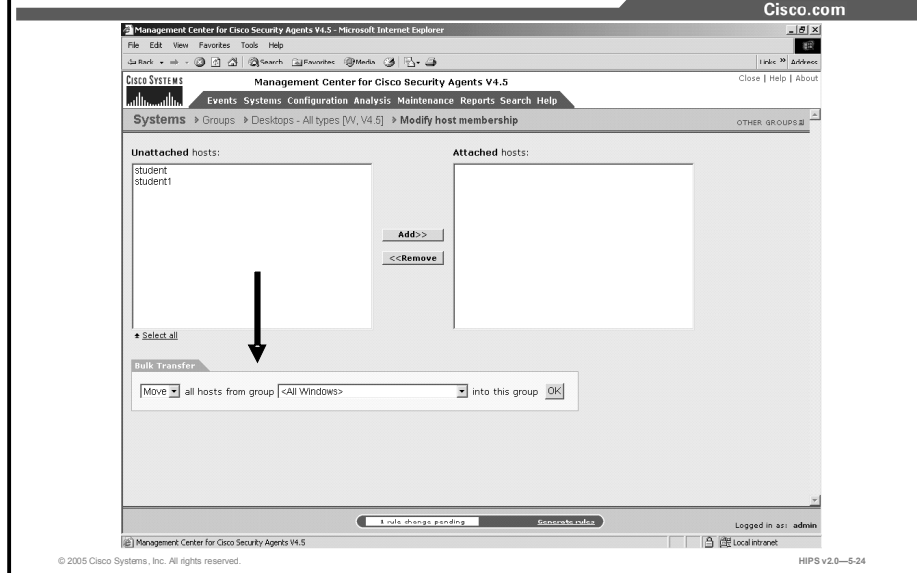
- To add a host to multiple groups, use the Hosts > Modify Group Membership link.
- To add multiple hosts to a single group, use the Group > Modify Host Membership link.
- To move or copy all hosts in one group to another group, use the Bulk Transfer feature accessible from the Group > Modify Host Membership 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 that are in this group (if any) in the box on the right. Hosts listed in the box on the left are not in the group.

- Step 4** To add a host to this group, select the host in the box on the left and click the **Add** button to move it to the box on the right. It is now a part of the group. To select multiple nonsuccessive items in a swap box, hold down the Control key as you select each item. To deselect a single item, hold down the Control 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.
- Step 5** To remove a host from a group, select the host that you wish to remove in the box on the right. Click the **Remove** button. The host moves to the left (unattached) box.

Adding Hosts to a Group (Cont.)



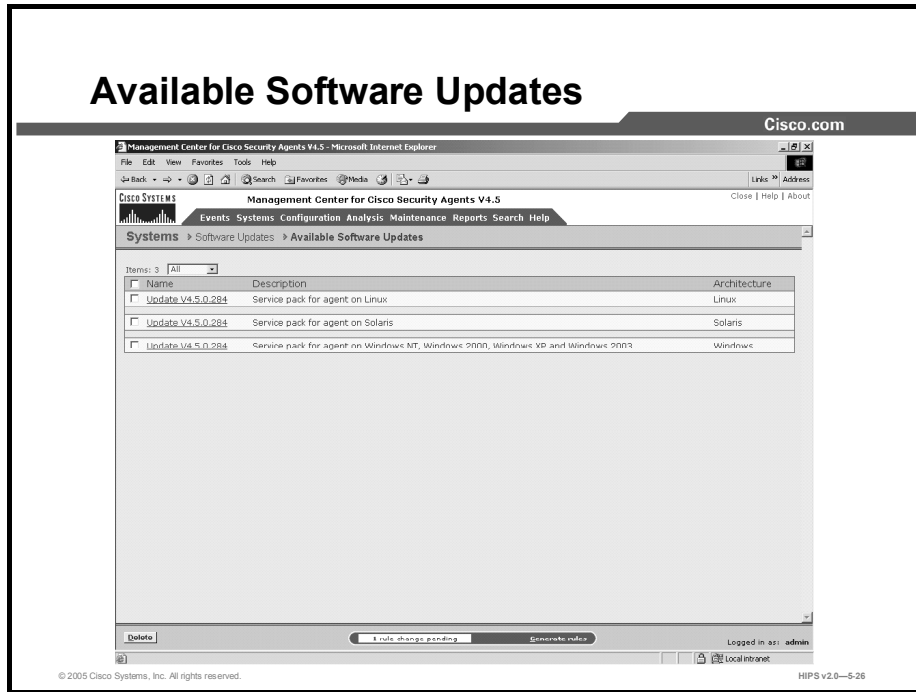
Use the Bulk Transfer feature to move or copy all hosts from the group you select in the available drop-down field, into the group you are currently viewing. When you click the OK button beside the group selection drop-down list, all hosts are moved or copied.

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.)

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 view all new and edited configurations first and then distribute them to the Agents.

Deploying Scheduled Software Updates

In this section you will learn about deploying scheduled software updates.



Cisco provides software updates via its website (<http://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.

The next time Agent systems poll in to the server, the Agent GUI prompts the user that there is a software update available (Windows and Linux only; Solaris 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.

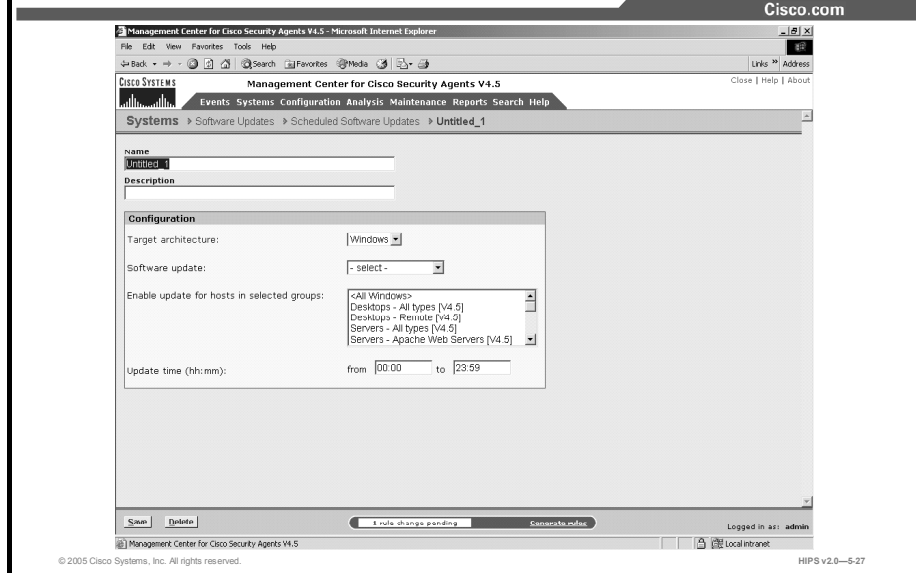
From the Available Software Updates page, you can click on a particular update and view the following information:

- Name and description of the software update
- 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 of the software update
- 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

Caution Always consider bandwidth availability and CSA MC utilization during software updates.

Scheduled Software Updates



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

To create scheduled software updates for distribution to Agent systems, do the following:

- Step 1** From the menu bar Systems drop-down list, move the mouse over **Software Updates**. A cascading menu with further selections appears. Select **Scheduled Software Updates**.
- 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 that you are distributing (Solaris, Linux, or Windows). When you select an OS, the available updates and selectable groups change accordingly.
- Step 6** From the Software Update drop-down list, select the Solaris, Linux, or Windows update that you want to distribute.
- Step 7** From the available groups in the Enable Update for Hosts in Selected Groups list, select one or more to distribute this update to.

Note To select multiple items in a list box, hold down the Control key as you select each item. To deselect a single item, hold down the Control key when you click on the item in question. Press the Shift key to select multiple successive items.

Step 8 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. Putting a time limit on the update may have an undesirable result. If you enter 10:00 to 11:00 (this would be AM), for example, 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 9 If the update in question allows for an automatic installation (an installation 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 two minutes of completion of the installation. 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 of whether the end user is present, if the machine is running and an automatic update is received, both the installation 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.

If the update is not automatic (and the end user has an Agent UI) a popup window gives the end user the ability to postpone the update.

Note To prevent any interaction by end users with the Agent installed on their system, you could use the Automatic Update feature in combination with the No Agent UI feature.

Step 10 Click the **Save** button.

You must generate rules to deploy software updates to Agents.

Summary

This topic summarizes the key points presented in this lesson.

Summary

Cisco.com

- **The use of groups will make administration of the CSA MC easier.**
- **CSA software is deployed with Agent kits, which include group membership and security policies.**
- **Access to Agent kits may be controlled by IP address.**
- **Host and group configuration information is available at the CSA MC.**
- **Software updates may be deployed with the CSA MC.**

© 2005 Cisco Systems, Inc. All rights reserved. HIPS v2.0—5-28

Lesson 6

Building Policies

Overview

The policies you create on the Cisco Security Agent Management Center (CSA MC) 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.

You will benefit by charting your security needs in advance rather than addressing problems as they are discovered. Because networks and network security are both dynamic entities, 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 lesson includes the following topics:

- Objectives
- Developing a Security Policy
- Policy Components
- Building Policies and Rule Modules
- Attaching rule Modules to Policies
- Summary

Objectives

Upon completing this lesson, you will be able to describe and configure policies and rule modules. This ability includes being able to meet the following objectives:

- Describe approaches to developing a security policy
- Discuss components of a policy
- Configure policies and rule modules
- Attach rule modules to policies

Developing a Security Policy

This topic introduces the development of a security policy.

Developing a Security Policy

Cisco.com

The following are alternative approaches to security policy design:

- **Permissive security model: Deny malicious actions and allow all other actions.**
- **Restrictive security model: Allow required actions and deny all other actions.**

© 2005 Cisco Systems, Inc. All rights reserved.HIPS v2.0-6.3

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

Before configuring your policies, you must understand exactly which network resources and services you want to protect and which threats you are most concerned about. The first step in planning a security policy is identifying the resources that your user community requires in order 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.

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

As you determine the network resources that are required by your user community, you can identify some of the threats posed to 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 e-mail 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, Post Office Protocol Version 3 (POP3), Internet Message Access Protocol (IMAP) for e-mail, and FTP.

You could take two 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 one hand, you might choose to enforce known good behaviors and selectively add targeted restrictions. This approach would be a more permissive security model. It 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 it may be suitable for highly secured environments. You could use both approaches, choosing the approach suited to different groups.

Developing a Security Policy (Cont.)

Cisco.com

- **Protect the application executables.**
- **Restrict the application processes.**
- **Protect application-specific data.**
- **Permit network access as required.**
- **Protect application registry keys.**

© 2005 Cisco Systems, Inc. All rights reserved.

HIPS v2.0--6-4

Once you understand how an application works, you can begin forming a policy to protect that application. You want to address five general areas for each resource that you are protecting. By addressing the security needs of these five areas, you can configure a well-formed policy to protect the resources that you are targeting.

As you build a policy to protect a designated resource, refer to the following steps, which will help you address each resource area:

- Protect the application executables.
- Restrict the application processes.
- Protect application-specific data.
- Permit network access as required.
- Protect application registry keys.

You must prevent writing to the application executables themselves to maintain the integrity of the executables. The only time that an executable should change is when you are upgrading the application. This type of rule, for example, would prevent a Trojan horse from naming itself “netscape.exe” to disguise itself as the Netscape executable.

Dictate what applications can and cannot do. You will likely want specific applications to write only to their own file types. To restrict an application, you must look at the files that the application needs to read and write to and then restrict the application to accessing those files only. 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 that 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, the CSA MC 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.

Restrict access to specified data by other applications. For server policies, you will want to protect information in certain directories on the server, allowing restricted access to specific files and blocking all outside access to other files. To correctly formulate this rule, you must examine which other applications (if any) need to access the application data. This type of rule would keep certain applications from retrieving sensitive data from a server, such as credit card information or a password file.

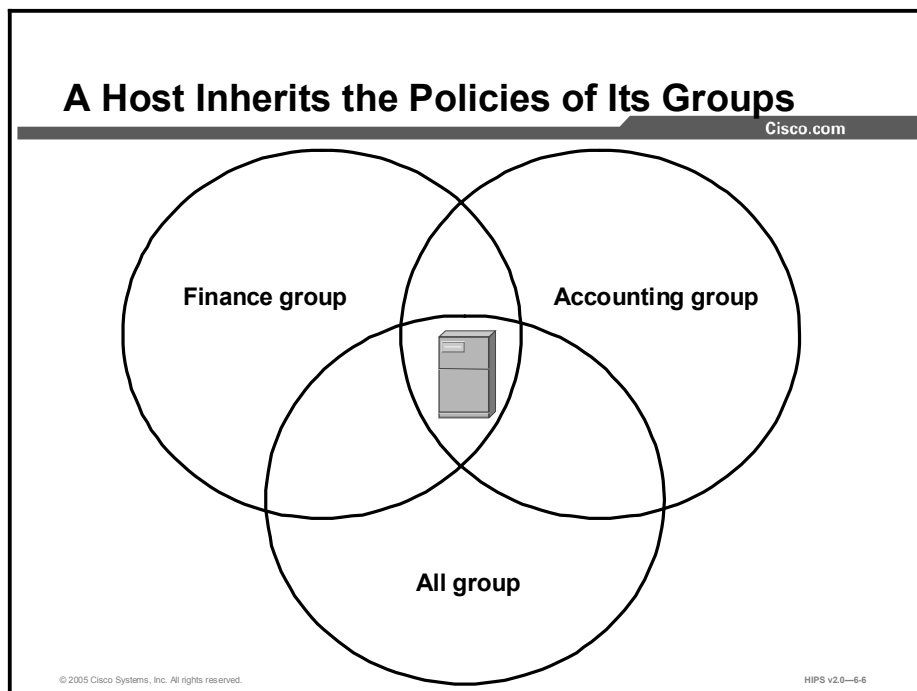
If an application requires network connectivity, you should specify which required network services must be enabled. Components that are “network visible” are especially vulnerable to attacks. Controlling what these network-accessible applications (and their spawned processes) can do is a critical part of a security policy.

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.

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

Policy Components

This topic describes policy components.



Cisco Security Agent (CSA) provides overall system protection by tying together the control of various system components while operating under the direction of assigned security policies.

You can attach multiple policies to a group. A host can belong to multiple groups, and the host then inherits policies from all of them. For example, a desktop can belong to the Finance group and inherit the Accounting group 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.

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

Rule Processing Order

Cisco.com

- **Priority 1 High Priority Terminate Process**
- **Priority 2 High Priority Deny**
- **Priority 3 Allow**
- **Priority 4 Query User (Default Terminate)**
- **Priority 5 Query User (Default Deny)**
- **Priority 6 Query User (Default Allow)**
- **Priority 7 Terminate Process**
- **Priority 8 Deny**
- **Priority 9 Default Action (Allow)**
- **Priority 10 Add process to application class**
- **Priority 11 Remove process from application class**
- **Priority 12 Monitor**

© 2005 Cisco Systems, Inc. All rights reserved.

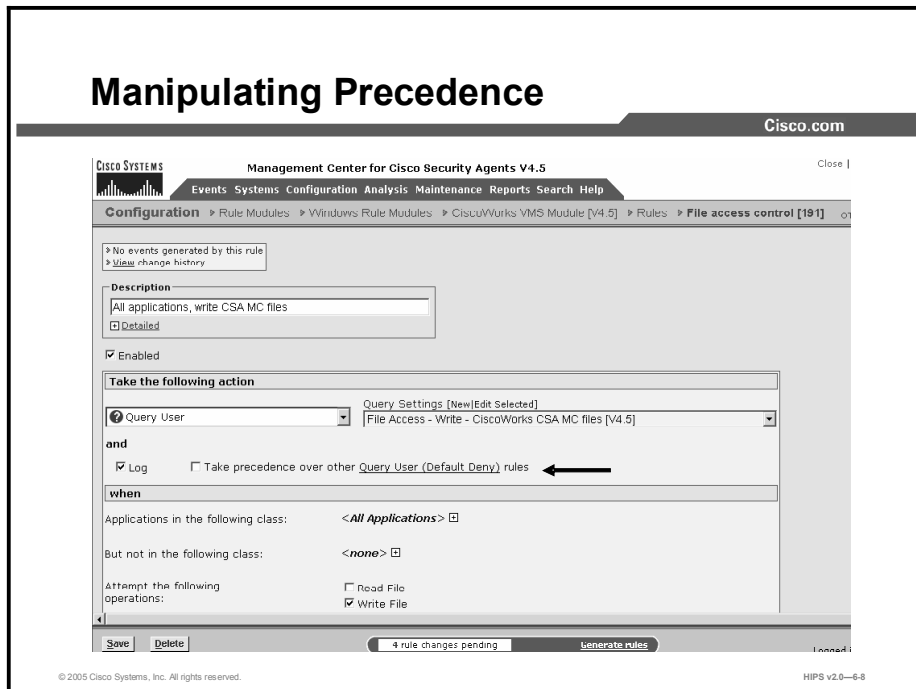
HIPS v2.0—6-7

When you configure certain rule types, you select an action for that rule (allow, deny, etc.). When you add your rule modules to policies, CSA MC orders individual rules from multiple modules according to action, in the following manner within each policy:

- Priority 1 High Priority Terminate Process
- Priority 2 High Priority Deny
- Priority 3 Allow
- Priority 4 Query User (Default Terminate)
- Priority 5 Query User (Default Deny)
- Priority 6 Query User (Default Allow)
- Priority 7 Terminate Process
- Priority 8 Deny
- Priority 9 Default Action (Allow)
- Priority 10 Add process to application class
- Priority 11 Remove process from application class
- Priority 12 Monitor

The priority listings beside each bulleted item indicate the manner in which CSA processes rules. All priority 1 enforcement rules (High Priority Terminate Process) are checked first and priority 8 enforcement rules (Deny) checked last, and that is only if no other higher-priority rules have already been triggered by a system action. Detection rules, such as priority 12 Monitor rules, are always checked, even in the presence of a higher-priority enforcement rule that governs the same resources that trigger first.

Manipulating Precedence



Most rule types provide a check box that allows you to manipulate how similar rule types are subordered within a policy. This check box, called Take Precedence Over Other <Action Type> Rules, is located in the rule configuration page. A rule with this check box selected is evaluated before similar rules that do not have this check box selected.

Manipulating precedence is discussed in detail in Lesson 7.

Making a Policy Mandatory

The screenshot shows the Cisco Management Center interface for configuring the CiscoWorks VMS policy. The 'Target Architectures' section is highlighted with a black arrow pointing to the 'Windows' checkbox, which is checked. The 'Attached Rule Modules' table lists several modules including 'CiscoWorks Application Classification Module', 'CiscoWorks Base Security Module', 'CiscoWorks CSA MC SQL Server module', and 'CiscoWorks Restrictive VMS Module'.

Name	Version	Description	Target
CiscoWorks Application Classification Module	4.5	Module classifying CiscoWorks applications	All v
CiscoWorks Base Security Module	4.5	Base security module for all systems running CiscoWorks	All v
CiscoWorks CSA MC SQL Server module	4.5	Module for SQL Server on the CSA MC system	All v
CiscoWorks Restrictive VMS Module	4.5	Module for systems running only the VMS bundle	All v

CSA MC provides three auto-enrollment architectural groups (Windows, Solaris, and Linux) that are mandatory for all hosts of a given operating system (OS) architecture. By providing group auto-enrollment for hosts, any policies that you attach to these groups also become mandatory by association. You might want to use these mandatory groups to apply policies that prevent some critical service from being inadvertently banned. For example, you could attach policies to prevent Domain Name System (DNS) or Dynamic Host Configuration Protocol (DHCP) from being disabled by an overly restrictive rule.

Building Policies and Rule Modules

This topic explains the configuration of CSA MC policies.

Building Policies and Rule Modules

Cisco.com

When configuring your own policies, configure items in the following manner:

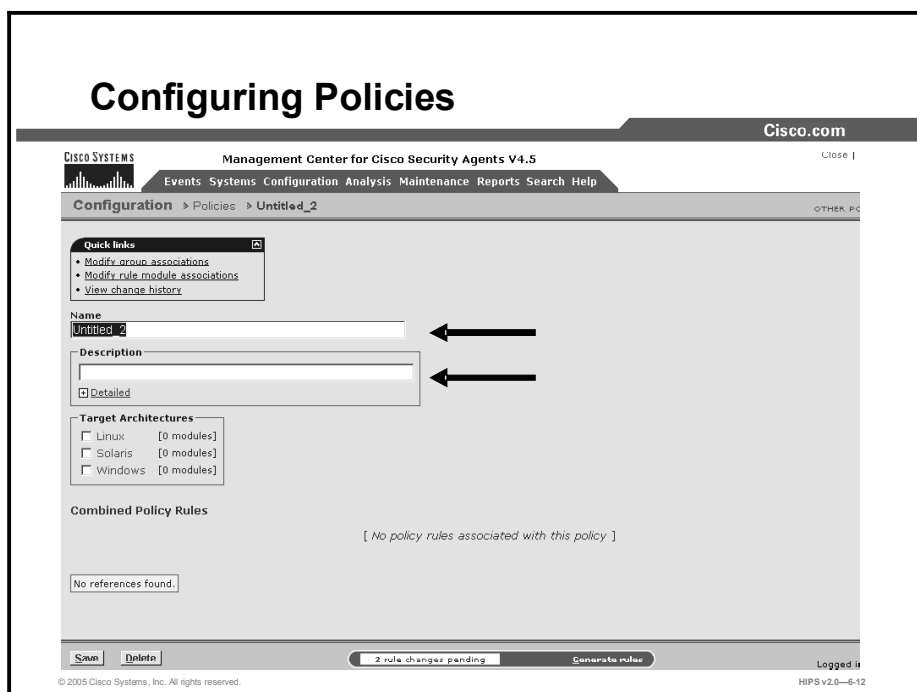
- **Decide what purpose the policy serves.**
- **Understand what tasks the rule modules that make up your policy must accomplish.**
- **Decide what rule types you must configure to accomplish the tasks that you have isolated.**

© 2005 Cisco Systems, Inc. All rights reserved. HIPS v2.0—6-11

When you configure a policy, you are combining multiple rule modules under a common name. That policy name is then attached to a group of hosts, and it uses the rules that make up the policy to control the actions that are allowed and denied on those hosts. You can have several different types of rules in a rule module and, consequently, within one policy.

The policy level is the common ground from which host groups acquire the rules that make up their security policy. If you are configuring your own policies, you should begin by understanding the purpose of your policy and how you must build your rule modules to meet your needs. It is recommended that you build your policies from the top down. In other words, configure items in the following manner:

- Decide what purpose the policy serves.
- Understand what tasks the rule modules that make up your policy must accomplish.
- Decide what rule types you must configure to accomplish the tasks that you have isolated.



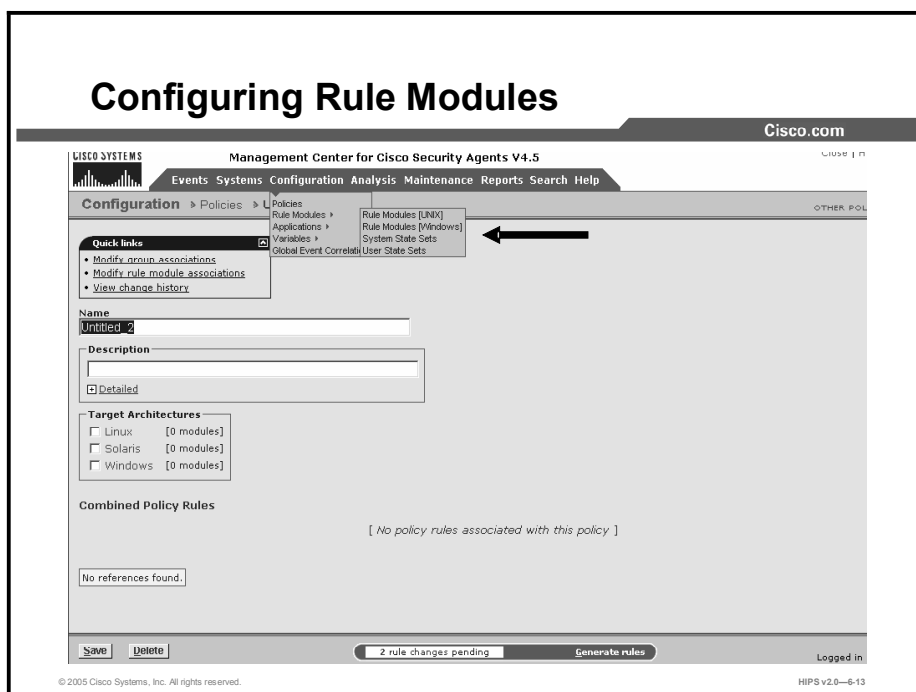
The policy level is the common ground by which host groups acquire the rules that make up their security policy. You can attach rule modules of differing architectures to the same policy. This way, you can configure task-specific, self-contained, inclusive policies across all supported architectures (Windows, Solaris, Linux) for software that is supported on all platforms.

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

To configure a policy, do the following.

- Step 1** Move the mouse over **Configure** in the menu bar of CSA MC and select **Policies** from the drop-down menu that appears. The policy list view appears.
- Step 2** Click the **New** button to create a new policy entry. This takes you to the policy configuration page.
- Step 3** In the available policy configuration fields, enter the following information:
 - **Name:** This is a unique name for this policy grouping of rule modules. Names are case insensitive, must start with an alphabetic character, can be up to 64 characters long, and can include alphanumeric characters, spaces, and underscores.
 - **Description:** This is an optional line of text that is displayed in the list view and helps you to identify this particular policy.
- Step 4** Click the **Save** button.

This policy is empty until you attach configured rule modules to it.



Rule modules are the building blocks of your policies. Modules are made up of several different types of rules.

Caution To maintain the integrity of the preconfigured policies and rule modules 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, create a new policy instead (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 rule module, do the following.

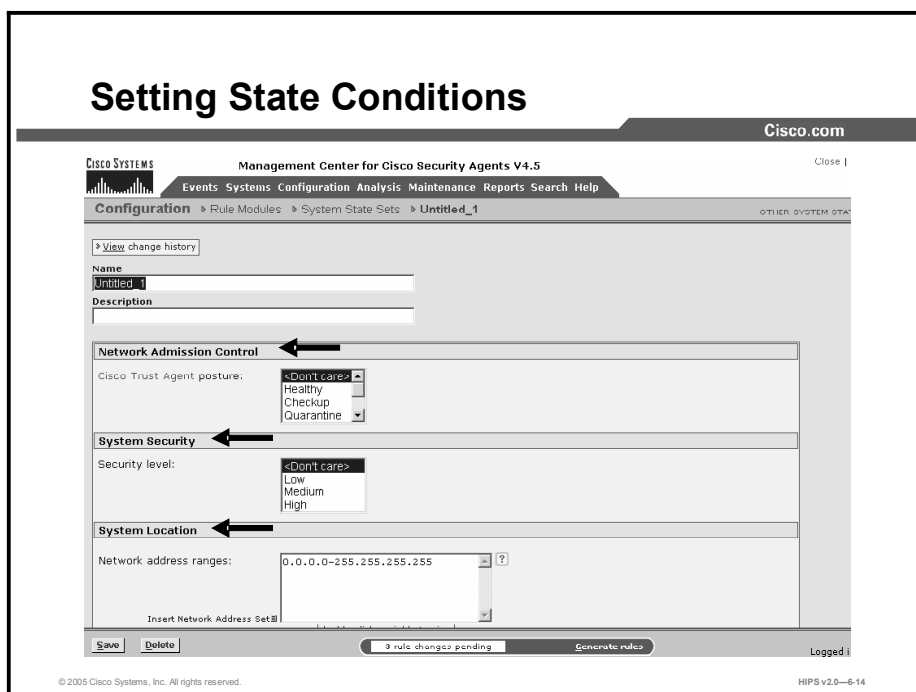
Step 1 Move the mouse over **Configuration > Rule Modules** in the menu bar. If you have not set OS admin preferences, you must select whether this is a Windows or a UNIX rule module from the cascading menu that appears. When you make a selection, the list of existing rule modules is displayed. CSA MC ships with several preconfigured modules.

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

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

Step 3 In the rule module configuration view, enter a unique name for your module. 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 policy list box when you are attaching modules to policies.

- Step 4** Enter a description of your module. This description is visible in the rule list view. Optionally, expand the Detailed field to enter a longer description.
- Step 5** Optionally, in the operating system box, you can select to target this module to a specific operating system within your Windows or UNIX classification.
- Step 6** Optionally, you can put this rule module into test mode. This way, you can have the rules within the test mode module operating in test mode while rules in other modules assigned to the same host are operating in a live mode. This is useful for testing new rule modules or changes to existing modules without having to turn off all protections for the hosts in question. You can also apply modes on the group level.
- Step 7** Click the **Save** button.
- Step 8** Now you add rules to your policy. Click the **Modify Rules** link at the top page.
- Step 9** Optionally, you can impose configured state conditions on rule modules. You can configure system state conditions or user state conditions.



System state and user state conditions let you write conditional rules based on the state of a system or the user of the system. Therefore, rules are only applied if the configured conditional settings are met.

System state parameters let you dictate conditions based on detected machine settings. When a machine is operating an Agent with a configured system state, the rules associated with that state apply only when the state parameters are met. They are conditional rules. Keep this in mind when assigning a state set to a rule module. For a more detailed description of system state setting options, read the configuration instructions here.

Configure a system state set by doing the following:

- Step 1** Move the mouse over **Configuration > Rule Modules** in the menu bar. Select **System State Sets** from the cascading menu that appears.
- Step 2** Click the **New** button to create a new system state.
- Step 3** Enter a unique name for your system state. You will select this name in the Rule Modules page. 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.
- Step 4** Enter a description.
- Step 5** In the Network Admission Control section, select one or more conditions. (Hold the Shift key down while you click the mouse to choose multiple, concurrent options. Hold the Control key down to select nonconcurrent options.) The Cisco Trust Agent posture state condition for a system ensures that corporate security requirements are met on that system. This feature works in conjunction with the capabilities of the Network Admission Control functionality.

Note Currently, the Cisco Trust Agent (an optionally installed product available from Cisco Systems) is only supported on Windows platforms.

The Cisco Trust Agent checks the status of a system and reports this status back to the Cisco Secure Access Control Server (ACS). Based on this status check, ACS returns a “posture state” that the Cisco Security Agent can act upon.

For example, if a machine is running antivirus software that is not up-to-date or is disabled, the Cisco Trust Agent can report this status to Cisco Secure ACS, which can then return an “Unknown” or a “Quarantine” state to the Cisco Security Agent. The Cisco Security Agent will then take action based on that posture state and enforce a stricter policy to protect that system or even quarantine the system from the network. Refer to your Cisco Secure ACS documentation for information on posture states and what they mean. Possible posture states are as follows:

- **<Don’t Care>**: This state is not provided by Cisco Secure ACS. All received posture states can match or not match this selection and the policy state is not affected. For UNIX states, this is currently the only valid posture state.
- **Healthy**: Host credentials are up-to-date, and the risk to the network from this host is low.
- **Checkup**: Host credentials are not quite up-to-date, but the risk to the network is low. The host should update credentials as soon as possible.
- **Quarantine**: Host credentials are out-of-date. The host is vulnerable to compromise and should be updated immediately. The risk to the network from this host is high.
- **Infected**: Host has been compromised. The risk to the network from this host is very high. The host should be cleaned immediately.
- **Unknown**: The posture of host cannot be determined due to an error.
- **Other**: This state is not provided by Cisco Secure ACS. If there is an incompatibility with posture state information received from ACS, it is seen as “Other” by the Cisco Security Agent. You can use this posture state as a criteria for enforcing a set of rules in the same manner that you use other criteria.

Step 6 In the System Security section, select one or more security level conditions. (Hold the Shift key down while you click the mouse to choose multiple, concurrent options. Hold the Control key down to select nonconcurrent options.) If the end user has an Agent UI, you can apply a security level condition, which allows the user to set the security sidebar on his or her UI to a specific level. This provides some degree of control to the user for managing false positives or controlling security when operating remotely or on the local network.

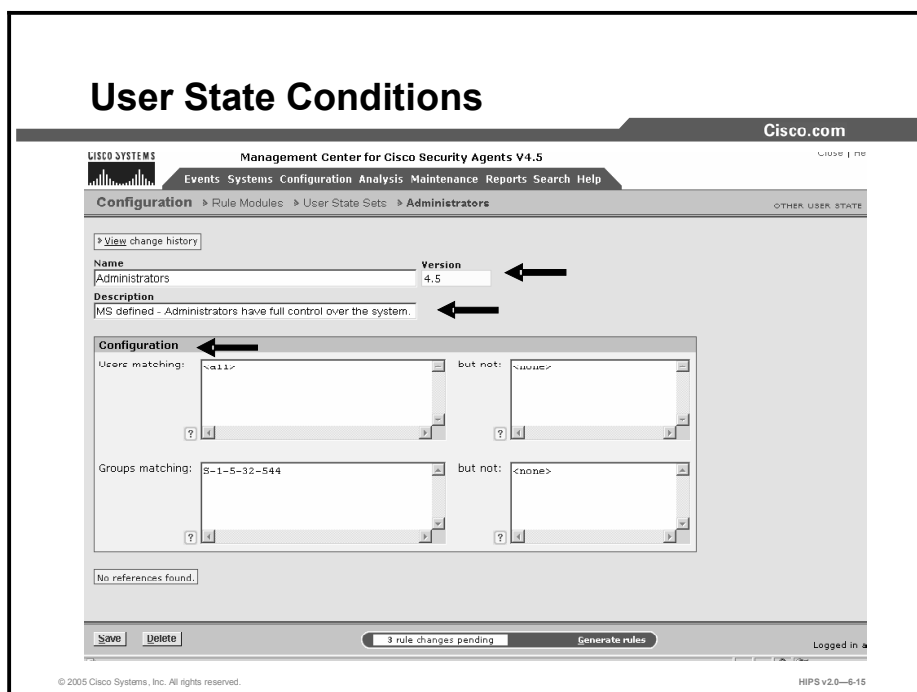
Step 7 In the System Location section, enter one or more addresses or address ranges in the Network Address Ranges field to create a state condition based on system address. By default, no restrictions are set here. If you enter address conditions here, the conditions apply if at least one interface matches what is specified. If you enter multiple ranges, only one address has to match for the system state to apply.

Step 8 In the System Location section, use the DNS Suffix Matching field to set a condition based on DNS suffix (the suffix of the DNS server that is used to resolve names). If any DNS server suffix (for example, cisco.com) matches an item specified here, the condition is applied. You can use the But Not field to make specific exclusions to DNS suffix-matching parameters that you configure.

- Step 9** In the Additional State Conditions section, click the **Add State** link to add one or more of the following additional states to this page. (Select an option from the drop-down menus that appear. Then use the drop-down menu to the right of your selected option to choose one of the following settings: <Don't Care>, Yes, No.)
- Select the Management Center Reachable option to set a state condition based on whether the Cisco Security Agent can communicate with the Management Center. Based on this condition, rules are applied or not applied.
 - Select the Installation Process Detected option to set a state condition to apply if an installation is in progress on a system. For example, perhaps you want to apply a less restrictive set of rules to allow an installation when it is detected on a system.
 - Select the Rootkit Detected option to set a state condition if a driver is seen attempting to dynamically load. Based on this condition, rules are applied or not applied.
 - Select the System Booting option to set a state condition to apply for the time frame in which the system is booting. Based on this condition, a set of designated rules apply only during boot time.
 - Select the Virus Detected option to set a state condition to apply if a virus is detected on a system. Based on that virus detection, a state condition setting can enforce a designated set of rules.

Step 10 Click the **Save** button.

Note The system states that you configure are additive. All specified state conditions are used as part of the requirement(s) to be met for the state to trigger.



User state parameters let you dictate conditions based on detected user and/or group settings. When a machine is operating an Agent with a configured user state, the rules associated with that state apply only when the state parameters are met. They are conditional rules. Keep this mind when assigning a user set to a rule module.

You should also keep in mind that the process of checking user states is an expensive one for the system. You should use these settings judiciously.

An example of when you might want to employ a user state is as a restriction dictating who can alter web server pages. The web server application itself should only serve pages, not edit them. You could use a setting here to ensure that only authenticated administrators using a specific application (such as FrontPage) are allowed to alter web server content.

Another example of appropriate user state setting usage is a situation in which groups of users are restricted from performing certain tasks that you only want to allow administrators to perform, such as suspending Agent security.

Configure a user state set by doing the following:

- Step 1** Move the mouse over **Configuration > Rule Modules** in the menu bar. Select **User State Sets** from the cascading menu that appears.
- Step 2** Click the **New** button to create a new user state.
- Step 3** Enter a unique name for your user state. You will select this name in the Rule Modules page. 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.
- Step 4** Enter a description.

- Step 5** In the Users Matching field, if you choose to set a condition based on user information, enter the user string data, using machine name or domain name and user account. For example, entries in this field might appear as follows:

```
Domain\jsmith  
jsmith\Administrator  
\Administrator  
Domain\*
```

You can use wildcards in the Users Matching and But Not fields.

- Step 6** Use the But Not field to make specific exclusions to user matching parameters that you configure.

- Step 7** In the Groups Matching field, if you choose to set a condition based on group information, an entry in this field might appear as follows:

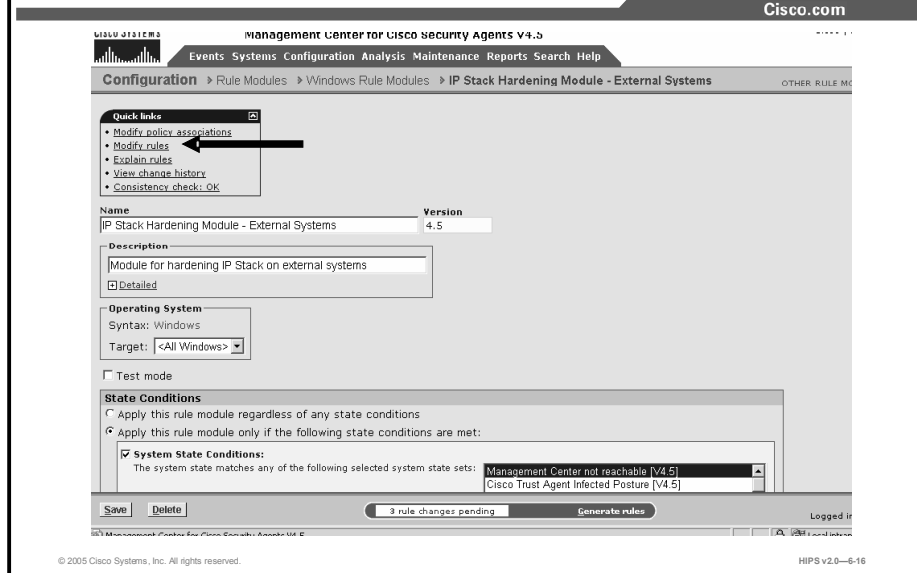
```
Administrator
```

You can also enter Security Identifier (SID) numerical classifications in the Groups Matching field. Using a SID rather than a group name is useful when writing states that will apply across international versions of operating systems. Group names may be different across languages, but a SID classification is always the same. You cannot use wildcards in the Groups Matching and But Not fields. If users belong to multiple groups, they only need to match one named group to meet the criteria of the user state.

Note It is recommended that you use group permissions rather than user permissions because group designations are more widely applicable.

- Step 8** Use the But Not field to make specific exclusions to group matching parameters that you configure.
- Step 9** Click the **Save** button.

Rule Module, Modify Rules



To add rules to a rule module, click the **Modify Rules** link at the top of the Rule Modules 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.

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

Rules List

The screenshot shows the Cisco Management Center for Cisco Security Agents V4.5 interface. The breadcrumb trail is: Configuration > Rule Modules > Windows Rule Modules > IP Stack Hardening Module - Internal Systems [V4.5] > Rules. The table below shows the following rules:

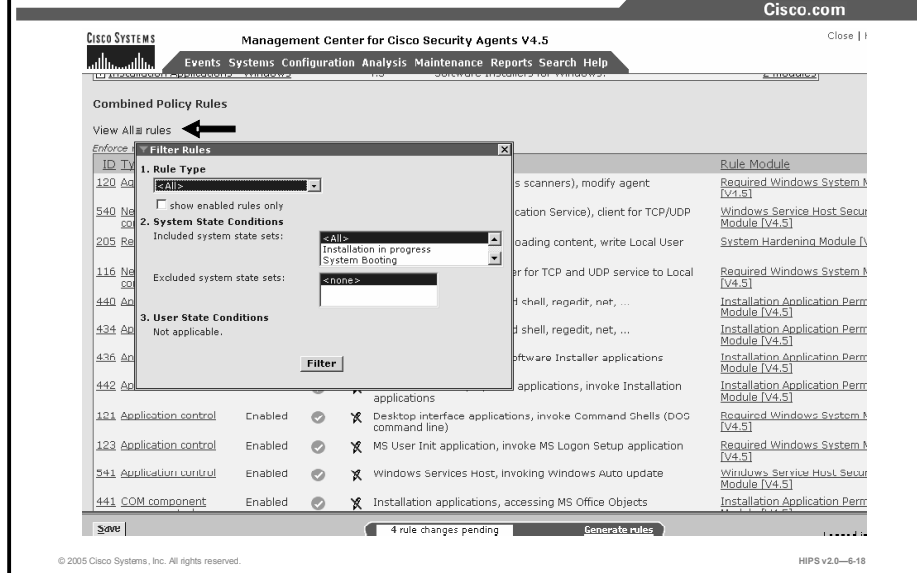
ID	Type	Events	Status	Action	Log	Description
453	Network shield		Enabled	✔	✗	Port scans, from Authorized port scanners
451	Network shield		Enabled	✗	✔	Port scans, from all hosts
452	Network shield		Enabled	✔	✔	TCP SYN flood, from all hosts
454	Network shield		Enabled	✔	✔	Ping scans, from external hosts
450	Network shield		Enabled	⚙	-	ICMP configuration and information messages, from all hosts

Below the table, there is an 'Add rule #' section with a 'Copy' button and a dropdown menu for the rule module, currently set to 'IP Stack Hardening Module - Internal Systems [V4.5]'. At the bottom of the interface, there are buttons for 'Delete', 'Enable', and 'Disable', a status indicator '4 rule changes pending', and a 'Generate rules' button. The footer includes '© 2005 Cisco Systems, Inc. All rights reserved.' and 'HIPS v2.0-6-17'.

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 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



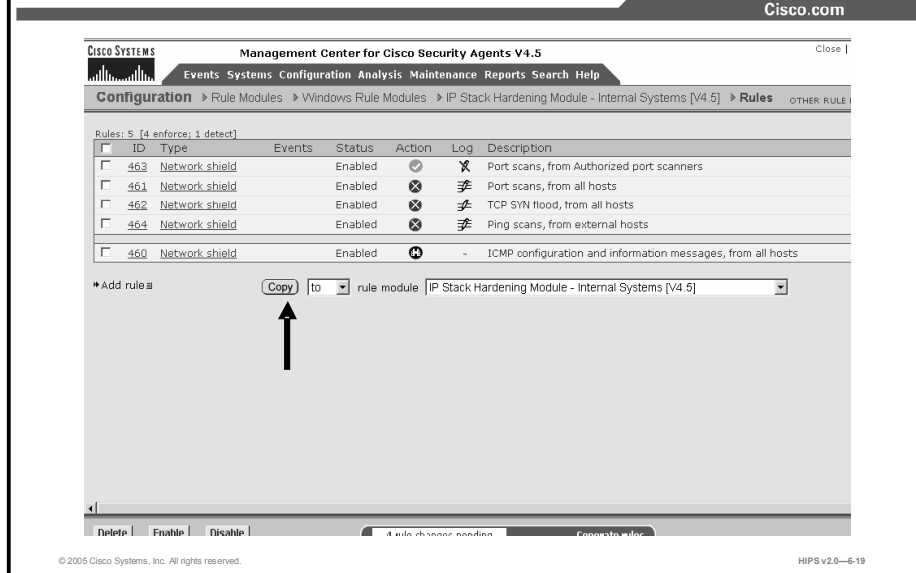
The Groups configuration page, Policy configuration page, and Rules configuration page display a table that lists either the rules attached to the group or the rules included in the module. On all three pages, there is a View All Rules item above the table. Clicking the All link here lets you filter your view of this rule list by selected rule type. When you click All, a popup appears listing the rule types present in the module or modules. Select a rule type from the popup, and that is now the only rule type displayed in the table. You can also view only enabled rules by selecting the Show Enabled Rules Only check box and then selecting the rule type you wish to view.

Note When you filter the rules display, other rules are not removed from the module; only your view of the module changes. You can revert to the entire summary view by selecting All from the same popup menu.

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

If you have user or system states applied to rule modules, you can also filter the display based on those settings. This is useful for viewing which rules are applied when particular states are active.

Copy Rules Between Modules



Use the Copy button in conjunction with the drop-down lists at the bottom of the Rule Modules page to copy selected rules to another rule module that you designate. Copying rules across modules is similar to cloning configurations. (You can also clone rules within policies using the Copy button, which is described in this section.)

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

- Step 1** From the Rule Modules page, select the check box for the rule or rules that you want to copy to another module.
- Step 2** Beside the Copy button, “to” is the default selection in the drop-down menu. (Do not change this for copying individual rules between modules.) From the Rule Module drop-down list, select the name of the module 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 module.

To clone rules within a module, repeat step 1 above. Then, rather than selecting another module in the Rule Module drop-down list, select the current module that you are in. Selected rules are cloned within the same module when you click the Copy button.

Select **from** in the drop-down menu beside the Copy button to copy all of the rules from the selected module (in the Rule Module drop-down list) to the current module.

Compare Rule Modules

The screenshot displays the 'Compare Rule Modules' utility in the Cisco Management Center. The main window is titled 'Compare Microsoft IIS Web Server [W, V4.5] and Microsoft SQL Server 2000 [W, V4.5]'. It compares two rule modules side-by-side. The left module is 'Microsoft IIS Web Server' (Version 4.5) and the right is 'Microsoft SQL Server 2000' (Version 4.5). The comparison table shows various attributes like 'Operating System' (All OS types) and 'Test mode' (No). Below the table, a detailed description of a rule is shown, including its 'Action' (Allow) and 'Application Classes' (IIS Web Server in Isolation Mode [V4.5]). The interface includes navigation buttons like 'Copy', 'Delete', and 'Generate Policy'.

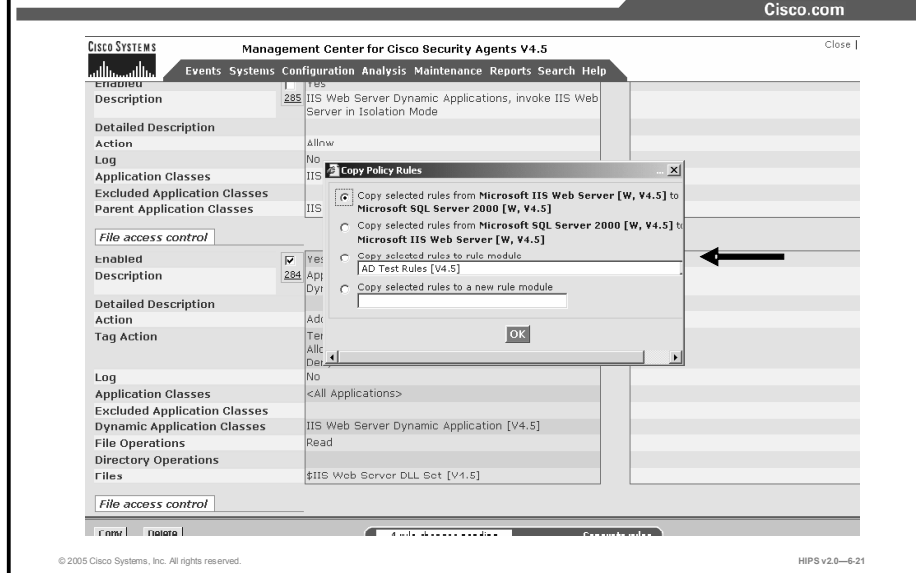
When you select the check box next to two items (you cannot compare more than two configurations at a time) and click the Compare button, CSA MC displays the configurations side by side and highlights the differences in red. Once you have examined how the configurations compare, you can select to merge specific rules, to copy rules to another module, or to copy rules to a new module. (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 have imported configurations or upgraded CSA MC. These processes can cause you to have duplicate or very similar configuration items. Comparing and merging configurations makes it easier to consolidate duplicate items. This Compare utility is also available for application classes and variables.

Here are some additional notes about the Compare utility:

- When you compare rule modules, the similar rules within those modules 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 module and no corresponding similar rule in the second module, there is nothing displayed beside that rule in the comparison.
- If you have rules in your modules 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 or deny actions. Logging and allow or 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.

Copy Rule Module Popup Box



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

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

View Change History

The screenshot displays the Cisco Management Center for Cisco Security Agents V4.5 interface. The page title is "View Change History". The breadcrumb navigation shows: Configuration > Rule Modules > Windows Rule Modules > IP Stack Hardening Module - External Systems. A "Quick links" menu on the left contains the following items: Modify policy associations, Modify rules, Explain rules, View change history (highlighted with an arrow), and Consistency check: OK. The main configuration area shows the "Name" as "IP Stack Hardening Module - External Systems" and the "Version" as "4.5". The "Description" is "Module for hardening IP Stack on external systems". The "Operating System" section shows "Syntax: Windows" and "Target: <All Windows>". The "State Conditions" section is set to "Apply this rule module only if the following state conditions are met:", with "System State Conditions:" checked. The selected system state sets are "Management Center not reachable [V4.5]" and "Cisco Trust Agent Infected Posture [V4.5]". At the bottom, there are "Save" and "Delete" buttons, a status bar indicating "3 rule changes pending", and a "Generate rules" button. The footer includes "© 2005 Cisco Systems, Inc. All rights reserved." and "HIPS v2.0-6-22".

At the top of each rule page, there is a View Change History link. Click this link to go to a page that lists all the changes that have been made to this rule. This View Change History link is also available for application classes, variables, rule modules, and policies.

Rule Explanation Page

Cisco.com

CISCO SYSTEMS
Management Center for Cisco Security Agents V4.5
Close |

Events Systems Configuration Analysis Maintenance Reports Search Help

Configuration
Rule Modules
Windows Rule Modules
IP Stack Hardening Module - Internal Systems [V4.5]
Explanation
OTHER

Explanation of rule module IP Stack Hardening Module - Internal Systems [V4.5]

! The following rules are applied only if the following conditions are met:

- the system state matches system state set `Management_Center_reachable [V4.5]`.

Network_shield

IP and Transport security checks

Irrespective of any other rules,
 Attempts to connect to any server and accept connections from any client whose address is contained in address ranges 0.0.0.0-255.255.255.255 using local addresses contained in address ranges 0.0.0.0-255.255.255.255 when detecting

- ICMP configuration message
- ICMP information message

will be monitored. An event will be logged when the rule is triggered.

460

In the absence of any applicable 'high priority deny' or 'high priority terminate process' rules,
 Attempts to connect to any server and accept connections from any client whose address is contained in address sets `Authorized_Port_C [V4.5]` using local addresses contained in address ranges 0.0.0.0-255.255.255.255 when detecting

- TCP/UDP port scan

will be allowed. No events will be logged when the rule is triggered.

463

In the absence of any applicable 'allow' or 'query' rules,
 Attempts to connect to any server and accept connections from any client whose address is contained in address ranges 0.0.0.0-255.255.255.255 using local addresses contained in address ranges 0.0.0.0-255.255.255.255 when detecting

- TCP/UDP port scan

will be denied. An event will be logged when the rule is triggered.

461

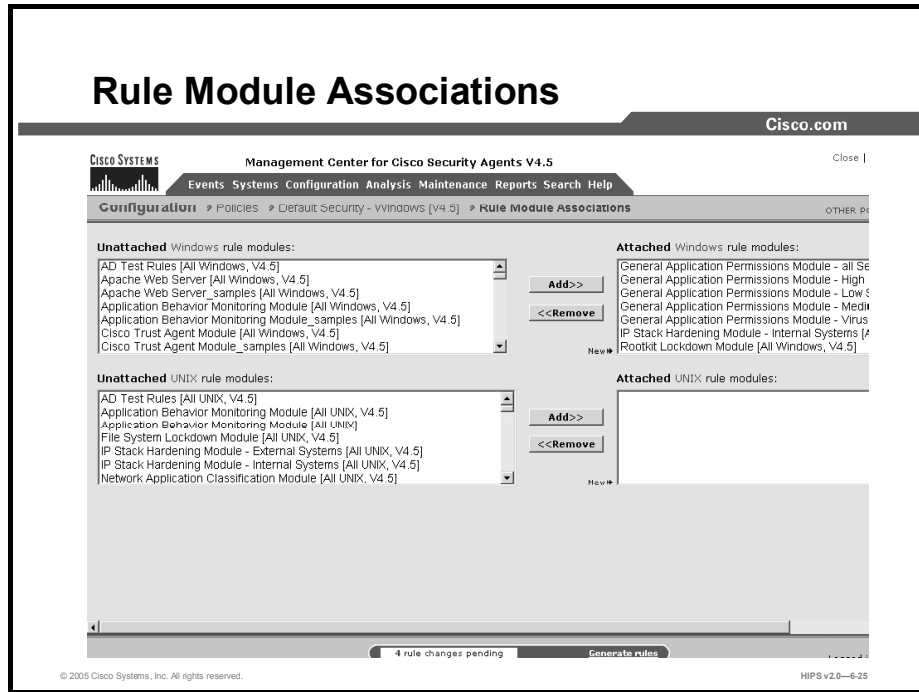
4 rule changes pending
Generate rules

© 2005 Cisco Systems, Inc. All rights reserved.
HIPS v2.0—6-23

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, Rule Modules, or Policy page takes you to this paragraph explanation.

Attaching Rule Modules to Policies

This topic explains how to attach rule modules to policies.



When you configure a rule module, you are combining access control rules and/or system correlation rules, and monitoring rules under a common name. That rule module name is then attached to a policy. That policy uses the rules that make up the module to control the actions that are allowed and denied on hosts.

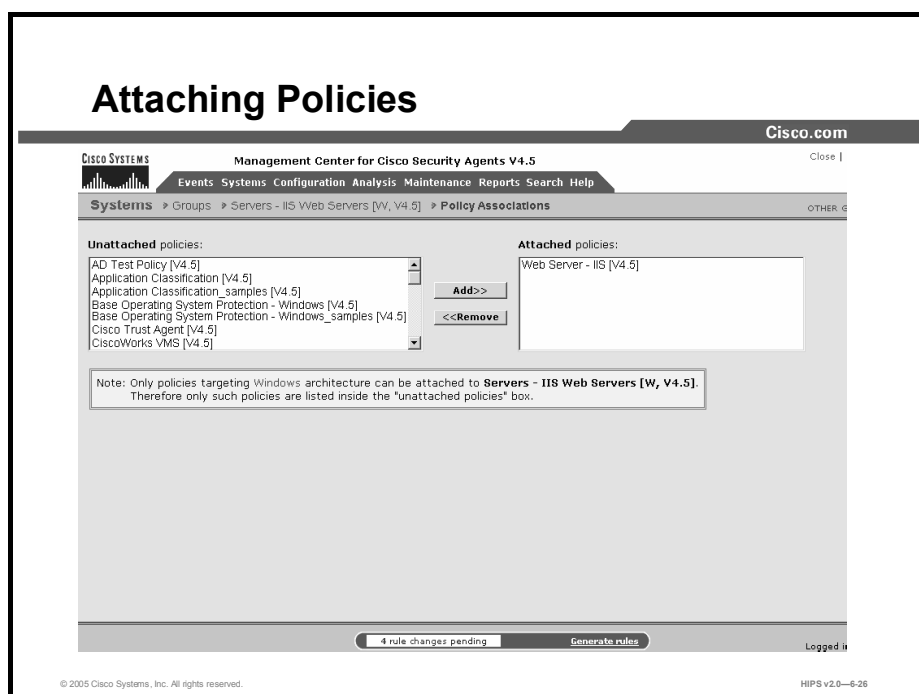
With CSA MC, you can attach a rule module to a policy by using the Modify Policy Associations link in the Rule Modules configuration page, or you can attach a policy to a rule module by using the Modify Rule Module Associations link in the Policy list view page.

To attach a rule module or rule modules to an existing policy using the Modify Policy Associations link in the rule module configuration page, do the following:

- Step 1** Attach a rule module to a particular policy by accessing that rule module's edit view. From Configuration in the menu bar, click on **Rule Modules** for the OS type that you want to access the list view for those modules.
- Step 2** From the rule module list view, click the link for the rule module that you want to attach to a policy. This brings you to that rule module'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 box on the left contains the policies that the rule module is not attached to. The box on the right contains policies that the rule module is attached to.
- Step 4** To add this rule module to an existing policy, select the rule module in the box on the left and click the **Add** button. The selected rule module moves to the box on the right and is now attached to the policy.

Note You can attach rule modules of differing architectures to the same policy. This way, you can configure task-specific, self-contained, inclusive policies across all supported architectures for software that is supported on all platforms. For example, Apache is a web server software product that supports Windows, Linux, and Solaris platforms. You can attach three OS-specific rule modules for Apache to one policy, and you only need to maintain that one Apache policy.

Caution In order to deploy rule modules to hosts, you must remember to attach the policy that the rule module is associated with to a group.



When you configure a policy, you are combining configured rule modules under a common name. That policy name is then attached to a group of hosts, and it uses the rules that make up the policy to control the actions that are allowed and denied on those hosts.

With CSA MC, you can attach a policy to a group by using the Modify Policy Associations link in the Group configuration page, or you can attach a group to a policy by 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 Associations link to attach one policy to multiple groups.)

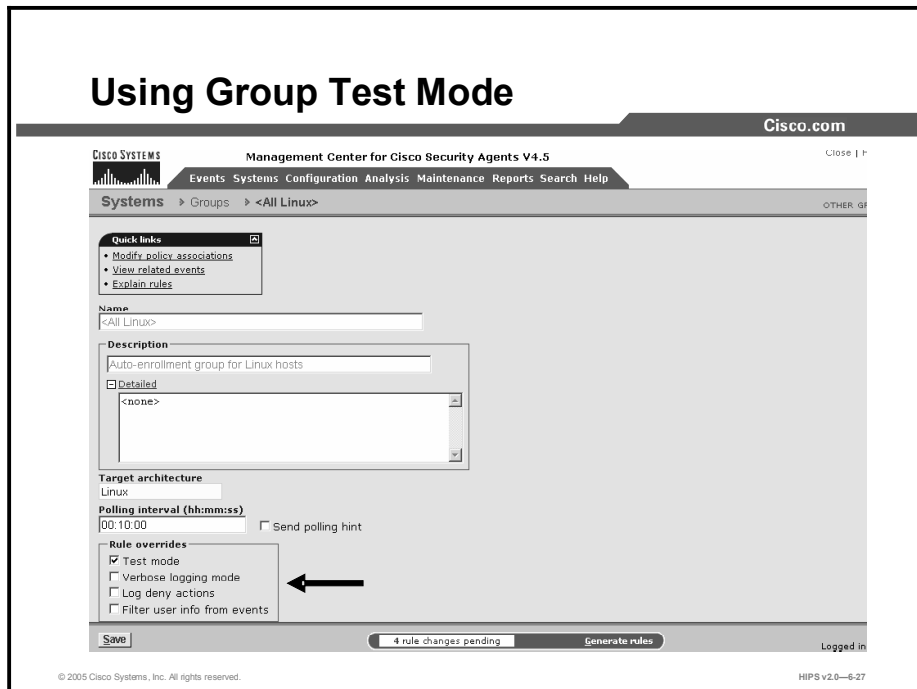
To attach a policy or policies to an existing group using the Modify Policy Associations link in the Group configuration page, do the following.

- 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 that 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 box on the left contains the policies that are not attached to this group. The box on the right contains policies that are attached to this group.
- Step 4** To add an existing policy to this group, select the policy in the box on the left and click the **Add** button. The selected policy moves to the box on the right and is now attached to the group.

Note To remove a policy from a group, select the policy in the box on the right and click the **Remove** button. It moves back to the box on the left. (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.

Note You can try out policies on host systems by selecting Test Mode for a group or for a particular rule module. If you select Test Mode and enable logging on rules attached to "test mode" groups, the Agent will log designated denied events that are triggered by policies but will not take any actions on those events.

Using Group Test Mode



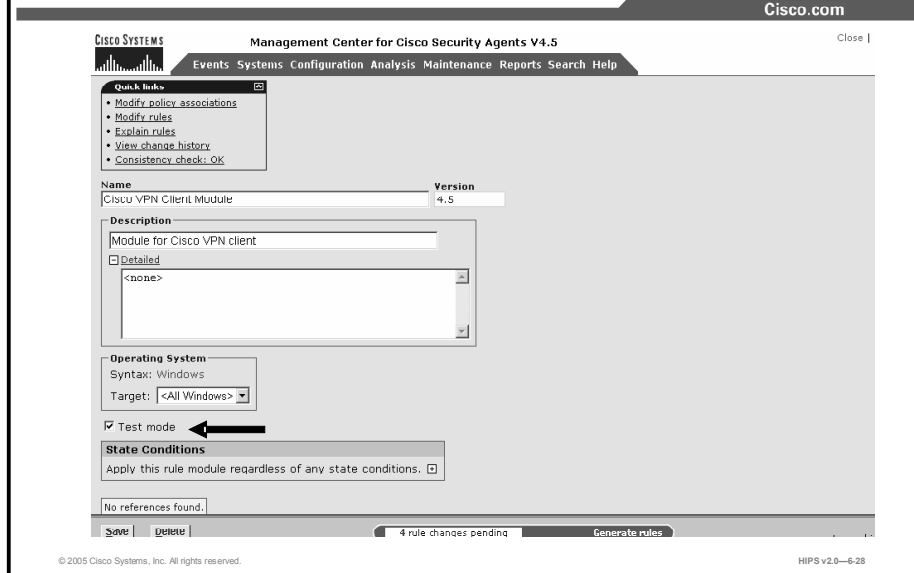
Test mode is useful when you are installing a new host or modifying a host configuration and you want to understand the ramifications without actually impacting host operation. When operating in test mode, the Agent will not deny any action or operation even if an associated policy says it should be 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 messages 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 of whether test mode is on or off.

In group test mode, you can turn on test mode in two places within the CSA MC. If it is enabled on the group level, all rules on hosts within test mode groups are in test mode. If a host belongs to a group with test mode selected, all policies associated with that host (not just the policies applied to the test group) are in test mode. (This is the case even if the host is part of another group that does not have test mode selected.) Therefore, test mode applies to the host as a whole, not to specific policies.

Using Rule Module Test Mode



You can also use test mode on the rule module level. This way, you can have the rules within the test mode module operating in test mode while rules from other modules assigned to the same host are operating in a live mode. This is useful for testing new rule modules or changes to existing modules without having to turn off all protections for the hosts in question.

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.

Generate Configuration

Warning :
The following policies are not attached to any hosts or groups:

- Instant_Messenger_samples
- Network_Ouarantine
- Samba_Server - Linux
- Samba_Server - Linux_samples
- Untitled_1
- Untitled_2
- Virus_Scanner - McAfee
- Virus_Scanner - McAfee_samples
- Virus_Scanner - Norton
- Virus_Scanner - Norton_samples
- Virus_Scanner - Trend
- Virus_Scanner - Trend_samples

Debug (will be removed)

4 changes since the last rule program generation:

Action	Time	Administrator
Create rule module 'Untitled_1'	12/10/2004 1:22:12 AM	admin
Create system state set variable 'Untitled_1'	12/10/2004 12:22:06 AM	admin
Create policy 'Untitled_2'	12/9/2004 11:31:24 PM	admin
Create policy 'Untitled_1'	12/9/2004 11:21:56 PM	admin

© 2005 Cisco Systems, Inc. All rights reserved. HIPS v2.0—6-29

The Generate Rule Programs view displays the status of all nondistributed 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 have 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.

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.)

Note Before you generate rule programs and distribute them to Agents, you can view all database changes, including the time that the changes were made and the administrator who made them by accessing the Audit Trail view from the Reports drop-down list.

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.

Summary

This topic summarizes the information that you learned in this lesson.

Summary

Cisco.com

- A well-balanced security policy must weigh business needs against security concerns.
- Rules can be configured with great flexibility and granularity.
- A host inherits the policies of all groups in which it has membership.
- When you add rule modules to policies, CSA MC orders individual rules from multiple modules according to action.
- CSA MC provides auto-enrollment for Windows, Solaris, and Linux systems.
- When you configure a policy, you are combining multiple rule modules under a common name.
- When configuring a rule module, you are combining access control rules and/or system correlation rules, and monitoring rules under a common name.
- Rule and policy changes do not take effect until the Generate Rules utility is run.

© 2005 Cisco Systems, Inc. All rights reserved.

HIPS v2.0—6-30

Lesson 7

Rule Basics

Overview

Rules are the foundation of your security policies. Cisco Security Agent Management Center (CSA MC) lets you create several rule types. Each rule type requires you to enter varying combinations of information using a specific syntax. This lesson explains the basics of how CSA MC rules work and shows how to configure a wide range of rules for Windows and UNIX systems.

The following topics are covered in this lesson:

- Overview
- Basics of Rule Construction and Functionality
- Rules Common to Windows and UNIX
- Windows-Only Rules
- UNIX-Only Rules
- Summary

Objectives

Upon completing this lesson, you will be able to configure rules for Windows and UNIX systems. This includes being able to meet the following objectives:

- Understand the basics of rule construction and functionality
- Configure rules common to Windows and UNIX systems
- Configure Windows-Only rules
- Configure UNIX-Only rules

Basics of Rule Construction and Functionality

This topic explains the basics of how CSA MC rules work.

Rule Basics

Cisco.com

- **File access control rules: Allow or deny based upon the following:**
 - The action you are allowing or denying
 - The application attempting to access the file
 - The operation (read, write) attempting to act on the file
- **Network access rules: Control access based upon the following:**
 - The action you are allowing or denying
 - The application attempting access
 - The direction (client, server) of the communication
 - The service that a system is attempting to use
 - The address that a system is attempting to communicate with

© 2005 Cisco Systems, Inc. All rights reserved.HIPS v2.0--7.3

Rules are the foundation of your security policies. Creation of each rule type requires you to enter information specifying the desired behavior.

- Use file access control rules to allow or deny the operations (read, write) that the selected applications can perform on files, depending on the following:
 - 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 rules to control access to specified network services according to the following:
 - 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 that a system is attempting to use
 - The address that a system is attempting to communicate with

Rule Basics (Cont.)

Cisco.com

- **Registry access control rules: Allow or deny according to the following:**
 - The action that you are allowing or denying
 - The application that is attempting to write to the registry keys and values
- **COM component rules: Allow or deny based upon the following:**
 - The action that you are allowing or denying
 - The application that is accessing the COM component

© 2005 Cisco Systems, Inc. All rights reserved.

HIPS v2.0—7-4

- Use registry access control rules (Windows only) to allow or deny writing to specified registry keys by selected applications according to the following:
 - The action that you are allowing or denying
 - The application that is attempting to write to the registry keys and values
- Use Component Object Model (COM) component access control rules (Windows only) to allow or deny access to specified COM components by selected applications according to the following:
 - The action that you are allowing or denying
 - The application that is accessing the COM component

Other types of policies shipped with the CSA MC provide event correlation and heuristic features that can be enabled on a per-group basis, such as port scan detection, SYN flood protection, the prevention of predictable TCP sequence numbers, and the blocking of malformed IP packets. These features are especially useful for network servers.

The following basic detection rules work as follows:

- Use various tagging rule types with “Add process to application class” or “Remove process from application class” selected to build application classes based on process behavior rather than executable name. Once applications are built, or “tagged,” they are used in other enforcement rules.
- Use rules such as NT Event Log and sniffer and protocol detection to log designated event types when they occur.
- By tying together the controlling and monitoring of various system functions and by operating under the direction of assigned policy rules, Agents provide overall system protection.

Rule Processing Order

Cisco.com

- **Priority 1 High Priority Terminate Process**
- **Priority 2 High Priority Deny**
- **Priority 3 Allow**
- **Priority 4 Query User (Default Terminate)**
- **Priority 5 Query User (Default Deny)**
- **Priority 6 Query User (Default Allow)**
- **Priority 7 Terminate Process**
- **Priority 8 Deny**
- **Priority 9 Default Action (Allow)**
- **Priority 10 Add process to application class**
- **Priority 11 Remove process from application class**
- **Priority 12 Monitor**

© 2005 Cisco Systems, Inc. All rights reserved.

HIPS v2.0--7.5

When you configure certain rule types, you select an action for that rule (allow, deny, and so on). When you add your rule modules to policies, CSA MC orders individual rules from multiple modules according to action, in the following manner within each policy:

- Priority 1 High Priority Terminate Process
- Priority 2 High Priority Deny
- Priority 3 Allow
- Priority 4 Query User (Default Terminate)
- Priority 5 Query User (Default Deny)
- Priority 6 Query User (Default Allow)
- Priority 7 Terminate Process
- Priority 8 Deny
- Priority 9 Default Action (Allow)
- Priority 10 Add process to application class
- Priority 11 Remove process from application class
- Priority 12 Monitor

The priority listings beside each bulleted item indicate the manner in which CSA processes rules. All priority 1 enforcement rules (High Priority Terminate Process) are checked first and priority 8 enforcement rules (Deny) checked last, and that is only if no other higher-priority rules have already been triggered by a system action. Detection rules, such as priority 12 Monitor rules, are always checked, even if a higher-priority enforcement rule that governs the same resources triggers first.

When you configure your access control rules, you must select an action for each rule. The following list describes all possible action types. Note that not all action types are available for all rules. Enforcement actions are as follows:

- **High Priority Terminate Process:** Select this action type to create a terminate rule that takes precedence over all other allow, terminate, deny, and query rules. This action denies the application access to the resource in question and also attempts to terminate the application process. Under the same circumstances, if the terminate is not high priority and a conflicting allow rule exists in the same policy, the allow takes precedence.
- **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 that you specify to run. Because the default action of all policies is allow, generally, you will only want to configure allow rules as exceptions to existing deny rules within policies.
- **Query User (Default Terminate):** Select this action type to prompt the user when the action that you indicate occurs. The user then has 5 minutes to answer Yes, No, or Terminate. By default, the action will be denied and the process will be terminated unless the user decides otherwise. (Query User options are not available for Solaris rules.)
- **Query User (Default Deny):** Select this action type to prompt the user when the action that you indicate occurs. The user then has 5 minutes to answer Yes, No, or Terminate. By default, the action will be denied unless the user decides otherwise. (Query User options are not available for Solaris rules.)
- **Query User (Default Allow):** Select this action type to prompt the user when the action that you indicate occurs. The user then has 5 minutes to answer Yes, No, or Terminate. By default, the action will be allowed unless the user decides otherwise. (Query User options are not available for Solaris rules.)

Note **Text used to query user:** If you are configuring a Query User rule, you must also configure query settings. The text you type into the query settings field is the same text that will appear in the query user popup box to explain what is occurring on the system to the user.

- **Terminate Process:** This action denies the application access to the resource in question and also attempts to terminate the application process.

Note All processes cannot be terminated safely (for example, winlogon). If it is not safe to terminate the process, the action will be denied but not terminated.

- **Deny:** Select this action type to create a rule that stops the specified application from running on systems. (When you select Deny for this rule, if users attempt to run the application in question, they are notified with a popup box explaining that the application is forbidden to run.)

Rule Processing Order (Cont.)

Cisco.com

Detection actions are as follows:

- **Add process to application class**
- **Remove process from application class**
- **Monitor**

© 2005 Cisco Systems, Inc. All rights reserved.

HIPS v2.0—7-6

Detection actions are as follows:

- **Add process to application class:** Use this for defining dynamic application classes. A dynamic application class is built based on an application's behavior rather than a specific application executable name. A process will be added to a dynamic class when the parameters of the rule (allow, deny, terminate) that dictate access to a resource are met.
- **Remove process from application class:** Use this action type to remove a dynamic application tag from a process. A process will be removed from a dynamic class when the parameters of the rule (allow, deny, terminate) that dictate access to a resource are met.

Note Dynamic classifications are part of an 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 reevaluated and classified on the system.

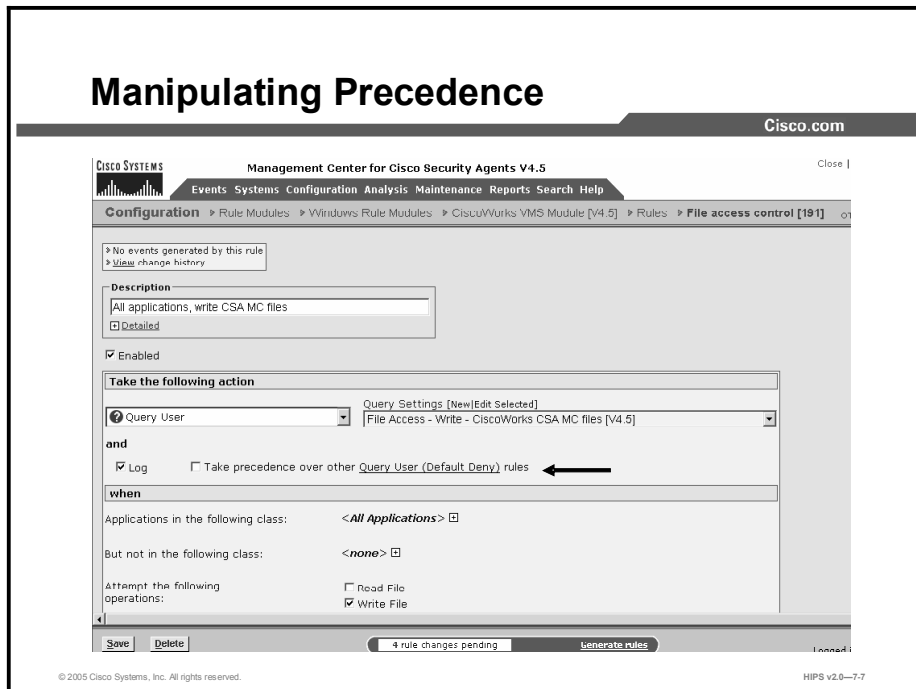
- **Monitor:** Most rule types provide a "Monitor" action. You may want to write monitor rules to produce events for certain resource accesses without having to write an explicit allow or deny rule about that resource. You can also write monitor rules in the presence of other similar rules dictating the availability (allow, deny, etc.) of a resource. For example, you can write a monitor rule for a resource and if the monitor rule is triggered, any subsequent rules about the resource in question will trigger as well. This allows you to configure general alerts about resources regardless of whether the resource access action is an allow or a deny.

For every rule module you configure, the default action of that rule is allow. All rule modules allow all system actions until you write a rule denying a specific action. Following that logic, you do not need to write allow rules unless they make exceptions to deny rules that you are

writing within a module or for monitoring purposes. If you do write a stand-alone allow rule, because the default action is allow, the allow rule itself is essentially irrelevant.

A useful model for configuring rules within modules is 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, it allows all system actions. First, write a deny rule, and then if you want to make any exceptions to that particular deny, write an allow rule. Next consider using query rules for access controls that allow the user to decide whether an action should be allowed or denied. Last of all, write any high priority rules you might need.

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, such as 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. There are cases, however, when the CSA MC ordering scheme causes policies to behave in an undesired manner. For this reason, most rule types provide a check box that allows you to manipulate how similar rule types are subordered within a policy. This check box, Take Precedence Over Other <Action> Rules, is located in the rule configuration page. A rule with this precedence check box selected is evaluated before similar rules that do not have this check box selected.

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

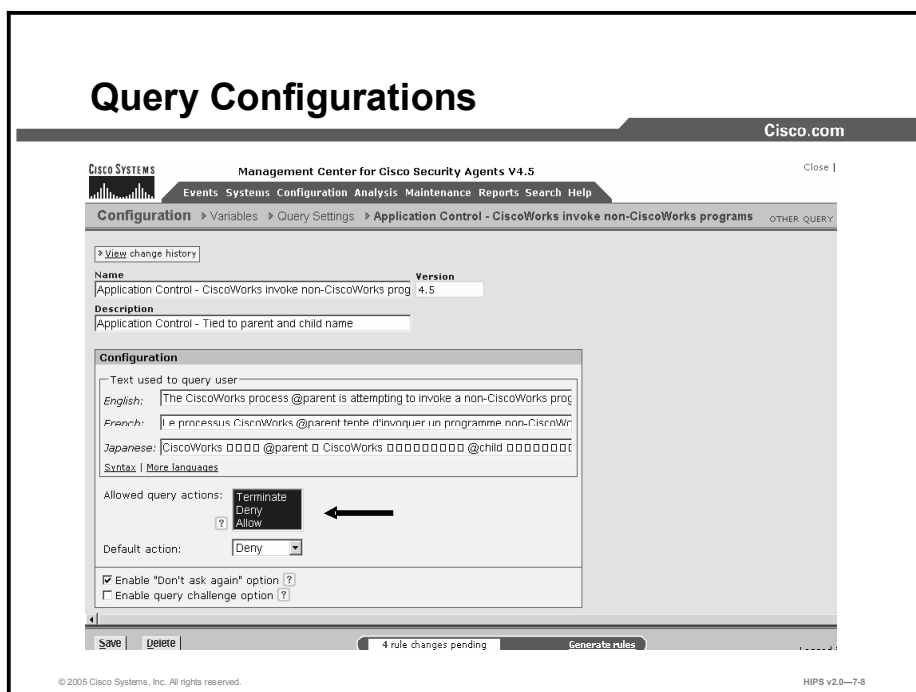
- Log, Deny, All applications, acting as a server, for TCP/1-65000
- No Log, Deny, All applications, acting as a server, for TCP/1900

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 Type> Rules check box in the TCP/1900 rule would allow you to designate its precedence as higher than other deny rules in the policy. This would suppress log messages for actions that 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 Type> Rules check box is a rule-ordering tool that you should rarely need. In most cases, the CSA MC automatic ordering of rules is sufficient. But if you are using this check box 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 check box On/Off * Log check box 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.



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, deny it, or terminate it. When you select to query the user, you also craft explanatory text to display to the user and you decide whether to allow, deny, or terminate the action by default if the query is not answered within five minutes. If the user is not logged in to the system, the default action is taken immediately.

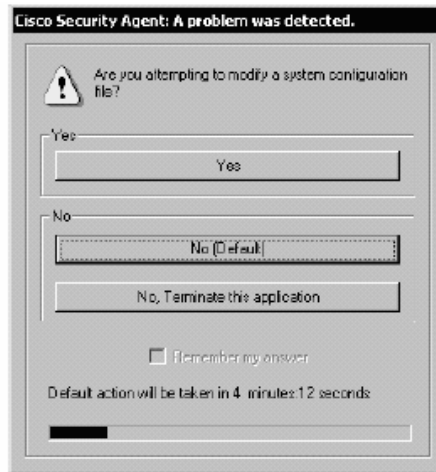
Query configurations offer a variety of settings. You can decide which buttons are displayed in the popup query box, which action is the default, what the content of the query text will be, and whether the answer given by the user is to be remembered.

When an action is attempted on a system where a query user rule is triggered, a popup box appears on the system where the resource is located. From the Query Settings page, accessible from the Configuration > Variables menu, you configure the query text and the query buttons that appear in the popup box that the end user will see.

Caution For Solaris rules, query user options are not available. For Windows and Linux Agents, if the Agent user interface (UI) is hidden for the group, no query user popup boxes are displayed. The default is immediately taken on all query user rules and heuristics that are present in the assigned policies.

Query User Popup Box

Cisco.com



In the query popup box, the user reads the information given on the attempted action and selects one of the following possible choices:

- **Yes:** Allows the application access to the resource in question.
- **No:** Denies the application access to the resource in question.
- **Terminate:** Denies the application access to the resource in question and also attempts to terminate the application process. (Some processes cannot be safely terminated, such as winlogon.)
- **Default action:** Of the buttons you decide to display, you also choose one of those buttons to be the default action. If the query is not answered by the user within 5 minutes, or if the user is not logged in to the system, the default action is taken immediately.
- **Don't ask again:** For each button option that appears on the query box, you can also decide to display a "Don't ask again" check box so that the user's query response is remembered. Users can select that check box when they respond to the query, and if they attempt the same action on the same resource, the response is remembered and they are not queried again.
- **Query challenge:** For added security, you can issue a query challenge on the query popup box. This ensures that the user sitting in front of the system is answering the query rather than a malicious remote user or program. To pass the challenge, the user enters the information displayed in a graphic on the popup box itself.

When you configure your query settings for the rule, the text you type into the Text Used to Query User field on the Query Settings page is the same text that appears at the top of the query user popup box. Therefore, write the text in a way that describes the system action that triggered the popup.

Caution With file access control rules, the query user popup 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 popup 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 It takes several seconds for buttons on the query user popup box to become active (selectable).

When users are queried, the Agent can remember the response permanently or temporarily. This way, if the same rule is triggered again, the action is allowed, denied, or terminated based on what answer was given previously, with no popup query box appearing again either permanently or for some period of time.

For example, consider a user who is queried as to whether an application can talk on the network, and the user responds by clicking the Yes button and a Don't Ask Again check box. The yes response is remembered permanently, and that response appears in the edit field in the Agent UI query response window. A second user is queried as to whether setup.exe can install software on the system, and the user responds by clicking the Yes button, but there is no Don't Ask Again check box, or it is there but the user does not select it. This response is remembered temporarily and does not appear in the Agent UI query response window.

If the user's response is only cached temporarily (for approximately an hour), the user can click the Clear button in the query response window to delete all temporarily cached responses. To clear permanent responses listed in the edit field, the user must select the response in the edit field and click the Delete key.

Note Permanent responses are remembered across reboots. Temporarily cached responses are not remembered across reboots. Also note, a query response is tied to the user who responded. On multiuser machines, multiple users may be asked the same question.

You should note how CSA MC manages rule priorities if multiple similar query rules need to be evaluated.

- Base Priority:
 - Allow
 - Deny Terminate/no challenge/no don't ask again/no logging
- Relative priorities for query options that are turned on are as follows (top to bottom):
 - Challenge/Don't ask again/Logging
 - Challenge/Don't ask again
 - Challenge/Log
 - Don't ask again/Log
 - Don't ask again
 - Log

Rules Common to Windows and UNIX

This topic discusses rules available for both Windows and UNIX policies. Rules that apply to Windows systems only and rules exclusive to UNIX policies are discussed separately later in the lesson.

Rules Available for Windows and UNIX

Cisco.com

- **Agent service control**
- **Agent UI control**
- **Application control**
- **Connection rate limit**
- **Data access control**
- **File access control**
- **Network access control**

© 2005 Cisco Systems, Inc. All rights reserved.HIPS v2.0—7-11

The following seven rules apply to both Windows and UNIX systems:

- Agent service control rules
- Agent UI control rules
- Application control rules
- Connection rate limit rules
- Data access control rules
- File access control rules
- Network access control rules

Add Rule Menu

The screenshot displays the Cisco Management Center for Cisco Security Agents V4.5 interface. At the top, there is a navigation bar with the Cisco logo and the text "Cisco.com". Below this, the main header reads "Management Center for Cisco Security Agents V4.5" with a "Close" button on the right. A secondary navigation bar includes tabs for "Events", "Systems", "Configuration", "Analysis", "Maintenance", "Reports", "Search", and "Help".

The main content area features a table of rules. Each row includes a checkbox, a rule ID, a rule name, a status, and a description. The rules listed are:

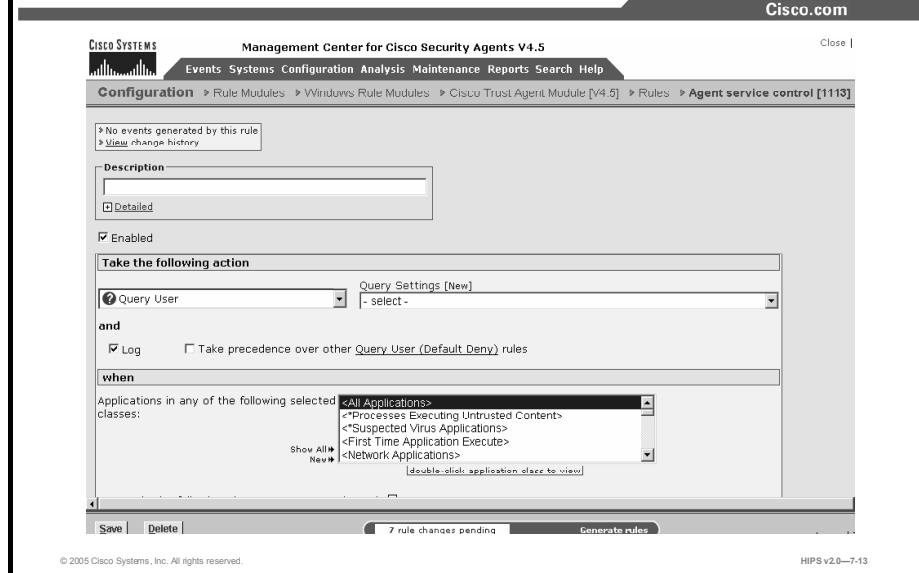
Rule ID	Rule Name	Status	Description
304	Network access control	Enabled	Allow the Cisco Trust Agent to communicate with its peers
305	Application control	Enabled	Permit the Cisco Trust Agent to run related trusted Applications
311	Application control	Enabled	Permit the Cisco Trust Agent to run Virus scanner apps
1112	Agent service control	Enabled	
306	File access control	Enabled	Query the user when an attempt is made to modify any Cisco Trust Agent
307	Application control	Enabled	Prevent the Cisco Trust Agent from running other applications
308	File access control	Enabled	Prevent the Cisco Trust Agent from writing files it should not
309	Network access control	Disabled	Prevent the Cisco Trust Agent from accepting network connections
310	Network access control	Enabled	Prevent the Cisco Trust Agent from making network connections

Below the table is the "Add Rule" section. It includes a "Copy" button, a "to" dropdown menu, and a "rule module" dropdown menu set to "Cisco Trust Agent Module [V4.5]". A list of available rule types is shown, with an arrow pointing to the "File access control" rule type.

At the bottom of the interface, there are buttons for "Delete", "Enable", and "Disable", along with a status indicator "6 rule changes pending" and a "Generate rules" button. The footer contains the text "© 2005 Cisco Systems, Inc. All rights reserved." and "HIPS v2.0-7-12".

To add rules to your module, click the **Add Rule** link at the bottom of the rule list. A popup list of the available rule types appears.

Agent Service Control Rules



Use the Agent service control rule to control whether administrators are allowed to stop Agent security. This is via a **net stop** command on Windows or via `/etc/init.d/csa stop` on UNIX. Stopping Agent security disables all rules until security is manually resumed or the system is rebooted.

If you use this rule to deny Agent service stops, the Agent service cannot be stopped on the system in question, and therefore Agents cannot be uninstalled.

Note Although Agents cannot be uninstalled by administrative users if this rule denies stopping the Agent service, this rule does not prevent Agent software updates from occurring.

You can also use the Agent service control rule to monitor, terminate, or tag a process that attempts to modify the Agent configuration. Here are the steps for configuring rule modules:

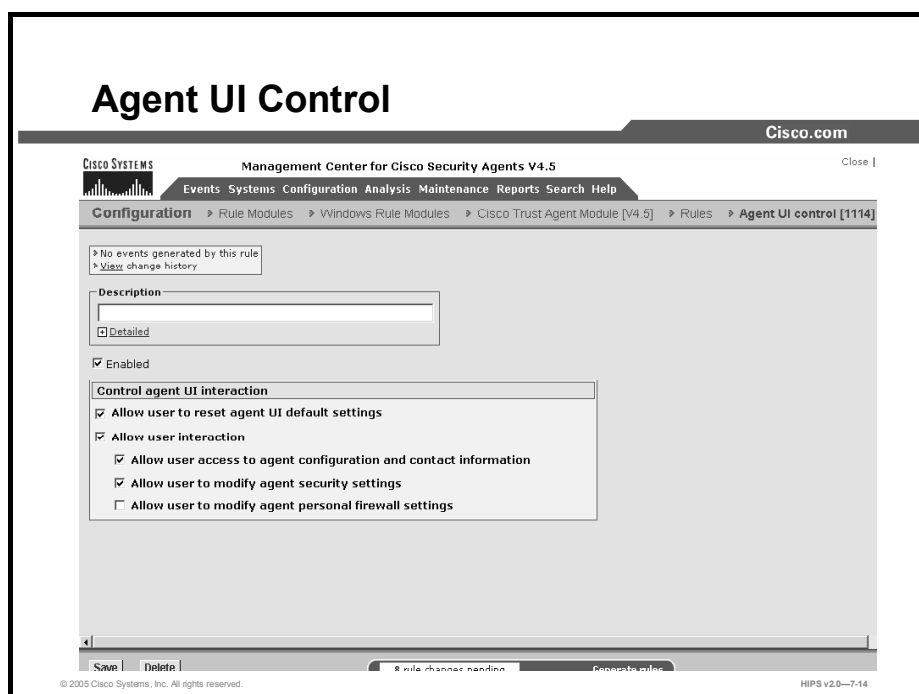
- Step 1** To add rules to your module, click the **Add Rule** link at the bottom of the rule list. A popup list of the available rule types appears.
- Step 2** Select the **Agent Service Control** rule. This takes you to the configuration view for this rule type.
- Step 3** In the Agent service control rule configuration view, enter the following information:
 - **Description:** Enter a description of this rule. This description appears in the list view for the module. Optionally, expand the Detailed field to enter a longer description.
 - **Enabled:** Use this check box to enable this rule within the module. It is enabled by default. By not selecting this check box, you can save this rule, but it will not be active in the module, and it will not be distributed to groups.

- Step 4** Select an action type from the drop-down list under Take the Following Action. (Note that not all action types are available for this rule.)
- Step 5** Select one of the following check boxes:
- **Log:** Enable this check box 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 Other <Action Type> Rules:** Enable this check box to manipulate rule precedence so that this rule is evaluated before other similar rules. Generally, you should not require this check box. Do not use it without understanding how it works.
- Step 6** In the section under When, you select applications that will be subject to this rule.
- **Applications in Any of the Following 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 deselect it and create your own application classes.

Note On UNIX systems, anyone with root access can stop the Agent service. To prevent this, but still allow 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 that allows only a UNIX Secured Management application class to stop the service.

- **But Not in any of the Selected Classes:** Optionally, select application classes here to exclude from the application class(es) you have selected in the included applications field. Note that the entry <None> is selected by default.
- **Attempt to Stop the Agent Service:** This check box 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 on Windows (or via `/etc/init.d/csa stop` on UNIX).
- **Attempt to Modify Local Agent Configuration:** The Cisco Security Agent has built-in global security policies that protect Agent binaries and data. (Note that this protection is only offered when the Agent service is running and is not stopped or in test mode.) Although you cannot turn these nonlogged, built-in rules off while the Agent is active, you can use this rule to monitor, terminate, or tag a process that attempts to modify the Agent configuration.

- Step 7** Click the **Save** button.



Use the Agent UI rule to control how the Agent user interface is displayed to end users. In the absence of this rule, end users have no visible Agent UI. If this rule is present in a module, you can select to display the Agent UI and one or more controls to the end user. These controls give the user the ability to change certain aspects of their Agent security.

Note This rule only applies to Windows and Linux platforms. The Agent UI is not supported on Solaris systems.

Optional controls are as follows:

- **Allow User to Reset Agent UI Default Settings:** On Windows, this is available from the Start > Programs > Cisco Systems menu. By selecting this option, users can reset Agent UI functionality to the original factory default settings, and all user-set controls are lost. This is useful on Windows platforms where different users with varying user Agent permission settings may log into the same machine.
- **Allow User Interaction:** Selecting this check box causes the end user to have a visible and accessible Agent UI, including a red flag in the system tray. With no other subsequent check boxes selected, the Agent UI contains a status view, a messages page to view Agent events, and the ability to clear persistent and temporary query user responses. If this rule is present in a module, but this check box is not selected, the end user will have no visible Agent UI. In the presence of two or more Agent UI control rules, these rules are combined, and selected check boxes take precedence over unselected check boxes. Add one or more additional controls as follows:
 - **Allow User Access to Agent Configuration and Contact Information:** Selecting this check box allows end users to enter contact information in the Agent UI. They also have access to detailed status information and a poll button to force manual polling if the MC becomes available.
 - **Allow User to Modify Agent Security Settings:** Selecting this check box provides end users with the ability to alter their security level by moving a sidebar between

Off, Low, Medium, and High (in accordance with policies) and to manage the classification of untrusted content. This check box also controls the Agent Off feature available from the Agent UI. Allowing this action (moving the sidebar to Off) permits all users (including nonadministrative users) to disable all rules on the Agent until they are reenabled by the user. 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 turned off.)

- **Allow User to Modify Agent Personal Firewall Settings:** Selecting this check box provides the end user with the ability to dictate which applications are allowed network access. They also gain a file protection capability by which they can enter the names of local files that network applications are not allowed to access on their system.

Hiding the Agent UI

Not enabling the Allow User Interaction check box in this rule has the following effects.

■ Software updates

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

■ Queries

- When no Agent UI is present, no query user popup boxes are 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

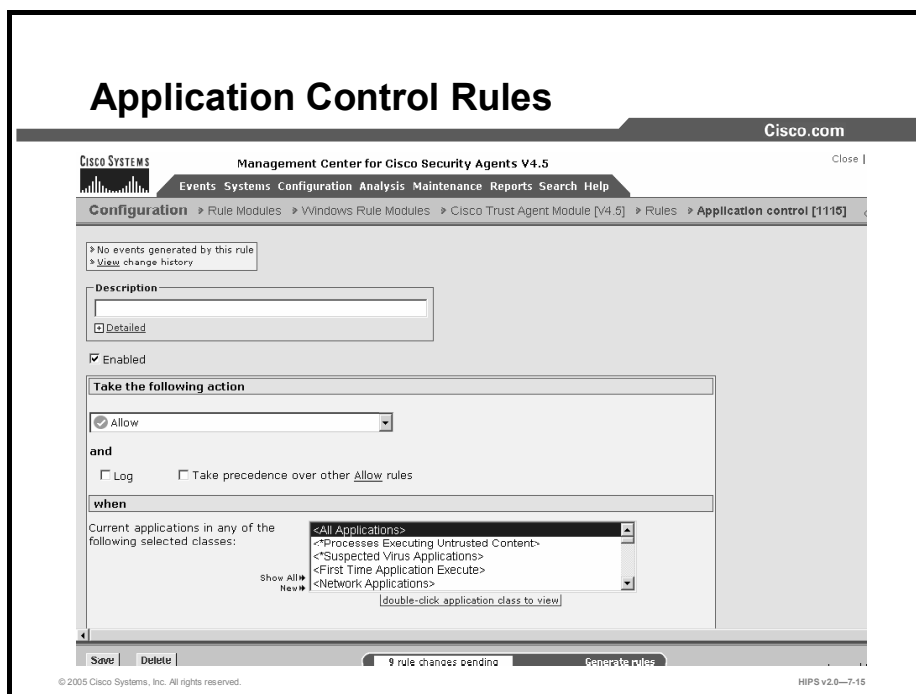
- No messages to inform users that actions have been denied and why
- No ability to clear the cache or reenabling logging
- No fast polling ability
- No end-user contact information sent to CSA MC

■ Hidden Agent UI feature notes

- If a host belongs to multiple groups with multiple policies, 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).

- When no Agent UI is present, no query user popup boxes are displayed. The default is immediately taken on all query user rules and heuristics that are present in the assigned policies.
- If an end-user system already has an Agent UI installed, when you deselect the Allow User Interaction check box and generate rules, the Agent UI disappears when the new rules are downloaded.

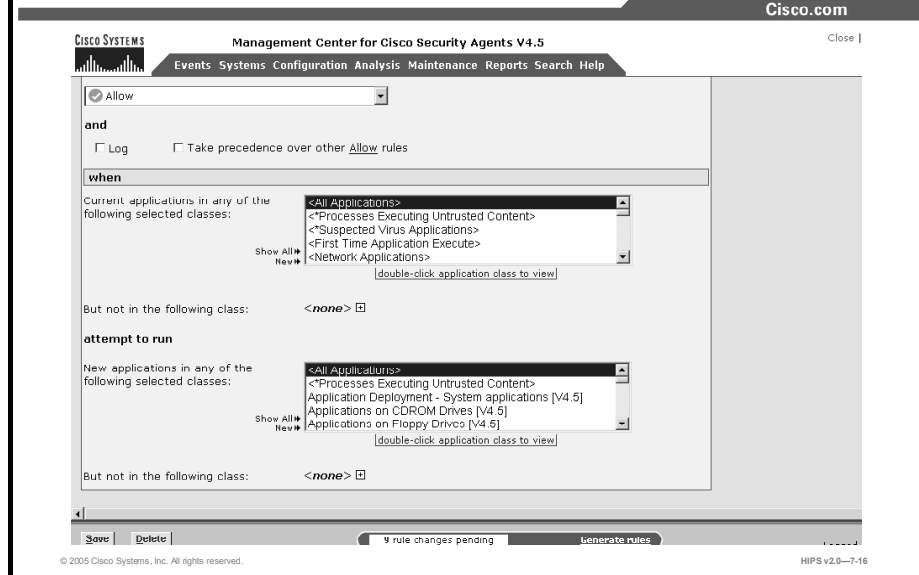


Use application control rules to control what applications can run on designated Agent systems. This rule type does not control what application can access what resources, as do other access control rules. This rule type can stop selected applications from running on systems. If you deny an application 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 was invoked by an application that downloaded content from the network.

- Step 1** To add rules to your module, click the **Add Rule** link at the bottom of the rule list. A popup 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.
- 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 module. Optionally, expand the Detailed field to enter a longer description.
 - **Enabled:** Use this check box to enable this rule within the module. (It is enabled by default.) By not selecting this check box, 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 an action type from the drop-down list under Take the Following Action.

Application Control Rules (Cont.)



Note Creating dynamic application classes from the application control rule differs 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.

Step 5 Select one of the following check boxes:

- **Log:** Enable this check box 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 Other <Action Type> Rules:** Enable this check box to manipulate rule precedence so that this rule is evaluated before other similar rules. Generally, you should not require this check box. Do not use it without understanding how it works.

Step 6 In the section under When, you select applications that will be subject to this rule.

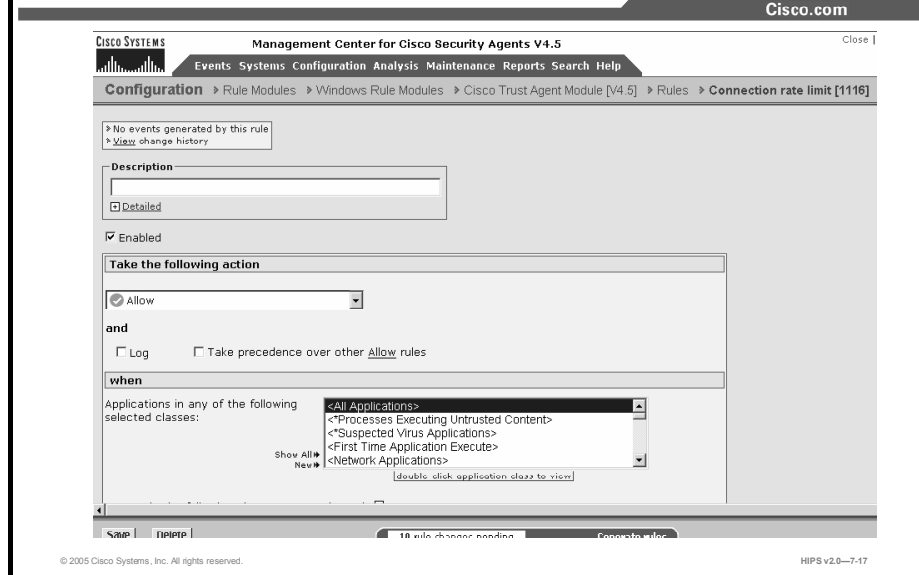
- **Current Applications in any of the Following Selected Classes:** If you want to control (allow or deny) an application that is running on a system no matter how it is invoked, allow All Applications to remain selected by default. Then you will select the application you want to control from the second 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.
- **But Not in the Following Class:** Optionally, select application classes here that you want to exclude from the application class(es) you have selected in the included applications field. Note that the entry <None> is selected by default.

- Step 7** In the attempt to run section, select the following options:
- **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 that you choose in the top field. If you select All Applications in the top application field, you cannot select All Applications in this second field. (If you did, all applications would be completely prevented from running on systems if this were a deny rule.)
 - **But Not in Any of the Selected Classes:** Optionally, select application classes here that you want to exclude from the application class(es) you have selected in the included applications field. Note that the entry <None> is selected by default.

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

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

Connection Rate Limit Rules



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, such as denial of service attacks (server connection rating limiting) and in preventing the propagation of denial of service attacks (client connection rate limiting).

Note These instructions are a continuation of Configuring Rule Modules.

Click the **Modify Rules** link at the top of the Rule Module page to go to the Rules page.

Step 1 To add rules to your module, click the **Add Rule** link at the bottom of the rule list. A popup 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 module. Optionally, expand the Detailed field to enter a longer description.
- **Enabled:** Use this check box to enable this rule within the module. It is enabled by default. By not selecting this check box, you can save this rule, but it will not be active in the module and it will not be distributed to groups.
- **Log:** Use this check box to enable logging within the module.

Connection Rate Limit Rules (Cont.)

The screenshot shows the Cisco Management Center for Cisco Security Agents V4.5 interface. The main window is titled "Take the following action" and contains the following configuration options:

- Take the following action:** A dropdown menu is set to "Allow".
- and:** There are two checkboxes: "Log" (unchecked) and "Take precedence over other Allow rules" (unchecked).
- when:**
 - Applications in any of the following selected classes:** A list box contains "<All Applications>", "<*Processes Executing Untrusted Content>", "<*Suspected Virus Applications>", "<*First Time Application Execute>", and "<Network Applications>". The "<All Applications>" entry is selected.
 - But not in the following class:** A dropdown menu is set to "<none>".
 - Attempt to act as a:** A dropdown menu is set to "server".
 - Communicating with:** A dropdown menu is set to "specific" hosts.
 - Under limit of:** A text input field contains "100" network connections.
 - In:** A text input field contains "5" minutes.

At the bottom of the window, there are "Save" and "Delete" buttons, a status bar indicating "19 rule changes pending", and a copyright notice: "© 2005 Cisco Systems, Inc. All rights reserved. HIPS v2.0-7-18".

Step 4 Select an action type from the drop-down list under Take the Following Action.

Note You cannot configure Query User Connection rate limit rules.

Step 5 In the section under When, you select applications that will be subject to this rule.

- **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 control. Note that the entry <All Applications> is selected by default. You can use this default, or you can deselect it and create your own application classes.
- **But Not in Any of the Selected Classes:** Optionally, select application classes here that you want to exclude from the application class(es) you have selected in the included applications field. Note that the entry <None> is selected by default.

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

Step 6 From the drop-down list for “Attempt to act as a,” 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’s connection, select **client**.

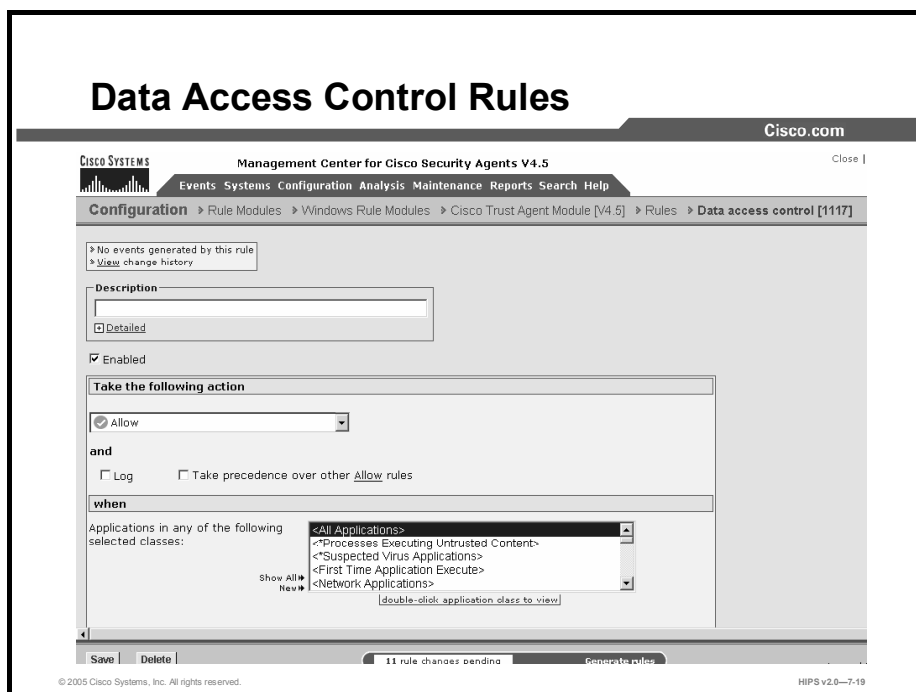
Step 7 In the next drop-down list, “Communicating with,” select specific hosts or all hosts. When the rate limit (set in the next box) is reached, you use this to determine whether all subsequent service requests are dropped or only those received or sent by a specific host. Selecting a specific host indicates that the host in question

exceeded the rate limit. Selecting all hosts indicates that the sum total of connections to and from all hosts exceeds the limit, and all hosts are blocked.

Step 8 By default, reasonable values are set for the number of network connections in a time frame. They define the number of connections that can normally be expected during a time frame, either by specific hosts or all hosts.

- If you select an action type other than allow, and the limit is exceeded (Over) in this time frame (an abnormal amount of connections that could represent an attack on the system), subsequent connection requests are dropped. (The dropped connections can be those received to and from individual, specific, hosts or to and from all hosts, according to this setting.)
- If you configure this as an allow rule, you are setting a connection limit Under which the number must remain.

Step 9 When you are finished configuring your connection rate limit rule, click the **Save** button.



Use data access control rules on web servers to detect malformed web server requests that 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 uniform resource identifier (URI) portion of the HTTP request.

An HTTP request has several elements:

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

The data access control rule examines patterns in the URI portion of the HTTP request. The preconfigured data sets group patterns to match based on these criteria:

- Functional associations of meta-characters (for example, "(" 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 CSA 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.

Caution 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.

Note These instructions are a continuation of Configuring Rule Modules.

Click the **Modify Rules** link at the top of the Rule Module page to go to the Rules page.

Step 1 To add rules to your module, click the **Add Rule** link at the bottom of the rule list. A popup 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.

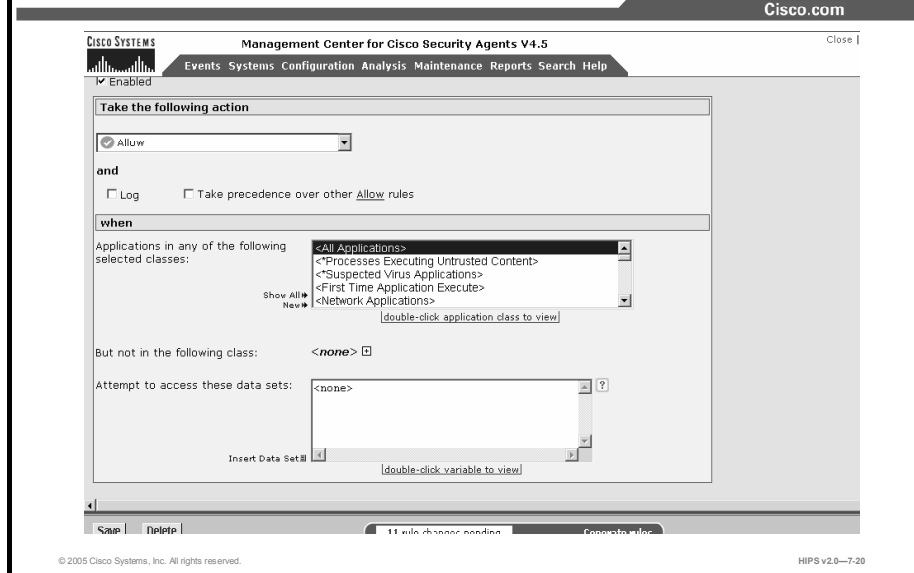
Step 3 Enter the following information for the rule:

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

Step 4 Select an action type from the drop-down list under Take the Following Action. Select one of the following check boxes:

- **Log:** Enable this check box 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 Other <Action Type> Rules:** Enable this check box to manipulate policy rule precedence so that this rule is evaluated before other similar rules. Generally, you should not require this check box. Do not use it without understanding how it works.

Data Access Control Rules (Cont.)



Step 5 In the section under When, you select applications that will be subject to this rule.

- **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 control.

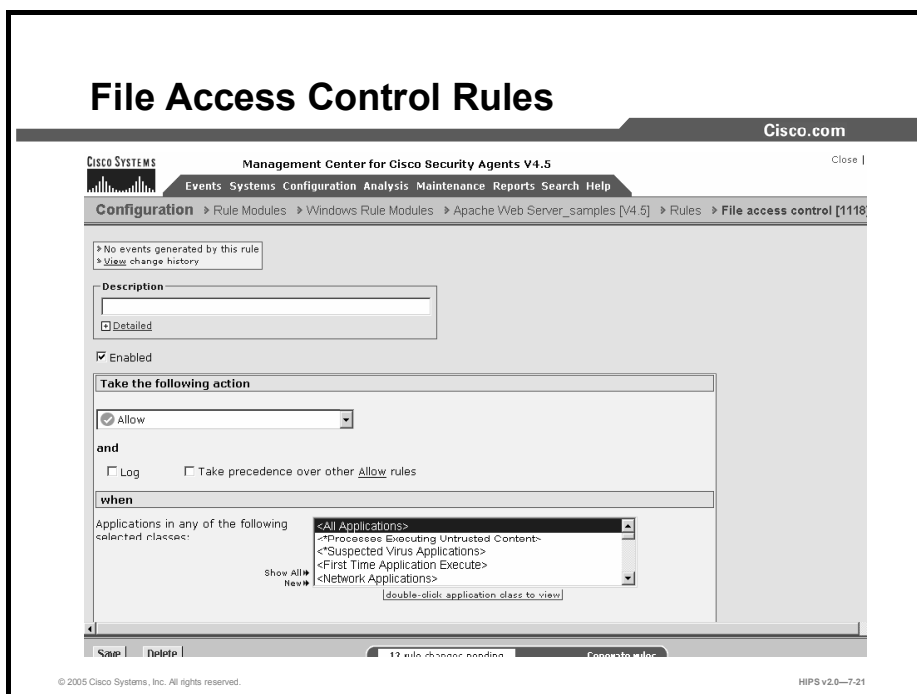
Note The entry <All Applications> is selected by default. You can use this default, or you can deselect it and create your own application classes.

- **But Not in Any of the Selected Classes:** Optionally, select application classes here to exclude from the application class(es) that you have selected in the included applications field. Note that the entry <None> is selected by default.

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

Step 6 Click the **Insert Data Set** link to enter a preconfigured data set in the Attempt to Access These Data Sets box. When you click this link, a list of the data sets that you have 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.

Step 7 When you are finished configuring your data access control rule, click the **Save** button.



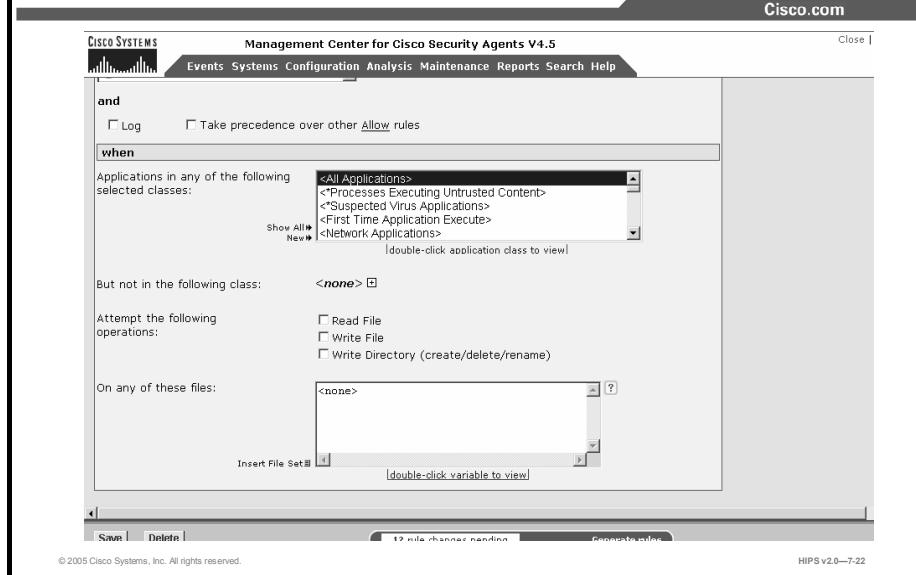
Use file access control rules to allow or deny operations (read, write) that selected applications can perform on files. File protection encompasses read/write access. Directory protection encompasses directory deletes, renames, and new directory creation.

Note These instructions are a continuation of Configuring Rule Modules.

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 module, click the **Add Rule** link at the bottom of the rule list. A popup 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 module. Optionally, expand the Detailed field to enter a longer description.
 - **Enabled:** Use this check box to enable this rule within the module. (It is enabled by default.) By not selecting this check box, you can save this rule, but it will not be active in the module and it will not be distributed to groups.
- Step 4** Select an action type from the drop-down list under Take the Following Action.

File Access Control Rules (Cont.)



Step 5 Select one of the following check boxes:

- **Log:** Enable this check box 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 Other <Action Type> Rules:** Enable this check box to manipulate rule precedence so that this rule is evaluated before other similar rules. Generally, you should not require this check box. Do not use it without understanding how it works.

Step 6 In the section under When, you select applications that will be subject to this rule.

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 control. Note that the entry <All Applications> is selected by default. You can use this default, or you can deselect it and create your own application classes.

- **But Not in Any of the Selected Classes:** Optionally, select application classes here to exclude from the application class(es) that you have selected in the included applications field. Note that the entry <None> is selected by default.

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

Step 7 In the check boxes next to Attempt the Following Operations, select the operations Read File and/or Write File that you are allowing or denying on the specified files. For directory protection, the actions you are allowing or denying are Create, Delete, and Rename.

Step 8 In the list box next to On Any of These Files, click the **Insert File Set** link to enter a preconfigured file set. When you click this link, a list of the file sets that you have configured appears here, allowing you to select one or more. Instead of file sets, you can list the literal files that 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.

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.

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.

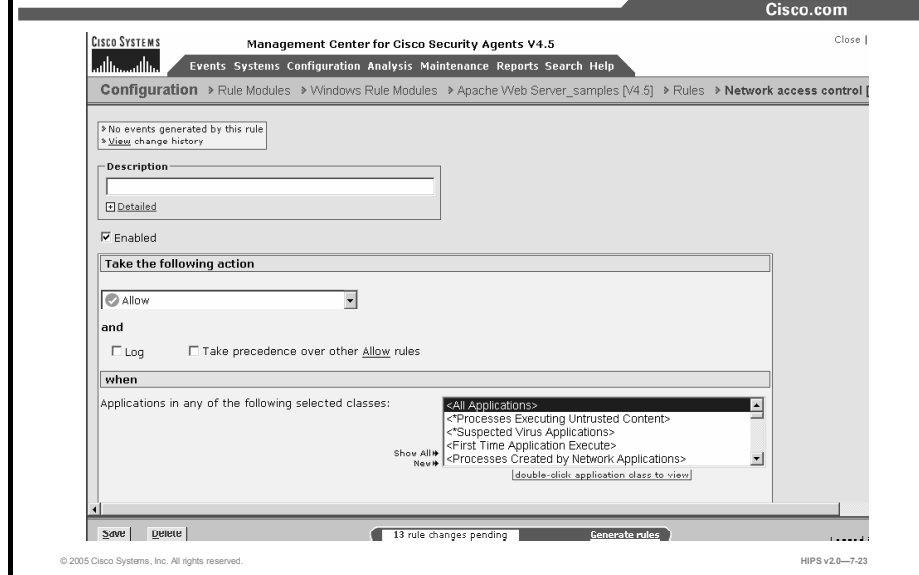
Caution Symbolic links: For UNIX, if you create a file access control rule to protect a symbolic link, *only* that symbolic link is protected. The underlying resource, unless also specified, is *not* protected. For example, a file access control rule written for /etc/hosts does not protect /etc/inet/hosts. 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 e-mail worm events, correlated virus scanner log messages, or files that were added manually by the administrator. This list updates automatically (dynamically) as logged quarantined files are received. To view the files that are added to the dynamically quarantined files list and to manually add files to be quarantined, click the **Manage Dynamically Quarantined Files** link on the Global Event Correlation page.

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 To distribute rules to the correct hosts, you must associate policies with groups and then generate rules.

Network Access Control Rules



Use network access control rules to control access to specified network services and network addresses. You can also use this rule type to listen for applications attempting to offer unknown or unsanctioned services.

Note The following instructions are a continuation of Configuring Rule Modules.

- Step 1** To add rules to your module, click the **Add Rule** link at the bottom of the rule list. A popup 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 module. Optionally, expand the Detailed field to enter a longer description.
 - **Enabled:** Use this check box to enable this rule within the module. (It is enabled by default.) By not selecting this check box, 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 an action type from the drop-down list under Take the Following Action.
- Step 5** Select one of the following check boxes:
- **Log:** Enable this check box 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 Other <Action Type> Rules:** Enable this check box to manipulate rule precedence so that this rule is evaluated before other similar

rules. Generally, you should not require this check box. Do not use it without understanding how it works.

Network Access Control Rules (Cont.)

The screenshot shows the Cisco Management Center for Cisco Security Agents V4.5 interface. The 'when' section is expanded, showing a list of application classes. The 'All Applications' class is selected. The 'But not in the following class' field is set to '<none>'. The 'Attempt to act as a' field is set to 'client'. The 'Communicating with host addresses' and 'Using these local addresses' fields are both set to '0.0.0.0-255.255.255.255'. The interface includes a 'Save' button, a 'Delete' button, and a 'Generate rule' button. The status bar indicates '13 rule changes pending'.

Step 6 In the section under When, you select the applications that will be subject to this rule.

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 control. Note that the entry <All Applications> is selected by default. You can use this default, or you can deselect it and create your own application classes.

- **But Not in Any of the Selected Classes:** Optionally, select application classes here to exclude from the application class(es) that you have selected in the included applications field. Note that the entry <None> is selected by default.

Step 7 From the drop-down list for “Attempt to act as a,” you can select server or client, or both, or select listener, depending on the direction or type of connection that you are controlling or listening for. If you are limiting a server’s contact with clients, select **server** here and enter the address(es) of client(s) in the host addresses field. If you are limiting a client’s contact with a server, select **client** here and enter the address(es) of server(s) in the host addresses field.

Step 8 For network services, enter the literal protocol/port number combination for the service that you want to control access to, or click the **Insert Network Service** link to enter a preconfigured network service variable. When you click this link, a list of the network service variables that you have 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.

Step 9 For Communicating with Host Addresses, enter the literal network address(es) for the clients or servers that you want to control access to, or click the **Insert Network Address Set** link to enter a preconfigured network address set variable here.

If you select server in the previous drop-down list, you enter client addresses here. If you select client in the list, you enter server addresses here. Note that you can use network address set variables.

- You can also use the following shorthand 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 shorthand 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 to network access.

Step 10 For Using These Local Addresses, enter the literal network address(es) for the local system addresses that you want to control (that is, control clients making connections from or control servers making connections to). You can also click the **Insert Network Address Set** link to enter a preconfigured network address set variable here.

The addresses or address ranges you enter here are used to control the host that initiates the network connection. For example, you could write a network access control rule that would only allow laptop users to connect to an internal network database if their connection is coming through a Virtual Private Network (a 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 Virtual Private Network 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 to network access (intra-box connections).

Use **@dynamic** in the addresses set field to indicate untrusted hosts that have been quarantined by CSA MC. Addresses are added to this list when they are seen as an untrusted host. (The built-in “Processes Communicating with Untrusted Hosts” is triggered in a rule.) This list updates automatically (dynamically) as logged quarantined addresses are received.

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

This rule is now part of your rule module. 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.

Caution In order to distribute rules to the correct hosts, you must associate policies with groups and then generate rules.

Note You can use multiple access control rules in one policy.

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

The “listener” option: You can use the listener option in a network access control rule to indicate what applications have the ability to be a server before they are allowed to accept a server connection. This is in contrast to the server option, which offers real-time per-connection control. The listener option can be used in a monitoring capacity to reveal any applications that are attempting to offer a network service. For example, if a system is already infected with a Trojan horse, that Trojan horse may be listening on a high-numbered port for a network server connection. A NACL listener rule would detect this before a server connection is achieved. You could then craft a subsequent NACL rule to deny the server connection.

Windows-Only Rules

The rules covered in this topic are available only for Windows policies.

Rules Available for Windows Only

Cisco.com

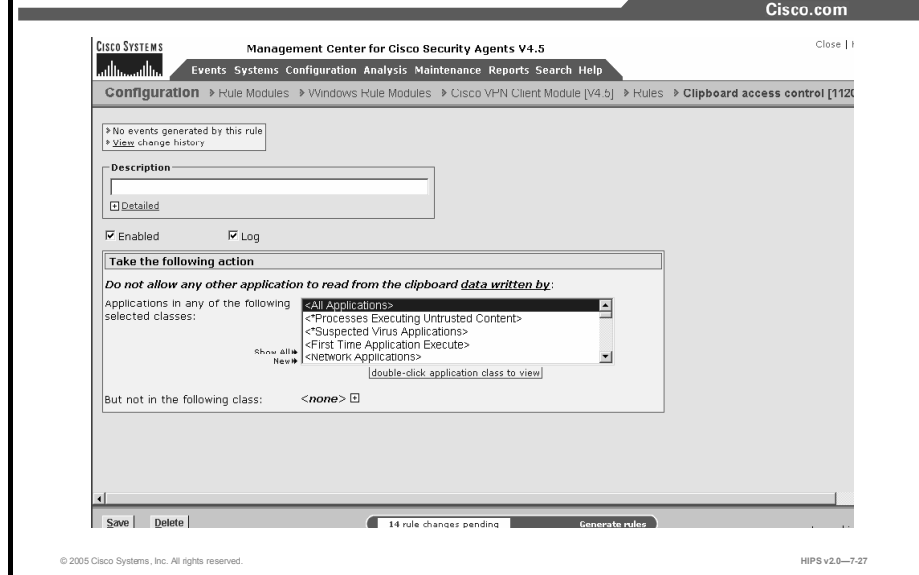
- **Clipboard access control**
- **COM component access control**
- **File version control**
- **Kernel protection**
- **NT Event Log**
- **Registry access control**
- **Service restart**
- **Sniffer and protocol detection**

© 2005 Cisco Systems, Inc. All rights reserved.HIPS v2.0—7-26

The following eight rules apply to Windows systems only:

- Clipboard access control rules
- COM component access control rules
- File version control rules
- Kernel protection rules
- NT Event Log rules
- Registry access control rules
- Service restart rules
- Sniffer and protocol detection rules

Clipboard Access Control



Use the clipboard access control rule to dictate which applications can access information that is written to the clipboard. When writing security policies, you may want to protect information from being accessed by other applications or network processes. To fully protect this information, you must consider preventing other applications from accessing protected information that may have been written to the clipboard.

These instructions are a continuation of Configuring Rule Modules.

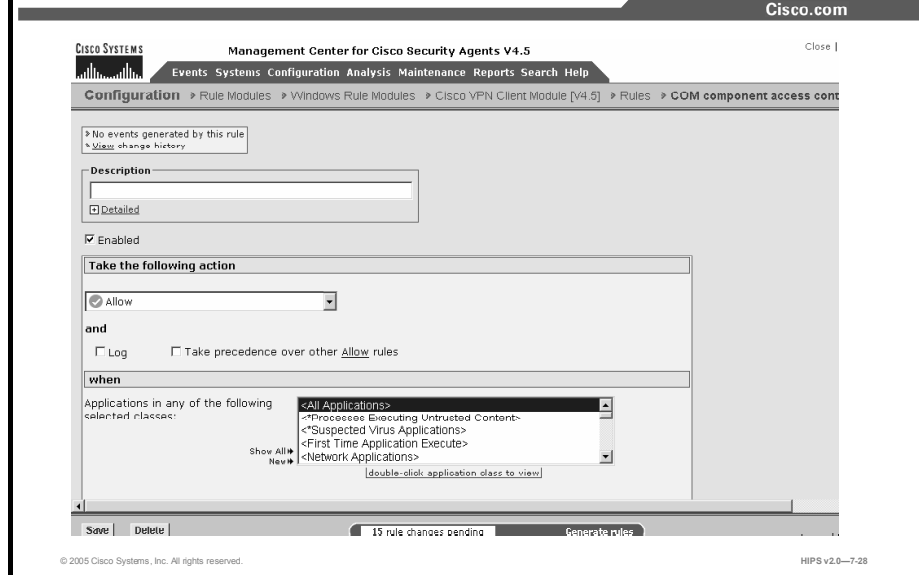
- Step 1** To add rules to your module, click the **Add Rule** link at the bottom of the rule list. A popup list of the available rule types appears.
- Step 2** Select the **Clipboard Access Control** rule. This takes you to the configuration view for this rule type.
- Step 3** Enter the following information for the rule:
- **Description:** Enter a description of this rule. This description appears in the list view for the module.
 - **Enabled:** Use this check box to enable this rule within the module. (It is enabled by default.) By not selecting this check box, you can save this rule, but it will not be active in the module and it will not be distributed to groups
- Step 4** Under “Do not allow any other application to read from the clipboard data written by,” select the applications that will be subject to this rule.
- **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 control. Note that the entry <All Applications> is selected by default. You can use this default, or you can deselect it and create your own application classes. When your rule is configured, currently selected application classes appear at the top of the list.

- **But Not in Any of the Selected Classes:** Optionally, select application classes here to exclude from the application class(es) that you have selected in the included applications field. Note that the entry <None> is selected by default.

Step 5 When you are finished configuring your clipboard access control rule, click the **Save** button.

Note If you are using the clipboard rule to restrict applications from accessing data on the clipboard, the system Print Screen functionality is also automatically disabled.

COM Component Access Control Rules



Use COM component access control rules to allow or deny applications access to 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 that can be used to create macros or utility scripts. Although this is useful functionality, it can be used maliciously by an inadvertently downloaded Visual Basic script.

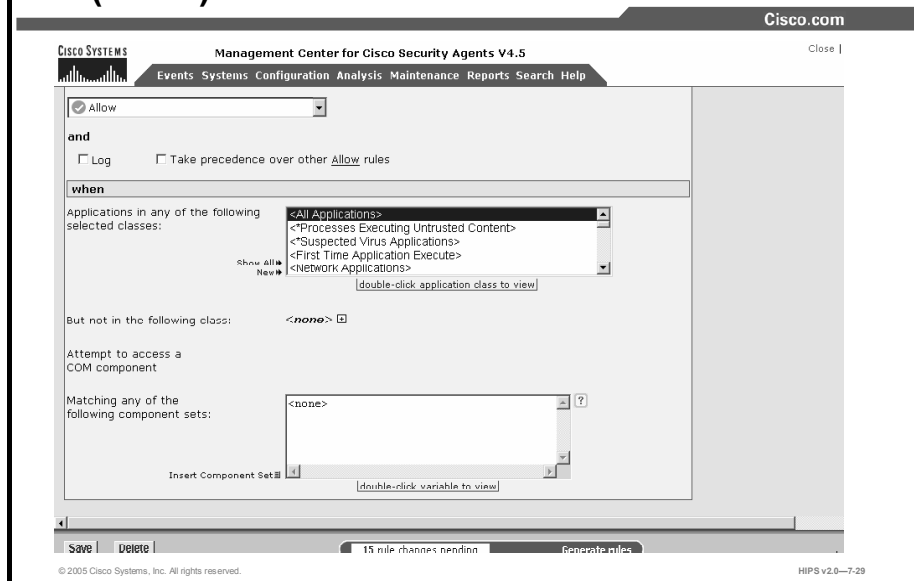
An example is the Mydoom virus, which propagated by using the "Outlook.Application" COM component to send itself to each entry in the local address book. Using the COM component access control rule, you can protect specific COM components. For example, you could create a rule that limits access to Office 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 that installs with each Cisco Security Agent. Running this utility extracts all COM component program identifiers (PROGIDs) and class identifiers (CLSIDs) for software running on the system in question.

These instructions are a continuation of Configuring Rule Modules.

- Step 1** To add rules to your module, click the **Add Rule** link at the bottom of the rule list. A popup 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.

COM Component Access Control Rules (Cont.)



Step 3 Enter the following information;

- **Description:** Enter a description of this rule. This description appears in the list view for the module. Optionally, expand the Detailed field to enter a longer description.
- **Enabled:** Use this check box to enable this rule within the module. (It is enabled by default.) By not selecting this check box, 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 an action type from the drop-down list under Take the Following Action.

Step 5 Select one of the following check boxes:

- **Log** Enable this check box 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 Other <Action Type> Rules:** Enable this check box to manipulate rule precedence so that this rule is evaluated before other similar rules. Generally, you should not require this check box. Do not use it without understanding how it works.

Step 6 Under When, select the applications that will be subject to this rule.

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 control. Note that the entry <All Applications> is selected by default. You can use this default, or you can deselect it and create your own application classes.

- **But Not in Any of the Selected Classes:** Optionally, select application classes here to exclude from the application class(es) that you have selected in the included applications field. Note that the entry <None> is selected by default.

Step 7 Under “Attempt to access a COM component,” choose the component sets that will be subject to this rule.

Click the **Insert Component Set** link to select one or more preconfigured COM component sets for this rule. If you do not want to use a COM component set variable, enter a literal PROGID or CLSID (one per line) here. Be sure to use the correct syntax, as shown here. CSA MC provides a utility for extracting PROGID and CLSID information from systems that are running Agent software.

PROGIDs, use the following syntax:

```
Outlook.Application
```

When entering CLSIDs (uppercase hexadecimals) 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.

File Version Control Rules

Cisco.com

- **File version control rules prevent users from running specified versions of applications on their systems.**
- **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.**
- **Users can get around a file version control rule by copying the file in question to a different filename. Therefore, users must work in cooperation for these rule types to be successful.**

© 2005 Cisco Systems, Inc. All rights reserved.

HIPS v2.0—7-30

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 that are not included in this rule to run unimpeded.

One particular example of where this type of rule would be beneficial is in the case described by a Microsoft Security Bulletin (MS01-020) giving details about an Internet Explorer (IE) vulnerability. This bulletin states, “Because HTML e-mails are simply web pages, IE can render them and open binary attachments in a way that is appropriate to their MIME types. However, a flaw exists in the type of processing that is specified for certain unusual MIME types. If an attacker created an HTML e-mail containing an executable attachment, then modified the MIME header information to specify that the attachment was one of the unusual MIME types that IE handles incorrectly, IE would launch the attachment automatically when it renders the e-mail message.”

Microsoft has a patch to correct this security problem, but the patch is available only for IE 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 a security problem that cannot be fixed, and you will want to prevent users from running these versions. The following configuration information uses the IE security bulletin as an example.

Note Users can get around a file version control rule by copying the file in question to a different filename. 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 filename in question.

File Version Control Rules (Cont.)

The screenshot shows the configuration page for a File Version Control rule in the Cisco Management Center. The page title is "File version control [1122]". The breadcrumb navigation is: Configuration > Rule Modules > Windows Rule Modules > Cisco YPN Client Module [V4.5] > Rules > File version control [1122]. The page contains several sections: a status message "No events generated by this rule" with a "View change history" link; a "Description" field with a "Detailed" checkbox; an "Enabled" checkbox which is checked; a "Take the following action" section with a dropdown menu set to "Allow", and checkboxes for "Log" (unchecked) and "Take precedence over other Allow rules" (unchecked); a "when" section with "An execution of the following" and a "File:" input field; and a "with version within these" section with a "Version range:" input field. A tooltip for the "File:" field lists: "- Enter filename (not path)", "- Allowed extensions : exe, dll, ocx", and "- No wildcards allowed". At the bottom, there are "Save" and "Delete" buttons, and a status bar with "13 file changes pending" and "13 operations in progress".

- Step 1** To add rules to your module, click the **Add Rule** link at the bottom of the rule list. A popup 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.
- 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 module. Optionally, expand the Detailed field to enter a longer description.
 - **Enabled:** Use this check box to enable this rule within the policy. (It is enabled by default.) By not selecting this check box, you can save this rule, but it will not be active in the module and it will not be distributed to groups.
- Step 4** Select an action type from the drop-down list under Take the Following Action.
- Step 5** Select one of the following check boxes:
- **Log:** Enable this check box 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 Other <Action Type> Rules:** Enable this check box to manipulate rule precedence so that this rule is evaluated before other similar rules. Generally, you should not require this check box. Do not use it without understanding how it works.
- Step 6** Enter the file that you are prohibiting under An Execution of the Following. (You will enter the exact version in the next field.) This field accepts file entries for .exe, .dll, and .ocx files. Enter only the filename here. No path is required.

For example: `iexplore.exe`

You cannot use wildcard entries in this field.

File Version Control Rules (Cont.)

The screenshot shows the configuration page for a File Version Control Rule in the Cisco Management Center. The rule is currently enabled. The action is set to 'Allow'. There are options for 'Log' and 'Take precedence over other Allow rules'. The 'when' section is set to 'An execution of the following'. The 'File' field is empty, and the 'Version ranges' field is set to '<none>'. A tooltip for the 'File' field provides instructions: 'Enter filename (not path)', 'Allowed extensions: exe, dll, ocx', and 'No wildcards allowed'. The interface includes a 'Save' button and a 'Delete' button.

Step 7 For the file that you entered in the previous field, enter the version or version range (using a dash to indicate range) that you are prohibiting.

For example:

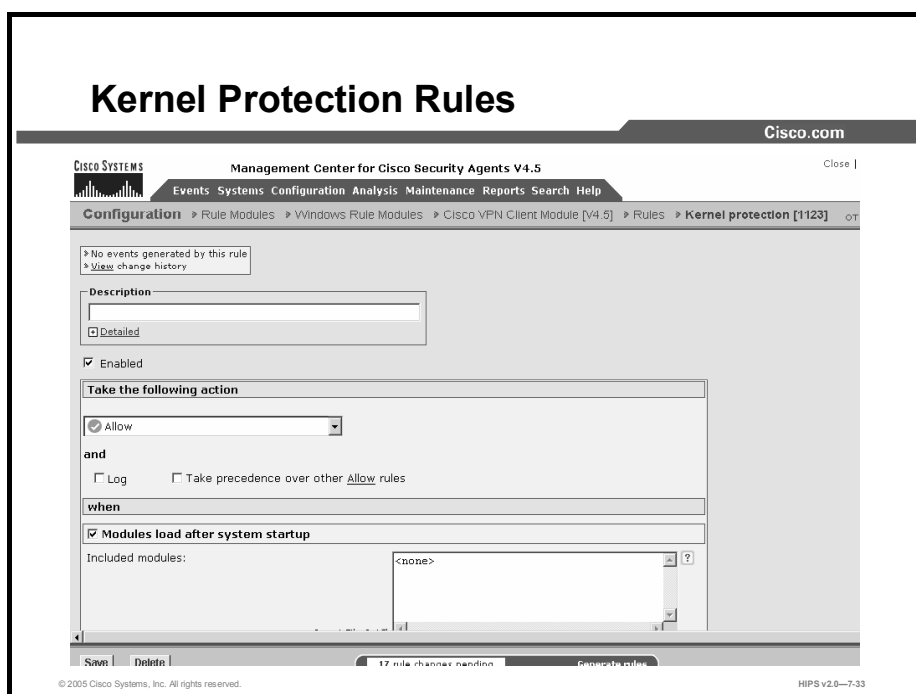
0-5.00.3314.2100 5.00.3314.2100-5.50.4522.1800

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 four values separated by dots.

Note When entering version numbers for Microsoft applications, refer to the Microsoft website. 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.



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 only to detect unauthorized access to the operating system at any time.

These instructions are a continuation of Configuring Rule Modules.

- Step 1** To add rules to your module, click the **Add Rule** link at the bottom of the rule list. A popup 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 module. Optionally, expand the Detailed field to enter a longer description.
 - **Enabled:** Use this check box to enable this rule within the module. (It is enabled by default.) By not selecting this check box, you can save this rule, but it will not be active in the module and it will not be distributed to groups.
- Step 4** Select an action type from the drop-down list under Take the Following Action.

Note Note that High Priority Deny, Allow, Deny, Add Process to Application Class, and Monitor are the only action types available. If you select Add Process to Application Class, the two classes you that are adding a given process to are either Authorized rootkit or Unauthorized rootkit.

Kernel Protection Rules (Cont.)

The screenshot shows the Cisco Management Center for Cisco Security Agents V4.5 interface. The main window is titled "Take the following action" and contains several configuration sections:

- Take the following action:** A dropdown menu is set to "Allow". Below it are checkboxes for "Log" (unchecked) and "Take precedence over other Allow rules" (unchecked).
- when:** A section with two checked checkboxes: "Modules load after system startup" and "Modules modify kernel functionality".
 - Included modules:** A text field containing "<none>" with a help icon (?) to its right. Below it is a link "Insert File Set" and a button "double-click variable to view".
 - Included module hashes:** A text field containing "<all>" with a help icon (?) to its right.
 - Included code patterns:** A text field containing "<all>" with a help icon (?) to its right.
- Note:** "The edit fields in this rule section are maintained by the Event Management Wizard."

At the bottom of the window, there are buttons for "Save" and "Delete", and a status bar indicating "17 rule changes pending" and "Generate rules".

Step 5 Under When, you have the following choices:

- **Modules Load 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.

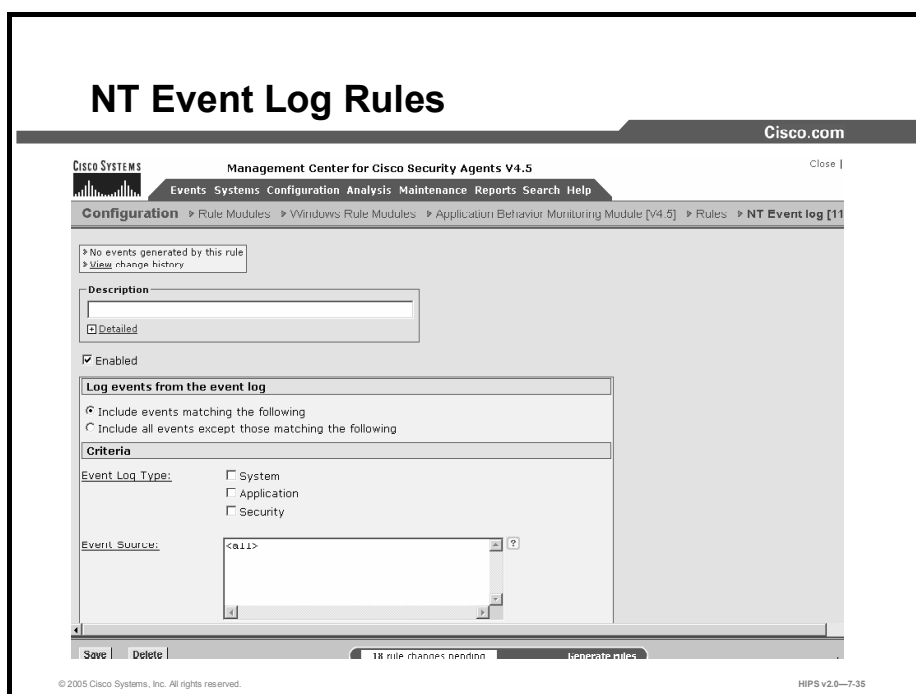
Caution This part of the rule only detects unauthorized access to the operating system and does not prevent it. Upon this detection, you can impose stringent network restrictions by selecting the Restrict Network Connectivity check box.

- **Modules Modify Kernel Functionality:** When unauthorized modules are detected, selecting this check box causes the system in question to log this event. You can use this detection to create dynamic rootkit application classes and to change the system state to a state that enforces a more restrictive policy.

Configure the edit fields in this rule using the wizard. 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. Create exceptions only for actions that 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 (for example, drivers) into this field.
- **Code Patterns to Be Excluded:** By default, this field contains <None>. The wizard enters code patterns (not inside any module) in this field.

Step 6 Click **Save** when finished.



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

These instructions are a continuation of Configuring Rule Modules.

- Step 1** To add rules to your module, click the **Add Rule** link at the bottom of the rule list. A popup 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.
- Step 3** Enter the following information:
 - **Description:** Enter a description of this rule. This description appears in the list view for the module. Optionally, expand the Detailed field to enter a longer description.
 - **Enabled:** Use this check box to enable this rule within the module. (It is enabled by default.) By not selecting this check box, you can save this rule, but it will not be active in the module and it will not be distributed to groups.
- Step 4** Select one of the following choices under Log Events from the Event Log:
 - **Include Events Matching the Following:** Select this radio button to specify the criteria for NT Event Log entries that 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 that 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 that are logged across multiple systems. You can also correlate NT events that are received from virus scanners that are running on Agent systems and quarantine contaminated files.

NT Event Log Rules (Cont.)

The screenshot shows the 'Criteria' configuration window in the Cisco Security Agent Management Center. The window is titled 'Management Center for Cisco Security Agents V4.5'. The 'Criteria' section is active, showing options for Event Log Type (System, Application, Security), Event Source (a text field containing '<all>'), Severity (Information, Warning, Error, Audit Success, Audit Failure), and Event Code (a text field containing '<all>'). The interface includes a 'Save' button and a status bar indicating '18 rule changes pending' and a 'Generate rules' button. Copyright information for Cisco Systems, Inc. is visible at the bottom.

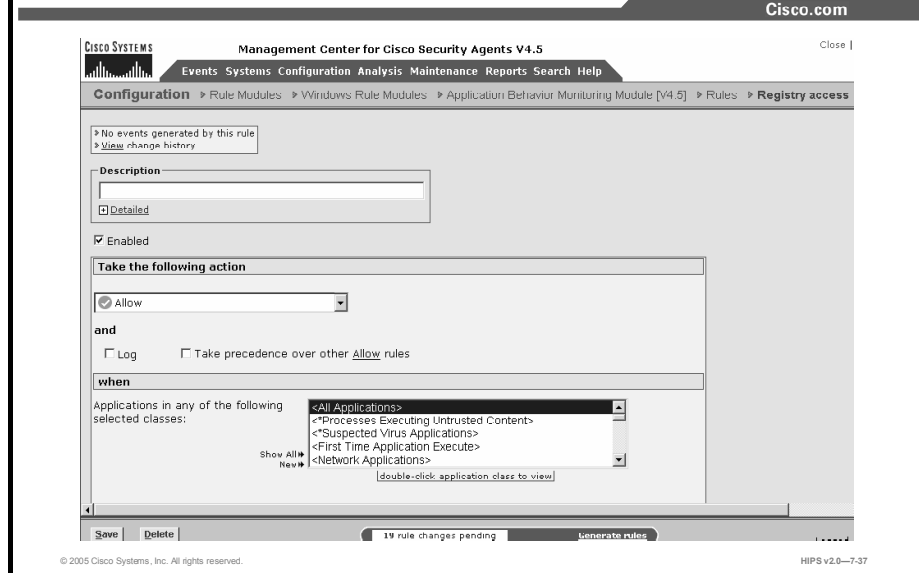
Step 5 You should select at least one criteria for the rule to have any effect.

- **Event Log Type:** Select one or more check boxes here to indicate which NT Event Log entries you want to appear in CSA MC Event Logs (if you choose the first radio button in the preceding step), or which entries you do not want to appear (if you choose the second radio button). The choices are System, Application, and Security.
- **Event Source:** In the text field, enter (one per line) event source parameters that 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 or more check boxes to filter the viewing of events according to severity. If you select no check boxes, all severity levels are included in the rule. The choices are Information, Warning, Error, Audit Success, and Audit Failure.
- **Event Code (Event ID):** In the text field, enter (one per line) event code parameters that 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 website: Q174074, Q299475, and Q301677.

Step 6 Click the **Save** button.

Note To receive messages logged by Norton AntiVirus, select the Application check box and enter **Norton AntiVirus** in the Event Source edit box.

Registry Access Control Rules



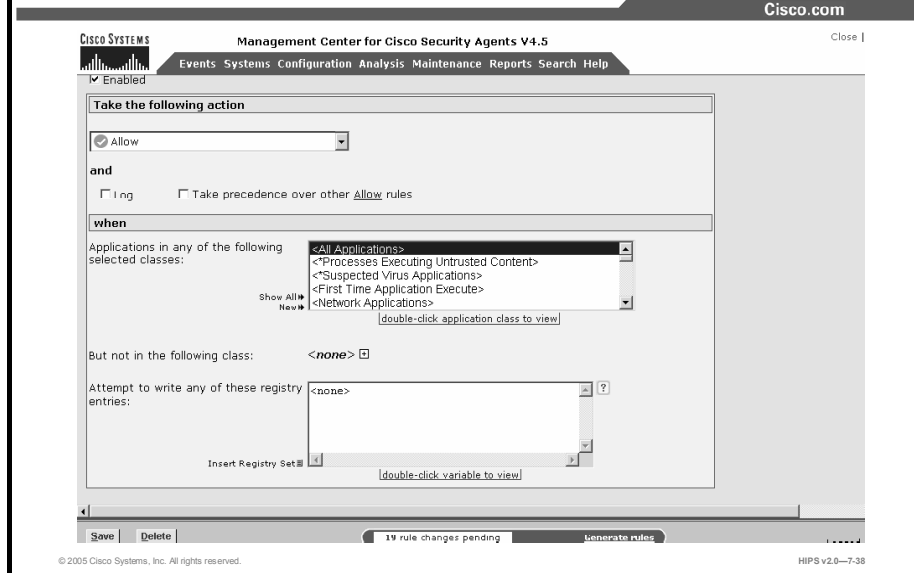
Use registry access control rules to allow or deny writing to specified registry keys by specified applications.

- Step 1** To add rules to your module, click the **Add Rule** link at the bottom of the rule list. A popup 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 module. Optionally, expand the Detailed field to enter a longer description.
 - **Enabled:** Use this check box to enable this rule within the module. (It is enabled by default.) By not selecting this check box, you can save this rule, but it will not be active in the module and it will not be distributed to groups.
- Step 4** Select an action type from the drop-down list under Take the Following Action.
- Step 5** Select one of the following check boxes:
- **Log:** Enable this check box 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 Other <Action Type> Rules:** Enable this check box to manipulate rule precedence so that this rule is evaluated before other similar rules. Generally, you should not require this check box. Do not use it without understanding how it works.

Registry Access Control Rules (Cont.)



Step 6 Under When, select the applications that will be subject to this rule.

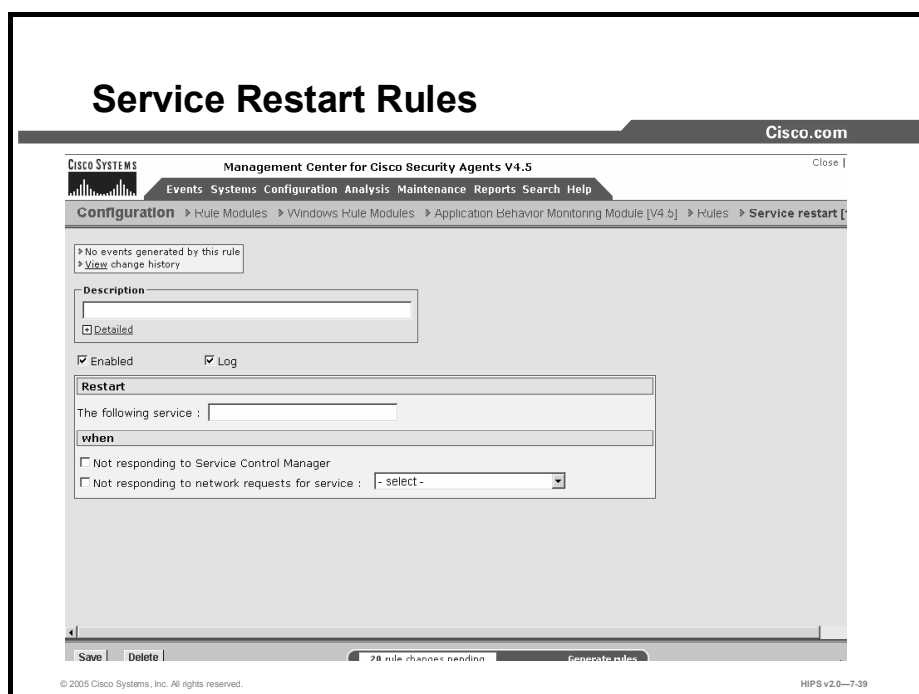
- **Applications in Any of the Selected Classes:** Select one or more preconfigured application classes here to indicate the application(s) whose access to the selected registry keys you want to control. Note that the entry <All Applications> is selected by default. You can use this default, or you can deselect it and create your own application classes.
- **But Not in Any of the Selected Classes:** Optionally, select application classes here to exclude from the application class(es) that you have selected in the included applications field. Note that the entry <None> is selected by default.

Step 7 Click the **Insert Registry Set** link to select one or more preconfigured registry sets for this rule.

Note You cannot enter registry literals here. You must create a registry set variable if you are not using preconfigured 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 rule module. It takes effect when the policy to which it is associated 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.



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

These instructions are a continuation of Configuring Rule Modules.

Step 1 To add rules to your module, click the **Add Rule** link at the bottom of the rule list. A popup 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.

Step 3 Enter the following information:

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

Step 4 Under Restart, enter a service that 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 2003 and 2000: Start > Settings > Control Panel > Administrative Tools > Services "Name" field
- On Windows NT: Start > Settings > Control Panel > Services "Service" field

Step 5 Select one or both of the following check boxes.

- **Not Responding to Service Control Manager:** The Windows Service Control Manager checks the status of system services and recognizes when a service is

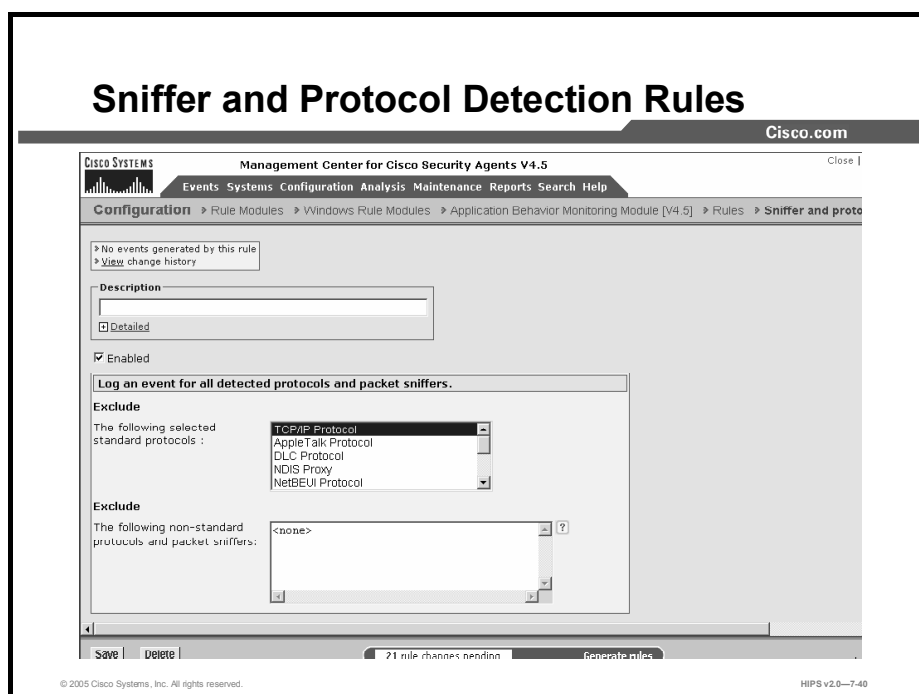
not responding. Selecting this check box causes the Cisco Security Agent to restart the specified service when it does not respond to the Windows Service Control Manager.

- **Not Responding to Network Requests for Service:** Select this check box and then choose a network service (such as HTTP) from the available drop-down 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 that is 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.



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 Internetwork Packet Exchange (IPX), AppleTalk, and NetBIOS Extended User Interface (NetBEUI), are used to provide distributed computing workgroup functions between server and clients or to allow 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 that is 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 that 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 that 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.

Step 1 To add rules to your module, click the **Add Rule** link at the bottom of the rule list. A popup 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.

Step 3 Enter the following information:

- **Description:** Enter a description of this rule. This description appears in the list view for the module. Optionally, expand the Detailed field to enter a longer description.
- **Enabled:** Use this check box to enable this rule within the module. (It is enabled by default.) By not selecting this check box, 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 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 is already excluded.

This is also where you should enter any packet sniffer programs that 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.

UNIX-Only Rules

The rules discussed in this topic are available only for UNIX policies.

Rules Available for UNIX Only

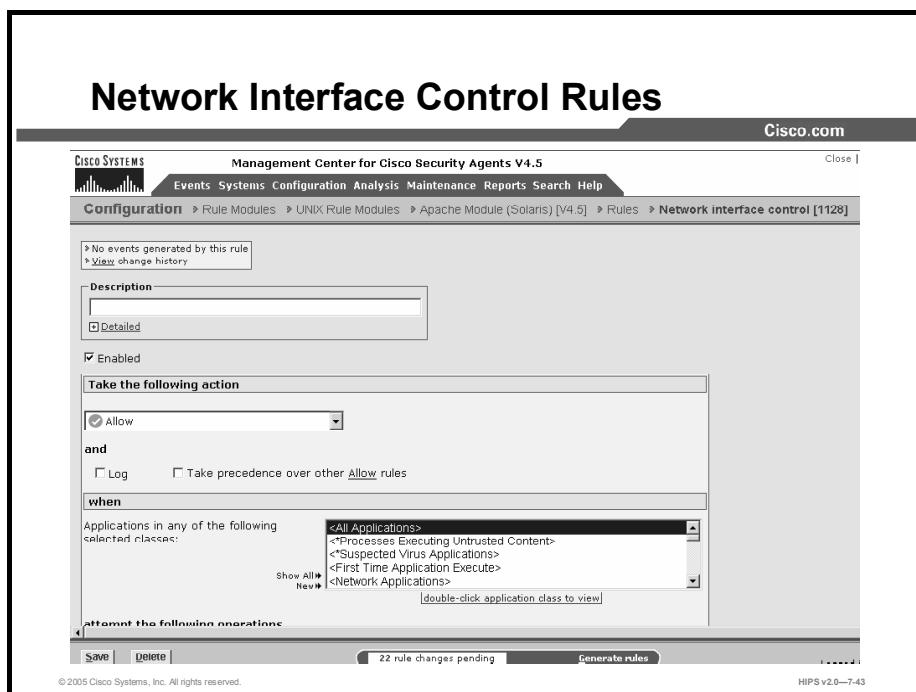
Cisco.com

- **Network interface control**
- **Resource access control**
- **Rootkit/kernel protection**
- **Syslog control**

© 2005 Cisco Systems, Inc. All rights reserved.HIPS v2.0—7-42

The following four rules are available only on UNIX systems:

- Network interface control rules
- Resource access control rules
- Rootkit/kernel protection rules
- Syslog control rules



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 that is 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 Rule Modules.

- Step 1** To add rules to your module, click the **Add Rule** link at the bottom of the rule list. A popup 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.
- Step 3** Enter the following information:
 - **Description:** Enter a description of this rule. This description appears in the list view for the module. Optionally, expand the Detailed field to enter a longer description.
 - **Enabled:** Use this check box to enable this rule within the module. (It is enabled by default.) By not selecting this check box, you can save this rule, but it will not be active in the module and it will not be distributed to groups.
- Step 4** Select an action type from the drop-down list under Take the Following Action.

Network Interface Control Rules (Cont.)

The screenshot displays the configuration page for a Network Interface Control Rule in the Cisco Management Center. The page is titled "Network Interface Control Rules (Cont.)" and is part of the "Management Center for Cisco Security Agents V4.5" interface. The interface includes a navigation bar with options like "Events", "Systems", "Configuration", "Analysis", "Maintenance", "Reports", "Search", and "Help". The main configuration area is divided into several sections: "Description" (with a text input field and a "Detailed" checkbox), "Enabled" (with a checked checkbox), "Take the following action" (with a dropdown menu set to "Allow"), "and" (with "Log" and "Take precedence over other Allow rules" checkboxes), "when" (with a list of application classes including "<All Applications>", "<Processes Executing Untrusted Content>", "<Suspected Virus Applications>", "<First Time Applications Execute>", and "<Network Applications>"), and "attempt the following operations" (with checkboxes for "Open a stream connection to the NIC driver" and "Put the NIC into promiscuous mode"). The interface also features a "Save" button, a "Delete" button, and a status bar indicating "22 rule changes pending".

Step 5 Select one of the following check boxes:

- **Log:** Enable this check box 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 Other <Action Type> Rules:** Enable this check box to manipulate policy rule precedence so that this rule is evaluated before other similar rules. Generally, you should not require this check box. Do not use it without understanding how it works.

Step 6 Under When, select the applications that will be subject to this rule.

Step 7 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) you want to control. Note that the entry <All Applications> is selected by default. You can use this default, or you can deselect it and create your own application classes.

- **But Not in Any of the Selected Classes:** Optionally, select application classes here to exclude from the application class(es) that you have selected in the included applications field. Note that the entry <None> is selected by default.

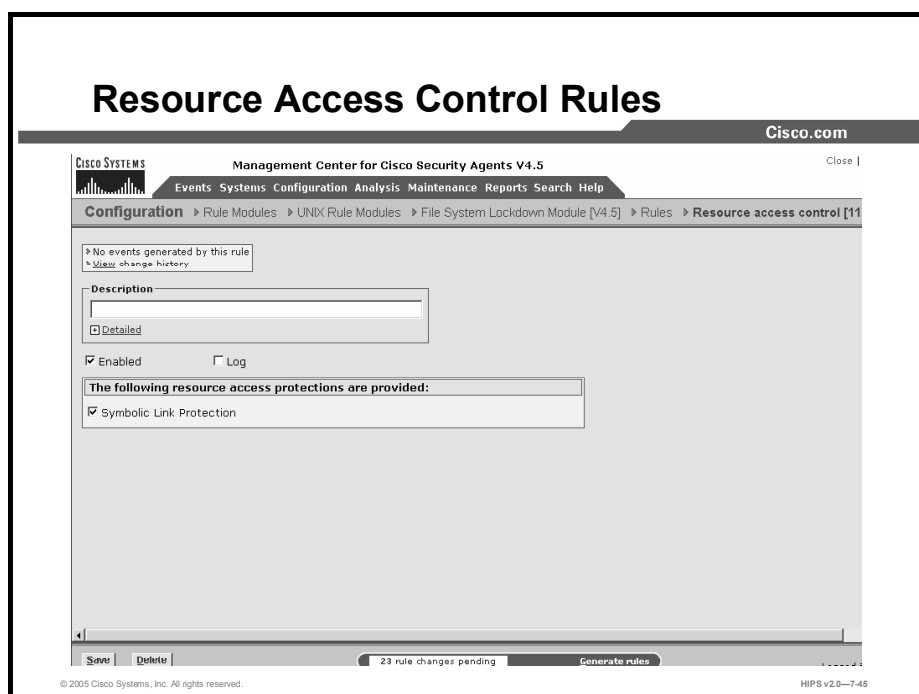
Step 8 Select one or more of the following check boxes under Attempt the Following Operations:

- Open a Stream Connection to the NIC Driver
- Put the NIC into Promiscuous Mode

Note If you select the Allow radio button, when you select to "Put the NIC into promiscuous mode," the "Open a stream connection to the NIC driver" check box 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" check box, the "Put the NIC into promiscuous mode" check box is also automatically selected. If you deny one, the other is automatically denied as well.

Step 9 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).



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 Rule Modules.

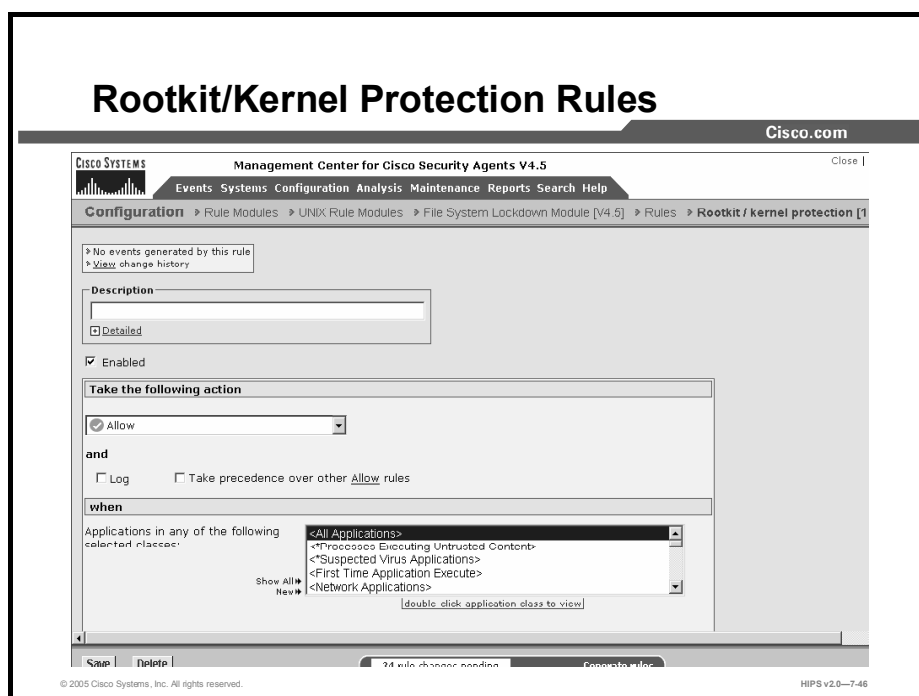
- Step 1** To add rules to your module, click the **Add Rule** link at the bottom of the rule list. A popup list of the available rule types appears.
- Step 2** Select the **Resource Access Control** rule. This takes you to the configuration view for this rule type.
- Step 3** Enter the following information:
- **Description:** Enter a description of this rule. This description appears in the list view for the module. Optionally, expand the Detailed field to enter a longer description.

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

Step 4 Select the **Symbolic Link Protection** check box 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.



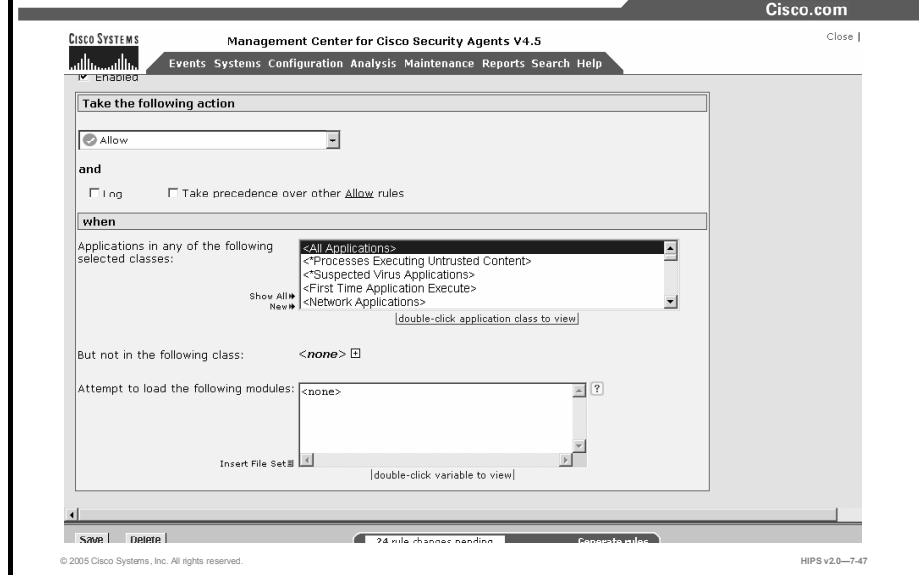
Use the rootkit/kernel protection rule to control unauthorized access to the operating system. In effect, this rule controls drivers that attempt to load dynamically after boot time. You can use this rule to specify authorized drivers that you are allowing to load any time after the system is finished booting.

These instructions are a continuation of Configuring Rule Modules.

- Step 1** To add rules to your module, click the **Add Rule** link at the bottom of the rule list. A popup 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.
- Step 3** Enter the following information:
- **Description:** Enter a description of this rule. This description appears in the list view for the module. Optionally, expand the Detailed field to enter a longer description.
 - **Enabled:** Use this check box to enable this rule within the module. (It is enabled by default.) By not selecting this check box, 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 an action type from the drop-down list under Take the Following Action.
- Step 5** Select one of the following check boxes:
- **Log:** Enable this check box 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 Other <Action Type> Rules:** Enable this check box to manipulate policy rule precedence so that this rule is evaluated before other

similar rules. Generally, you should not require this check box. Do not use it without understanding how it works.

Rootkit/Kernel Protection Rules (Cont.)



Step 6 Under When, select the applications that will be subject to this rule.

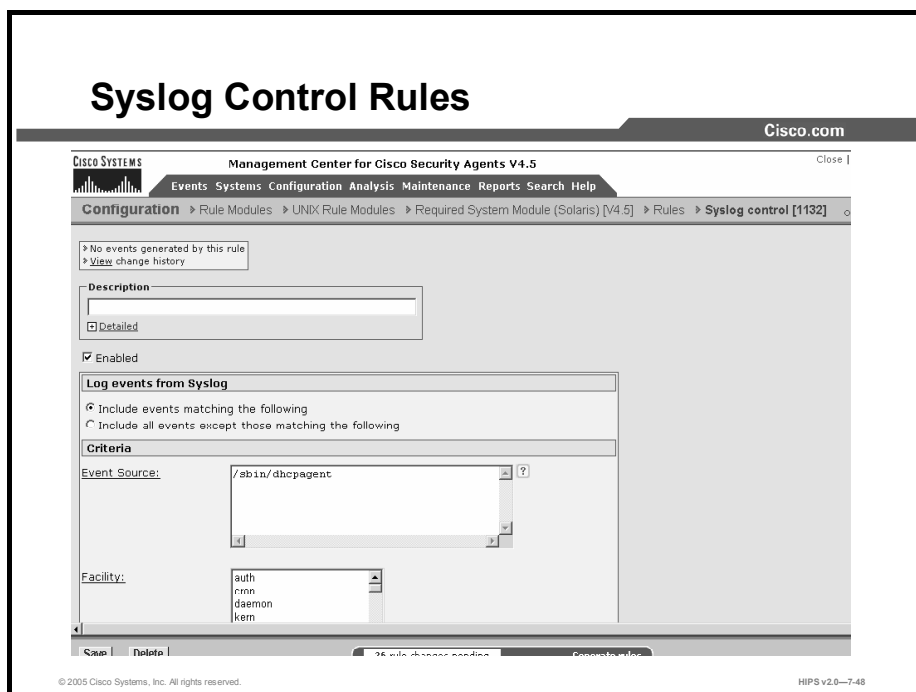
Step 7 Applications in Any of the Selected Classes: Select one or more preconfigured application classes here to indicate the application(s) whose access to the selected resource(s) you want to control. Note that the entry <All Applications> is selected by default. You can use this default, or you can deselect it and create your own application classes.

- **But Not in Any of the Selected Classes:** Optionally, select application classes here to exclude from the application class(es) that you have selected in the included applications field. Note that the entry <None> is selected by default.

Step 8 By default, the Attempt to Load the Following Modules field contains <None>, which indicates all drivers. Enter the names of drivers you want to specify for this the rule and therefore allow, deny, or monitor the loading of at any time.

Caution If you enter file sets that use content-matching constraints, via the Insert File Set link, the content-matching constraints are ignored.

Step 9 Click the **Save** button.



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

These instructions are a continuation of Configuring Rule Modules.

- Step 1** To add rules to your module, click the **Add Rule** link at the bottom of the rule list. A popup 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.
- Step 3** Enter the following information:
- **Description:** Enter a description of this rule. This description appears in the list view for the module. Optionally, expand the Detailed field to enter a longer description.
 - **Enabled:** Use this check box to enable this rule within the module. (It is enabled by default.) By not selecting this check box, you can save this rule, but it will not be active in the module and it will not be distributed to groups.
- Step 4** Select an option under Log Events from Syslog:
- **Include Events Matching the Following:** Select this radio button to specify the criteria for syslog entries that 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 that 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.

Syslog Control Rules (Cont.)

The screenshot shows the 'Criteria' configuration window in the Cisco Management Center. The 'Event Source' field is populated with '/sbin/dhcpagent'. The 'Facility' list box contains 'auth', 'cron', 'daemon', 'kern', and 'local0'. Under 'Priority', the 'Warning' checkbox is checked. The 'Message Pattern' field contains '<a.11>'. At the bottom of the window, there are 'Save' and 'Delete' buttons, a status bar indicating '20 rule changes pending', and a 'Generate rules' button.

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

- **Event Source:** In the text field, enter (one per line) event source parameters that 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 that you want to appear in CSA MC Event Logs (if you choose the first radio button in Step 4), or which entries you do not want to appear (if you choose the second radio button).
- **Priority:** Select one or more check boxes by which to filter the viewing of events according to priority. If you select no check boxes, all priorities are included in the rule.
- **Message Pattern:** In the text field, enter (one per line) message patterns that 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.

Note On Linux platforms, the default syslogd does not embed the facility or priority level in the syslog messages. Using a different syslogd, such as syslog-ng, with correct message formatting, it is possible to use the facility and/or priority levels to report these events. Therefore, if syslog-ng is used, the message template must take the following form:

```
template("%DATE $HOST $PROGRAM: [ID 0 $FACILITY.$LEVEL] $MSG\n")
```

Note For example, the entry for content recorded into /var/log/messages would appear as follows:

```
destination d_1 {file("/var/log/messages" create_dirs(yes)
template("$DATE $HOST $PROGRAM: [ID 0 $FACILITY.$LEVEL] $MSG\n");
};
```

General Syslog Rule Configuration Examples

Consider this example: Configure a syslog rule to log warning messages such as the one listed here:

```
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:

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

- Facility: daemon
- Event Source: /sbin/dhcpagent
- Priority: Warning check box
- Message Pattern: <all>

Here is a second example: Configure a syslog rule to log failed su root attempts such as the one listed here:

```
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 these criteria:

- Facility: auth
- Event Source: su
- Priority: Alert and Above check box
- Message Pattern: root

For a final example: Configure a syslog rule to include all events but exclude all lockstat-related messages such as the one listed here:

```
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 these criteria:

- Facility: kern
- Event Source: <all>

- Priority: all check boxes
- Message Pattern: lockstat

Summary

This topic summarizes the information that you learned in this lesson.

Summary

Cisco.com

- **Rules for both Windows and UNIX:**
 - Agent service control
 - Agent UI control
 - Application control
 - Connection rate limit
 - Data access control
 - File access control
 - Network access control
- **Rules for Windows only:**
 - Clipboard access control
 - COM component access control

© 2005 Cisco Systems, Inc. All rights reserved. HIPS v2.0—7-50

Summary (Cont.)

Cisco.com

- File version control
- Kernel protection
- NT Event Log
- Registry access control
- Service restart
- Sniffer and protocol detection
- **Rules for UNIX only:**
 - Network interface control
 - Resource access control
 - Rootkit/kernel protection
 - Syslog control

© 2005 Cisco Systems, Inc. All rights reserved. HIPS v2.0—7-51

Lesson 8

System Correlation Rules

Overview

This lesson explains the system correlation rules and rule modules and how they can be configured using the Cisco Security Agent Management Center (CSA MC). This lesson contains the following topics:

- Objectives
- System API Control Rule
- Network Shield Rule
- Buffer Overflow Rule
- E-Mail Worm Protection Rule Module
- Installation Applications Policy
- Global Events
- Summary
- Lab Exercise

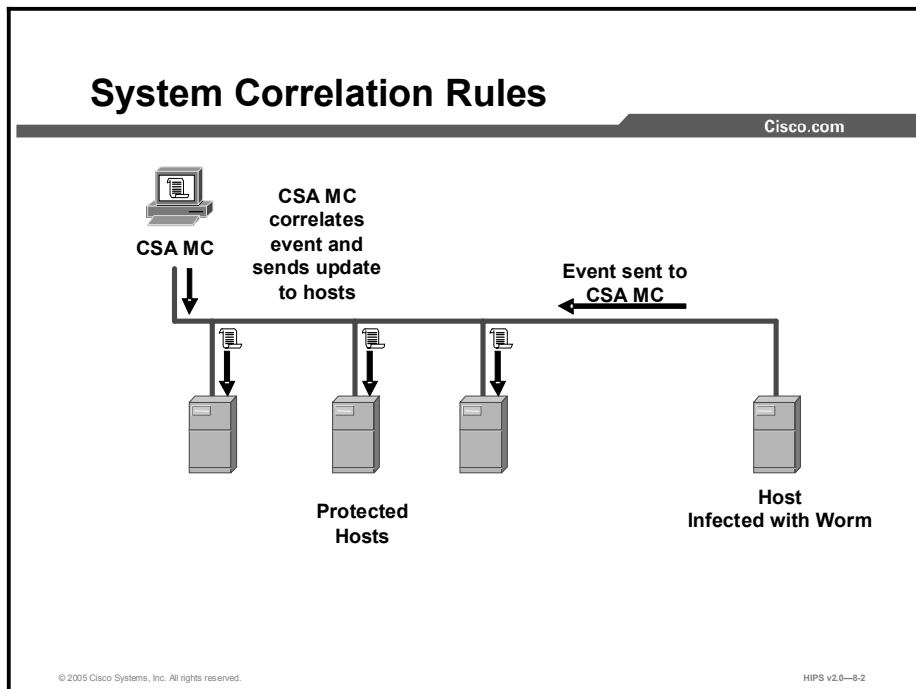
Objectives

Upon completion of this lesson, you will be able to perform the following tasks:

- Define and understand individual rules that when added to your policies allow CSA MC to categorize processes and correlate events across multiple systems
- Understand event correlation and heuristics
- Understand and configure the system API control rule
- Understand and configure the network shield rule
- Understand and configure the buffer overflow control rule
- Understand and configure the e-mail worm protection rule module
- Understand and configure the Installation Applications policy
- Understand and configure global event correlation

About System Correlation Rules

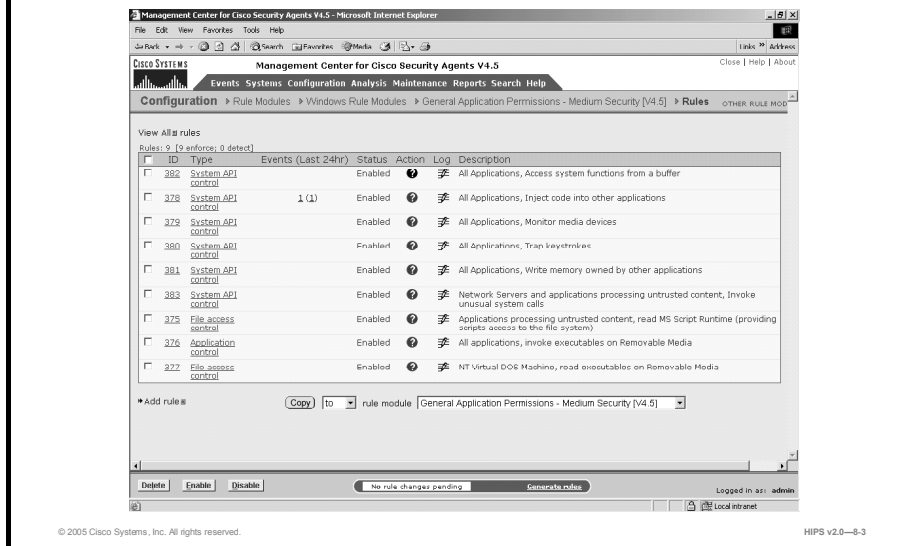
This topic describes system correlation rules.



Management Center for Cisco Security Agents provides rule modules and individual rules that when added to your policies allow CSA MC to categorize processes and correlate events across multiple systems. When these rules are triggered by one or more system actions across a network, the MC registers this occurrence and automatically builds application classes and sends out new process categories to Cisco Security Agents. In some cases, the MC can prevent actions from executing on any additional systems. The Cisco Security Agent also uses heuristics to detect and terminate suspicious activities on systems, such as buffer overflows and password stealing attempts. Use the rules and rule modules described in this lesson just as you would other rules to protect systems.

Event Correlation and Heuristics

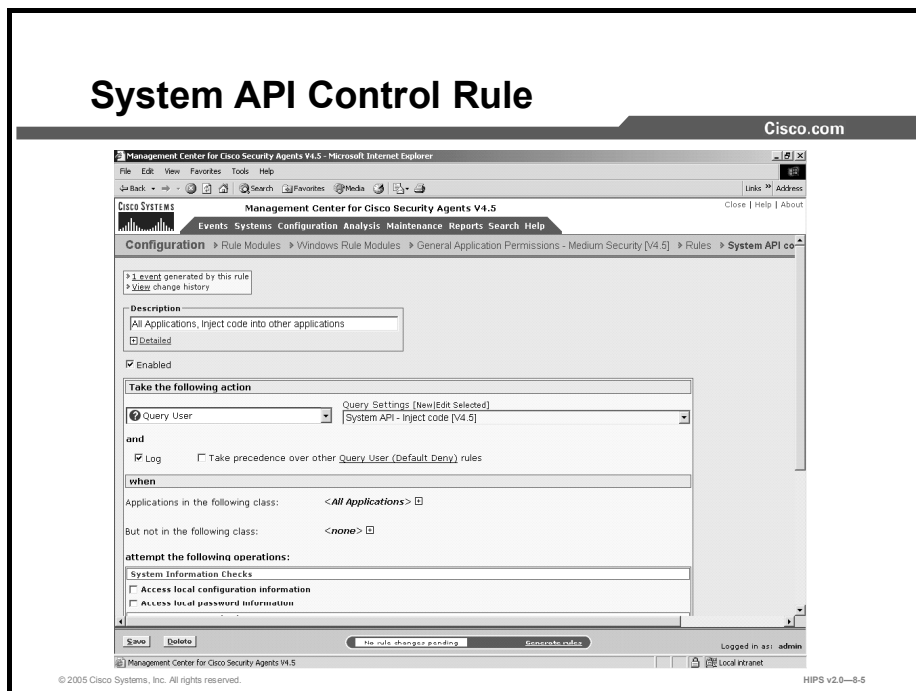
Cisco.com



The network shield rule, which controls SYN flood protection and port scan detection, and the system API control rule are some examples of preconfigured rules that you can add to your modules in the same way that 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 General Applications Permissions module shown in the figure.

System API Control Rule

This topic discusses the uses of the system API control rule.



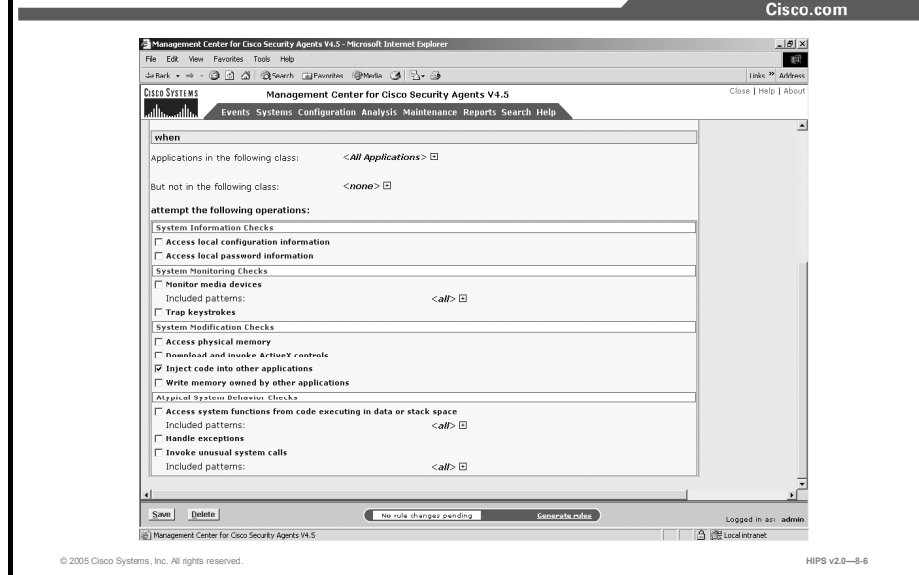
The system API control rule detects several forms of malicious programming code that is installed on a system by an unsuspecting user. Users may think that they are running some other type of program, or they may inadvertently install it as a result of some other activity such as reading an attachment to an e-mail message. Once installed, these malicious programs (for example, Trojans) may allow others to access and virtually take over a system across the network. Other errant programs 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.

Using a service restart rule in conjunction with a system API control rule could be useful, especially in the case of server systems. This way, if you are forced to click the Terminate button when queried by a triggered rule, and you subsequently terminate the application in question, a service restart rule will cause the application to restart automatically.

Use the system API control rule in a policy to detect and prevent errant programs from performing malicious acts on individual systems and networks.

System API Control Rule (Cont.)



The system API control rule lets you enable several different types of system detection.

■ System Information Checks

- Access Local Configuration Information: Detect applications that attempt to read system registry settings.
- Access Local Password Information: Detect applications that attempt to steal local system passwords.

■ System Monitoring Checks

- Trap Keystrokes: Detect applications that attempt to capture system keystrokes.
- Monitor Media Devices: This check box lets you control which applications can monitor media devices on the system. Media device “inputs” can be exploited by Trojans that can, for example, turn on the microphone on a system and covertly listen to a conversation.
 - Patterns to be included: Use the wizard from the Event Log message in question to include particular devices in a system API allow rule. You must specify media devices as “device\port.” For example, plantronics\microphone.

Note Monitor Media Devices is not supported on Windows NT systems. It is also not supported for parallel port media devices on any operating system.

■ System Modification Checks

- Access Physical Memory: Detect applications that attempt to access physical memory directly while bypassing virtual memory restrictions.
- Download and Invoke ActiveX Controls: Detect applications that download ActiveX controls and immediately attempt to execute them. 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.

Note The preceding rule may be unnecessary if system web browser settings are configured with a "High" security level that would restrict the downloading of ActiveX controls.

- Inject Code into Other Applications: Detect applications that are attempting to write code to space owned by other applications, for example, injecting a malicious .dll into a privileged process.
- Write 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.
- Atypical System Behavior Checks
 - Access System Functions from Code Executing in Data or Stack Space: Although this behavior is sometimes exhibited by downloaded or executable content (for example, license-checking software), this may be symptomatic of a buffer overflow attack.
 - Patterns to be included: Use the wizard from the Event Log message in question to include a particular pattern in a system API allow rule when you are seeing buffer overflow events that you believe are harmless.
 - Handle Exceptions: Detect processes running exception-handling routines. This typically occurs due to bugs in the application software. But this may be a sign of an attack if this occurs with an application that does not generally exhibit this behavior.
 - Invoke Unusual System Calls: Use this check box to detect processes invoking system calls that are rarely used. In normal system operation, many system calls are either never used or may only be used infrequently by a specific system application that is performing a service. Attempting to exploit undetected flaws in these unusual system calls is a common attack vector for malware.
 - Patterns to be included: Use the wizard from the Event Log message in question to include a particular module in a system API allow rule when you are seeing events that you believe are harmless.

You also have the ability to select specific application classes to exclude from the various system API control rules that you designate. For example, in some cases, debuggers may perform actions that can be misconstrued as malicious behavior. Therefore, you would want to create an application class and select it as an exclusion to one or more system API control rule features.

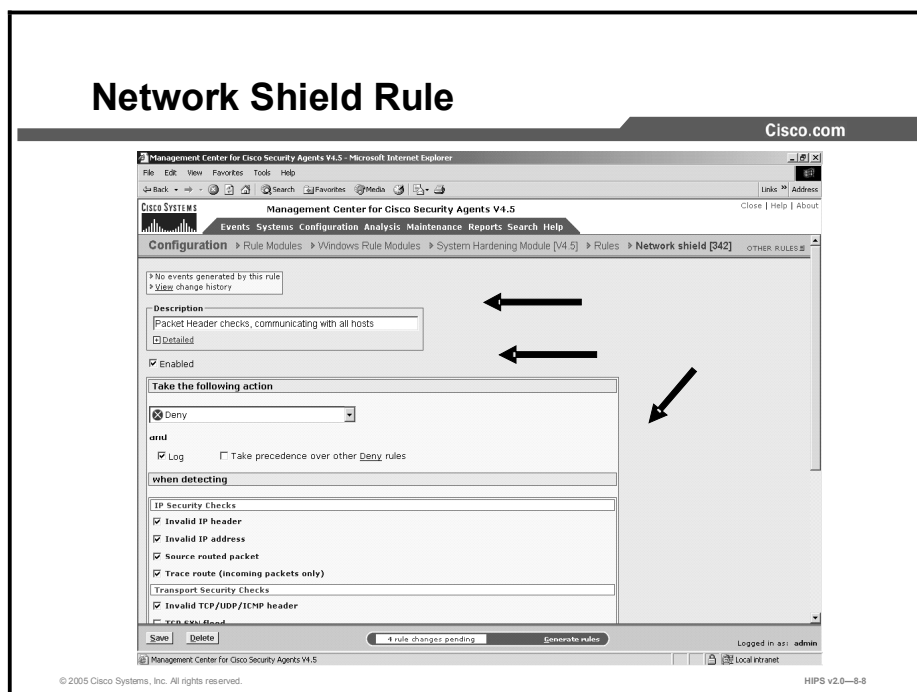
Replicate Feature

When you make rule changes and click the Save button for rule types that contain multiple check boxes, such as system API control rules, a Replicate link appears beside the "Saved changes" message at the top of the rule page. (Network shield and buffer overflow rules also provide this feature.) Click on **Replicate** to access a popup box. From this box, you select other policies that contain system API control rules and choose to propagate the same change(s) you made on the current page to system API control rule pages in other policies. If the change you make to one system API control rule page is a change you need to make to all system API

control rules in all your policies, this is a quick way to propagate those changes on a wide or even global scale.

Network Shield Rule

The network shield rule provides network protocol stack hardening capabilities. This topic describes the network shield rule and how to configure it.



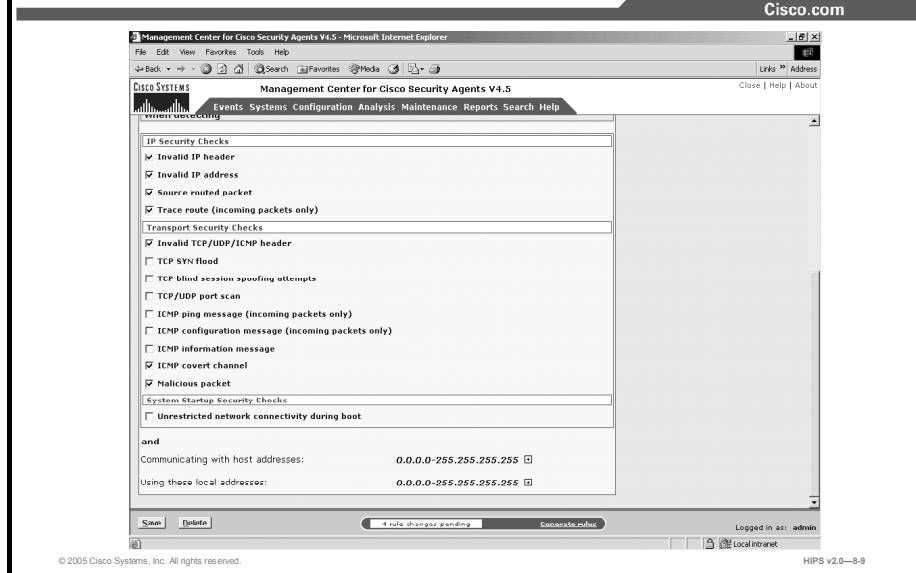
- Step 1** In the network shield rule configuration view, enter a description of this rule. This description appears in the list view for the module. Optionally, expand the Detailed field to enter a longer description.
- Step 2** Select the **Enabled** check box. Use this check box to enable or disable this rule within the module. (It is enabled by default.) By not selecting this check box, you can save this rule, but it will not be active in the module and it will not be distributed to groups.
- Step 3** Select an action from the Take the Following Action drop-down menu.

Note You may choose to add the system process to the Processes Communicating with Untrusted Hosts application class, which causes the remote host IP address to be sent to the MC for global correlation. This may result in the address being added to the @dynamic address list for quarantining. Because IP addresses can be spoofed, using this capability for this rule type is not recommended. It is more applicable for NACL-based rules, when you are sure that you are communicating with the address. (such as an established TCP connection).

- Step 4** Select the Log check box 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.
- Step 5** Select the **Take Precedence Over Other <Action Type> Rules** check box to manipulate rule precedence so that this rule is evaluated before other similar rules.

Generally, you should not require this check box. Do not use it without understanding how it works.

Network Shield Rule (Cont.)



The following steps describe each check box available in the When Detecting configuration section of the network shield rule.

Caution You cannot use network shield rules in policies that have user state conditions set. If you attempt to attach a rule module that contains a network shield rule to a policy that has user state settings, you will be notified of a configuration error.

IP Security Checks

- Step 6** (Optionally) Select the Invalid IP Header check box to cause 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. If you configure this as a deny rule, the following occurs: If any of these checks fail, the packet is dropped; 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.)
- Step 7** (Optionally) Select the Invalid IP Address check box to have CSA determine invalid IP addresses. IP addresses are determined to be invalid for these reasons: if the source address is a multicast address, if the TCP connection is to a broadcast address. You can select this check box as part of a deny rule to protect against these types of attacks.
- Step 8** (Optionally) Select the Source Routed Packet check box to detect IP options that 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.
- Step 9** (Optionally) Select the Trace Route check box to detect the mapping of network topology via trace route.

Transport Security Checks

- Step 10** (Optionally) Select the Invalid TCP/UDP/ICMP Header check box 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 includes verifying that certain fields have valid values and that certain combinations of TCP flags are legal. This defeats attacks such as a Christmas Tree scan.
- Step 11** (Optionally) Select the TCP SYN Floods check box to identify SYN flooding. 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 (such as a nonexistent 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. Detecting and preventing SYN floods stops this attack from succeeding.

Note The preceding rule type is not available for UNIX policies, as the UNIX OS already provides this protection.

Note If you enable the TCP SYN Floods check box, you cannot enter address restrictions in the address field for this rule. You must use all addresses.

- Step 12** (Optionally) Select the **TCP Blind Session Spoofing Attempts** check box. If you configure this as a deny rule, this causes agents to make TCP sequence numbers unpredictable. A server that accepts 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.

Note The preceding rule type is not available for UNIX policies, as the UNIX OS already provides this protection.

Note If you make any changes to the TCP Blind Session Spoofing Attempts feature, these changes are not enforced until after the Agent system(s) is rebooted.

- Step 13** (Optionally) Select the TCP/UDP Port Scan check box. 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. Configure this rule to log an event 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. In most cases, you should apply port scan detection to servers and end-user systems in your enterprise. Configure this rule as a deny rule to prevent unauthorized port scans, effectively cloaking a system on the network. Denying port scans causes a system not to respond to connectivity tests and not to respond to service requests with connectivity error messages. A system generally sends out error messages when a remote

machine sends a request for a service that 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. If you are running an allowed service on a system and you are denying port scans, connection requests to this service are honored and your machine is viewable for the service that you are offering.

Note If you select the Correlate Network Scans check box in the Global Event Correlation page (covered later in the lesson), when scans are detected and denied across several machines, CSA MC correlates these events and generates an additional event to warn of this correlation. This correlation only occurs when deny rules are triggered.

- Step 14** (Optionally) Select the ICMP Ping Message check box. This option works like the TCP/UDP port scan feature, but for ping scans.
- Step 15** (Optionally) Select the ICMP Configuration Message check box. If you configure this rule as a deny, this feature restricts messages that can change the configuration of a machine. For example, a redirect can be used to cause routing tables to be updated.
- Step 16** (Optionally) Select the ICMP Information Message check box. Some ICMP messages may be used to gather information about a machine in an attempt to attack it. This data, when obtained, can be used to gather system information that can be used to exploit the system. If you configure this rule as a deny, this feature restricts messages that report back on system or network configuration.
- Step 17** (Optionally) Select the ICMP Covert Channel check box. Configuring this rule as a deny causes Agents to drop unsolicited echo responses. The Cisco Security Agent validates that the echo response data matches the echo request data. This way, ping cannot be used as a transport for communications.
- Step 18** (Optionally) Select the Malicious Packet check box. Configuring this rule as a deny causes Agents to block packets that are technically legal but are known exploits against protocol stacks (for example, UDP packet storm or RF poison).

System Startup Security Checks

- Step 19** (Optionally) Select the Unrestricted Network Connectivity During Boot check box. Configuring this feature as a deny prevents nonessential network connections during system startup. This check is automatically disabled when the Agent service starts and policies (including those that govern allowed network connections) are enforced. This protects the system from network-based attacks at boot-time before the Agent service has started.

Note The preceding rule type is not available for UNIX policies.

Note If you select the Unrestricted Network Connectivity During Boot check box, you cannot enter address restrictions in the address field for this rule. You must use all addresses.

Note You cannot use a rule that has the Unrestricted Network Connectivity During Boot check box selected in policies with rule modules that have system and/or user state conditions set.

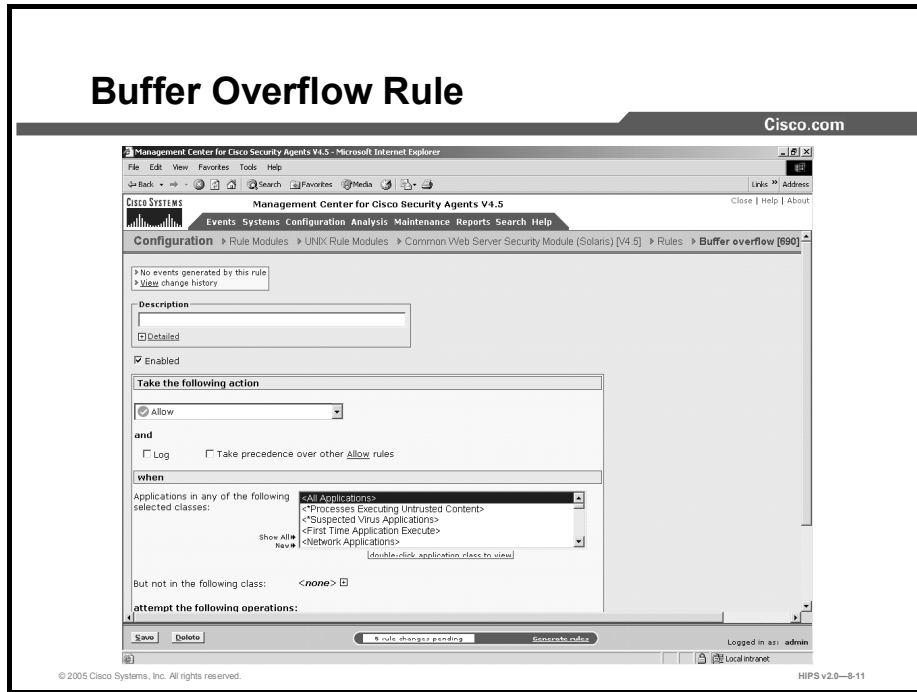
Step 20 (Optionally) Enter specific addresses for those check box features in this rule that support addressing parameters in the And Communicating with Host Addresses section.

Step 21 By default, the Using These Local Addresses field indicates all local addresses on the Agent system. You would want to use this to identify specific network interfaces, if necessary.

Step 22 Click the **Save** button when finished.

Buffer Overflow Rule

This topic describes and explains how to configure the buffer overflow rule on the CSA MC.



A buffer overflow is what happens when two conditions are met: First, an application is coded in such a way that it trusts that all users of that application will provide the application with reasonable and expected data. Second, the application is provided larger quantities of data than it is capable of handling correctly. When these events come together, an application can behave in unexpected and unintended ways.

For applications with special privileges, this can result in external users gaining access to machine resources and privileges that 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 undesirable parties gaining access to your system.

In the case of UNIX operating systems, three distinct types of buffer overruns can occur, based upon the type of memory space involved: stack, data, and heap.

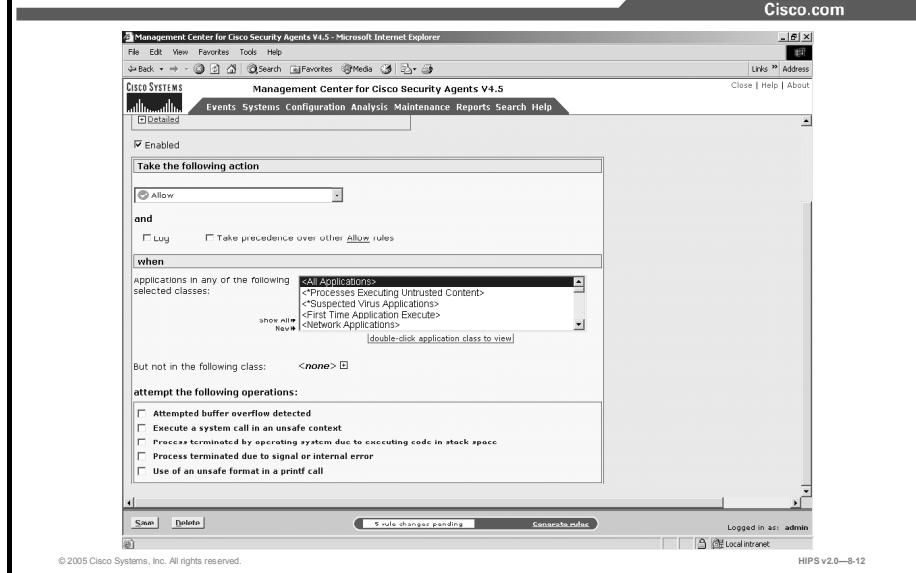
- Stack space is used to store data and information that is local to the piece of code currently being executed in an application. It contains stored control flow information for the application.
- Data space is used to store data with fixed sizes that needs to be shared among different parts of an application. 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 data sets, and is frequently visible to numerous subcomponents of an application.

Note The buffer overflow rule is UNIX-specific. Some corresponding Windows functionality is available from the system API control rule page.

Configure the buffer overflow rule as follows:

- Step 1** To add rules to your policy, click the **Add Rule** link at the bottom of the rule list. A popup 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.

Buffer Overflow Rule (Cont.)



- Step 1** Enter a description in the Description field for this rule. This description appears in the list view for the module. Optionally, expand the Detailed field to enter a longer description.
- Step 2** Select the Enabled check box to enable this rule within the module. (It is enabled by default.) By not selecting this check box, you can save this rule, but it will not be active in the module and it will not be distributed to groups.
- Step 3** Select an action type from the Take The Following Action drop-down list.
- Step 4** Select the Log check box 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.
- Step 5** Select the Take Precedence Over Other <Action Type> Rules to manipulate rule precedence so that this rule is evaluated before other similar rules. Generally, you should not require this check box. Do not use it without understanding how it works.
- Step 6** Select one or more preconfigured application classes from the Applications in Any of the Following Selected Classes list box.

Note The entry <All Applications> is selected by default. You can use this default, or you can deselect it and create your own application classes.

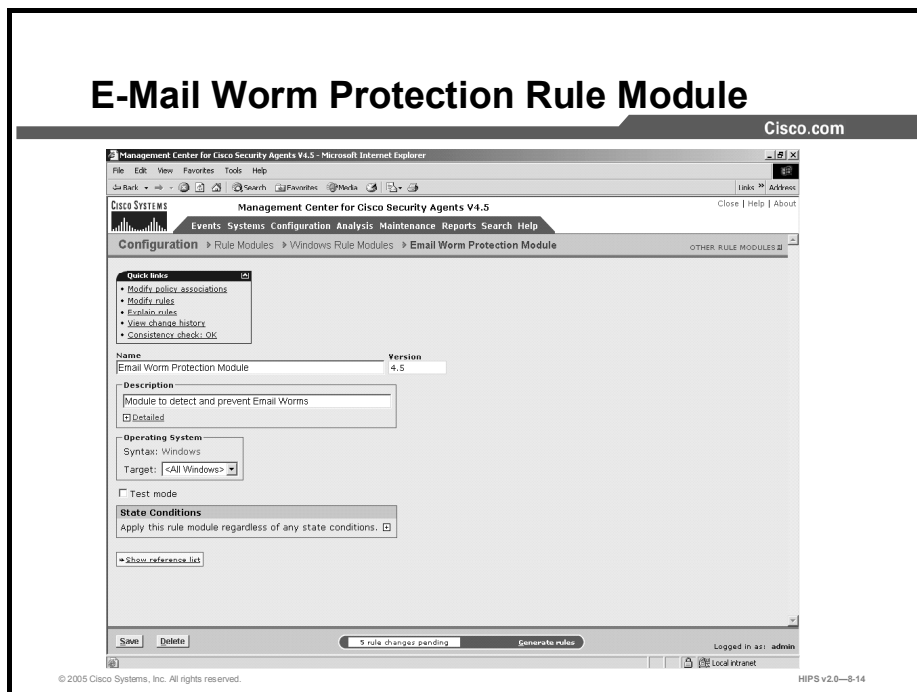
- Step 7** Optionally, select application classes from the But Not in the Following Class link to exclude from the application class(es) that you have selected in the included applications field.

Note The entry <None> is selected by default.

- Step 8** Select the Attempted Buffer Overflow Detected check box to detect buffer overflow conditions that occur in UNIX executables. This feature provides protection from stack buffer overflows for a number of commonly used libc routines. As a large number of attacks on UNIX systems are based upon buffer overflow attacks, enabling this feature is recommended.
- Step 9** Select the Execute a System Call in an Unsafe Context check box to prevent certain system calls (such as those that 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.
- Step 10** Select the Process Terminated by Operating System Due to Executing Code in Stack Space check box to enable the "noexec_user_stack" system variable for all processes or for processes added to the <*Processes requiring OS Stack Execution Protection>. This check box monitors the execution of instructions from stack memory. This only provides logging.
- Step 11** Select the Process Terminated Due to Signal or Internal Error check box to cause the Agent to monitor when processes are killed on a system by either another process or an internal error occurring on the system. The only action type available when this check box is enabled is Monitor.
- Step 12** Select the Use of an Unsafe Format in a printf Call check box to prevent use of the '%n' *printf() format qualifier. Numerous attacks use 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 that you designate. If you select an application in the available list beside a check box 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.

E-Mail Worm Protection Rule Module

This topic describes and explains how to configure the e-mail worm protection rule module using the CSA MC.



E-mail worms are some of the most commonly spread and costly attacks affecting corporate networks today. Worms easily infect systems, passing undetected through most security software until virus scanner vendors provide 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.

When a worm of this type is received through e-mail and executed by unsuspecting users, it generally attempts to send copies of itself to all entries in the e-mail address book of the user. In doing this, the worm modifies registry keys, writes its own script files, and modifies existing files. This not only makes file recovery difficult but can also cause users to invoke the virus again when they attempt to open these infected files.

The Cisco Security Agent ships with a preconfigured e-mail worm protection rule module. You must have this module deployed to take advantage of e-mail worm network event correlation and quarantine capabilities.

The e-mail worm protection module works through a combination of steps, including dynamically building an application class through the detection of a suspicious action occurring on a system. If this suspicious action detection is seen by the CSA MC as occurring on more than one system, a quarantine of the detected malicious process will also occur. More specifically, the detection and tagging of a virus or e-mail worm occurs through two sets of rule types. In fact, these rules can be used more widely to identify and stop any type of virus, not only e-mail; however, this does require some different parameters to be set in the first group of rules.

The e-mail worm protection rule module works this way:

- The first set of rules are written to deny or terminate processes or to query the user when a set of actions are attempted. Those actions are something along the lines of “a process that downloaded content over the network is now attempting to access an e-mail COM component, such as the address book.” This action is suspicious. It is either denied or terminated, or the user is queried about it.
- If the action is denied or terminated (automatically or by the user), the second set of rules tags the offending process and adds the process to the dynamically built “Suspected Virus Applications” class. Once a process is in this class, other rules prevent all processes that are dynamically added to this class from accessing any resources on a system. If these processes are seen on more than one system, it also quarantines the processes in question.

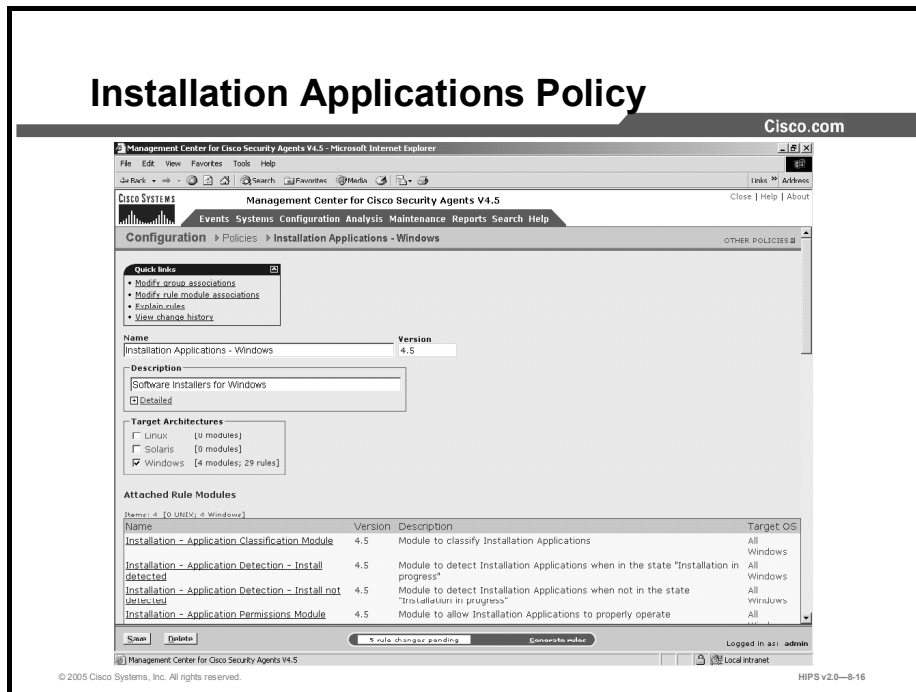
The methodology used in the e-mail protection module can be applied to any virus type that you are protecting against. By altering the parameters of the first rule set (in this example, downloaded content accessing the COM component for the e-mail application address book), you can configure parameters to categorize any process as suspicious and subsequently stop any type of errant action.

E-Mail Worm Event Correlation

If you select one of the options to add dynamically quarantined files to the list in the Global Event Correlation page, 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 filename that the worm was written into. If at least two Agents report worms writing to the same filename 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.

Installation Applications Policy

This topic describes and explains how to configure the Installation Applications policy using the CSA MC.

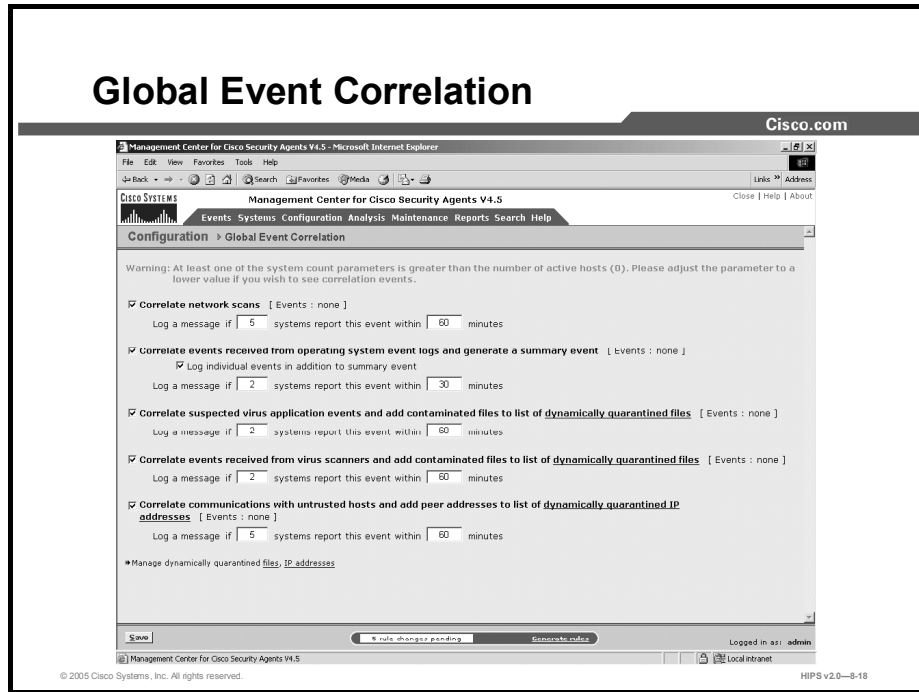


You can apply a preconfigured policy to systems to detect when a software installation occurs and to add the install process or processes to a dynamically built application class. If dynamic tagging occurs, another set of rules would apply to the install processes added to this dynamic class. You may need to use this installation detection rule module in order to enforce a less strict set of rules on the system while an approved software installation is occurring. Under normal conditions, other rules on the system may prevent the install from occurring.

This built-in dynamic application class is called Installation Applications. The rule module may build this application class under the following circumstances: A set of rules determines that setup.exe is detected on a system and it is added to the dynamic built-in Installation Applications class. As a result, a System State installation condition is triggered and a new policy is applied to the system. The system should automatically return to its original policy when the install exits. If this does not occur, the user can manually indicate when the installation is complete and return the system to the initial stricter policy. The installation state may also time out, and the system then automatically returns to its initial policy.

Global Events

This topic describes and explains 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 Configuration > Global Event Correlation on the menu bar, provides the following capabilities:

- **Correlate Network Scans:** With this check box 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. The threshold and time frame for correlating network scans are values you can configure.

Note You must have a network shield rule with port scan detection and ping scan enabled in a policy deployed to the Agent(s) in question for these event types to be detected and logged.

- **Correlate Events Received from Operating System Event Logs and Generate a Summary Event:** With the Log Individual Events in Addition to Summary Event check box 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 two systems log the same NT event type within 30 minutes, a correlated summary event is logged. If you do not enable these check boxes, NT event correlation does not take place, but individual NT events are logged in accordance with the NT Event Log rule that you have configured.

Note 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.

Note In this case, there is an additional check box (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 check box, but you do enable the Correlate Events check box, 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 Suspected Virus Application Events and Add Contaminated Files to List of Dynamically Quarantined Files:** With this check box enabled, when processes are added to the dynamic <Suspected Virus Applications> application class and this event is logged across multiple Agent systems, these events are correlated and the contaminated file that triggered the event is added to a dynamic list of quarantined files that CSA MC maintains. If you have a rule configured to stop dynamically quarantined files in a deployed policy, no further Agents can access the contaminated file. If you do not enable this check box, suspected virus correlation does not take place, but individual virus events are logged.

Note You must have a corresponding policy deployed to the Agent(s) in question for these event types to be detected and logged.

- **Correlate Events Received from Virus Scanners and Add Contaminated Files to List of Dynamically Quarantined Files:** With this check box enabled, 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 the list of quarantined files. If you have a rule configured to stop access to dynamically quarantined files in a deployed policy, no further Agents can receive the contaminated file.

Note This feature works with Norton, McAfee, and Trend 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 the name of the antivirus software in the Event Source field. The threshold and time frame for correlating events received from virus scanners are values that you can configure.

Note To view the files that are added to the dynamically quarantined files list, click the numbered link beside dynamically quarantined files. It takes you to the pertinent Event Log messages. Read the messages there to locate the names of quarantined files. You can also click the Manage Dynamically Quarantined Files link at the bottom of the page.

- **Correlate Communications with Untrusted Hosts and Add Peer Addresses to List of Dynamically Quarantined IP Addresses:** With this check box enabled, when processes are added to the dynamic <Processes communicating with Untrusted Hosts> application class and this event is logged across multiple Agent systems, these events are correlated and the untrusted peer address that triggered the event is added to a dynamic list of

quarantined IP addresses that CSA MC maintains. If you have a rule configured to stop dynamically quarantined IP addresses in a deployed policy, no further Agents can communicate with this peer address. If you do not enable this check box, untrusted host correlation does not take place, but individual untrusted host events are logged.

Note You must have a corresponding policy deployed to the Agent(s) in question for these event types to be detected and logged.

Note To view the IP addresses that are added to the dynamically quarantined addresses list, click the numbered link beside dynamically quarantined IP addresses. This takes you to the pertinent Event Log messages. Read the messages there to locate the quarantined IP addresses. You can also click the Manage Dynamically Quarantined IP Addresses link at the bottom of the page.

Manage Dynamically Quarantined Files and IP Addresses

You can use the @dynamic token in the File Set text field and in the Network Address Set text field to control access to files and addresses that have been quarantined by CSA MC. Files are quarantined as a result of suspected virus application events, correlated virus scanner log messages, or files that were added manually. This list updates automatically (dynamically) as logged quarantined files are received. Addresses are quarantined as a result of communication with a suspected untrusted host (this updates dynamically) or by being added manually.

To view the files that are added to the dynamically quarantined files list and to manually add files to be quarantined, click the **Manage Dynamically Quarantined Files** link on the Global Event Correlation page. Add and remove files from this list using the provided buttons on the bottom of the window that appears.

To view the addresses that are added to the dynamically quarantined IP addresses list and to manually add addresses to be quarantined, click the **Manage Dynamically Quarantined IP Addresses** link on the Global Event Correlation page. Add and remove IP addresses from this list using the provided buttons on the bottom of the window that appears. The “Source” column in this window describes how the address was added to the list (manually by the administrator or through a correlation event).

Summary

This topic summarizes what you have learned in this lesson.

Summary

Cisco.com

- **Management Center for Cisco Security Agents provides rule modules and individual rules that when added to your policies allow CSA MC to categorize processes and correlate events across multiple systems.**
- **The system API control rule detects several forms of malicious programming code that is installed on a system by an unsuspecting user.**
- **The network shield rule provides network protocol stack hardening capabilities.**
- **A buffer overflow occurs when two conditions are met: First, an application is coded in such a way that it trusts that all users of that application will provide the application with reasonable and expected data. Second, the application is provided larger quantities of data than it is capable of correctly handling. When these events come together, an application can behave in unexpected and unintended ways.**

© 2005 Cisco Systems, Inc. All rights reserved. HIPS v2.0—5-19

Summary (Cont.)

Cisco.com

- **The e-mail worm protection module works through a combination of steps, including dynamically building an application class through the detection of a suspicious action occurring on a system. If this suspicious action detection is seen by the MC as occurring on more than one system, a quarantine of the detected malicious process will also occur.**
- **You can apply a preconfigured policy to systems to detect when a software installation occurs and to add the install process or processes to a dynamically built application class.**
- **The CSA MC lets you enable correlation functions for particular types of events.**
- **You can use the @dynamic token in the File Set text field and in the Network Address Set text field to control access to files and addresses that have been quarantined by CSA MC.**

© 2005 Cisco Systems, Inc. All rights reserved. HIPS v2.0—5-20

Lesson 9

Defining Application Classes

Overview

This lesson explains the application classes shipped with the Cisco Security Agent Management Center (CSA MC) and provides instructions for creating new static and dynamically defined application classes.

This lesson contains the following topics:

- Objectives
- About Application Classes
- Configuring Static Application Classes
- Dynamic Application Classes
- Summary
- Lab Exercise

Objectives

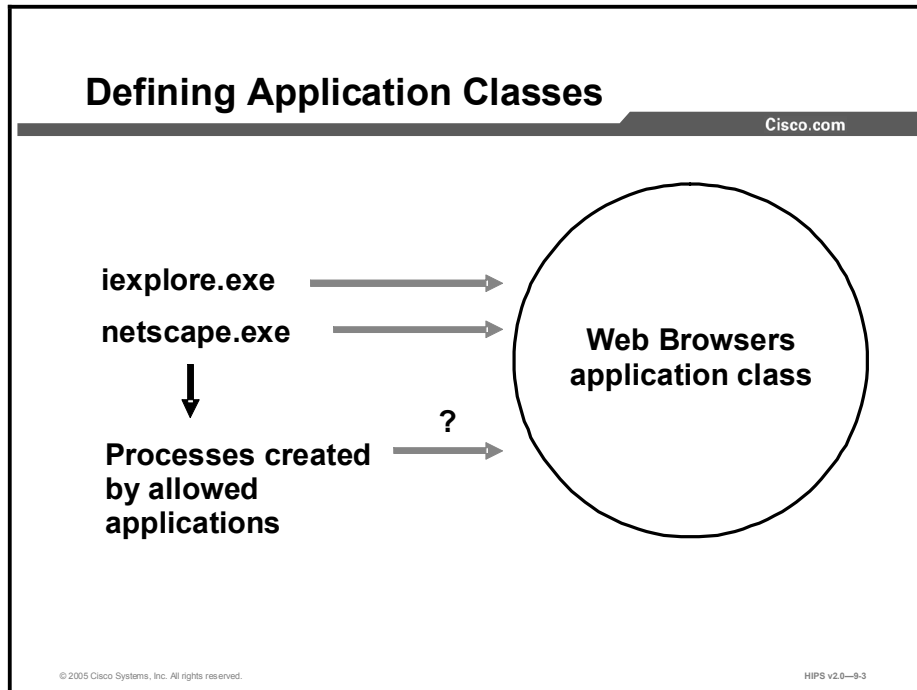
Upon completion of this lesson, you will be able to perform the following tasks:

- Explain the use of application classes in creating security policies
- Discuss the preconfigured application classes included in the CSA MC
- Configure a static application class
- Create a dynamic application class and an application-builder rule

About Application Classes

When you create rules, you must decide which applications are performing the operations that 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; for example, you could enter `netscape.exe` and `iexplore.exe` under the heading of Web Browsers. Then you would 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 that 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 that you define as part of the application class.

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 reevaluated 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 that satisfy both of the following conditions:

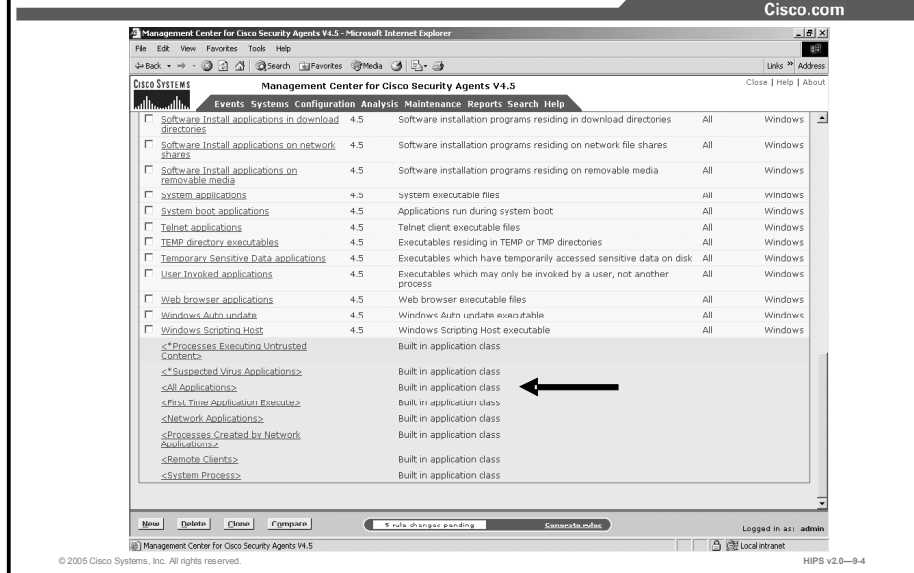
- The script begins with an interpreter string (for example, `#!/bin/bash`).
- The script is executed directly on a command line (for example, `"$foo.sh"`).

Therefore, if you have an application class "foo.sh," a process satisfying these two conditions becomes a member of that application class.

Note A shell may be launched by various methods that do not meet those conditions, for example, `"$. foo.sh"`, or `"$ cat foo.sh | /bin/sh"`. Note also that if you happen to have an application class for a script's interpreter, such as `/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.

Built-in Application Classes



CSA MC ships with several built-in application classes. Those application classes appear inside brackets in the rule application class selection list boxes. Some built-in classes are also marked with asterisks, which indicates that the built-in class is configurable. You can view all application classes in the Application Class list page.

The following application classes are among those included with CSA MC:

- **First Time Application Execute:** This includes the first invocation of any application that has never been observed to execute on the system.
- **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 that viruses are propagated.
- **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 bulleted points).
- **Server (TCP based):** This application class includes all processes that have accepted an inter-box connection on a nonephemeral port.
- **Server (UDP based):** This application class includes all processes that have accepted an inter-box connection on a nonephemeral port.
- **Processes Monitoring the Keyboard:** This includes all processes that continuously monitor keystrokes over an extended period of time.

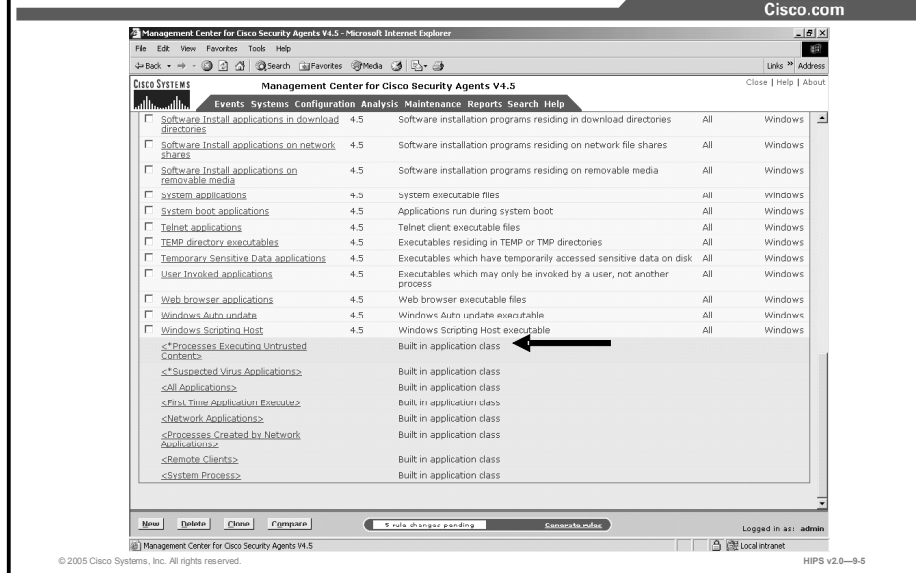
- **Processes with Elevated Privileges:** This application class is only available for UNIX rule types. It includes processes that have elevated user privileges for users other than root, such as ping. Using such processes is a common way to attempt a system break-in.

Note This elevated privilege designation does not apply to processes when the user is logged in as root.

- **Recently Created Untrusted Content:** This includes executables that are newly created by <Processes Writing Untrusted Content> and are immediately invoked.
- **Remote Clients:** 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 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.

Configurable Built-in Application Classes



The Management Center for Cisco Security Agents also ships with built-in application classes that are built by policy rules. These application classes appear in the rule application class selection list boxes inside brackets with asterisks before them (*). This means that you should only use them in conjunction with a rule module that dictates the parameters that cause processes to become classified as one of these application types. CSA MC ships with preconfigured policies to define these classes. You can change these policies, if necessary.

The following configurable built-in application classes are among those included in CSA MC:

- **Authorized Rootkit:** This is intended to identify modules loading after boot time or to identify modules attempting to modify kernel functionality. Processes are classified as belonging to this application category when a kernel protection rule detects and tags the process as an authorized rootkit.

Note If a rootkit gets tagged as both authorized and unauthorized (as a result of tagging rules), an authorized rootkit tag takes precedence over an unauthorized tag.

- **Installation Applications:** This includes processes installing software.
- **Processes Communicating with Untrusted Hosts:** This is intended to capture the IP addresses of hosts that are viewed as violating security policies or exhibiting malicious behavior. Being classified as belonging to this application category causes a host to be quarantined from the network.
- **Processes Copying Untrusted Content:** This is intended to identify processes that copy executables that need to be treated as untrusted and tracked.
- **Processes Executing Untrusted Content:** This includes any downloaded executable or any process that is interpreting downloaded content.
- **Processes Requiring Kernel Only Protection:** This is intended to remediate interoperability issues with the user component in CSA and other third-party software

products. Processes in this class will not enforce Component Object Model (COM) component checks and some buffer overflow checks.

- **Processes Requiring OS Stack Execution Protection:** This application class is only available for UNIX rule types. This is intended to enable native Solaris operating system stack execution protection emulation. This enables additional buffer overflow protection.
- **Processes Requiring Security Level <High, Medium, Low>:** Move applications into this dynamic class to programmatically change the Agent security level based on the current running state of the system. If the Agent security level is low, for example, and a virus is detected on the system, this can trigger a system state policy that will automatically move the security level to high. On a high setting, you may enforce a rule that denies the virus-infected system from making outgoing network connections.
- **Processes Writing Untrusted Content:** This is intended to identify processes that write executables that need to be treated as untrusted and tracked. You could use it to identify a network application that downloads an executable and saves it to disk. The process is the network application and the untrusted content is the downloaded executable.
- **Suspected Virus Applications:** This application class includes processes that are dynamically defined as being suspect by specified, exhibited behavior. Being classified as belonging to this application causes a quarantine message to be sent to CSA MC.
- **Unauthorized Rootkit:** This is intended to identify modules loading after boot time or to identify modules attempting to modify kernel functionality. Processes are classified as belonging to this application category when a kernel protection rule detects and tags the process as an unauthorized rootkit. Then, for example, a system state can take effect as a result of a process being classified as an unauthorized rootkit.

Preserving Application Process Classes

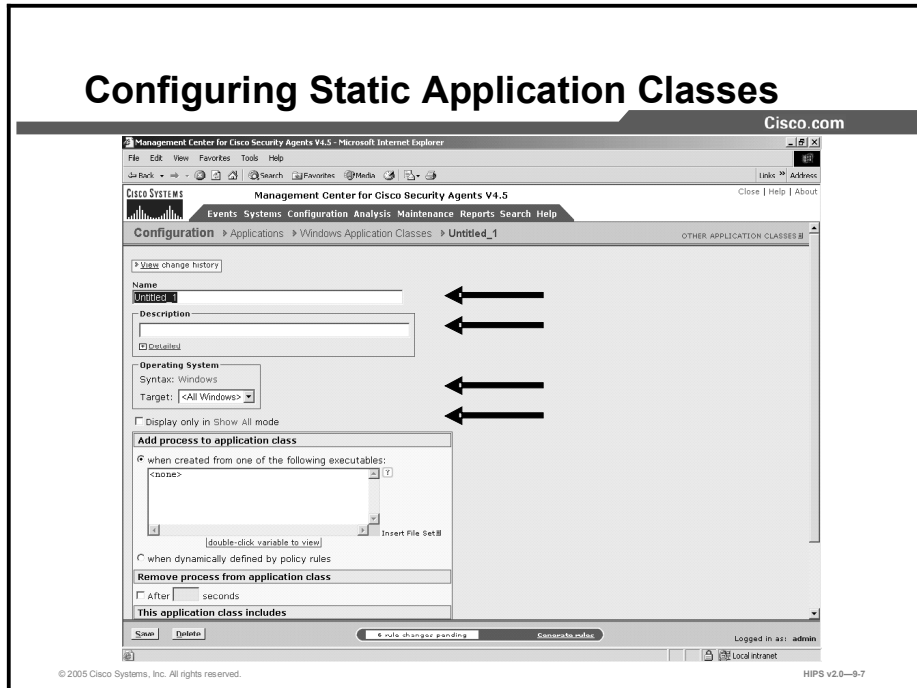
You should be aware that all application process classes are preserved when your policies are changed if those processes (application classes) are used in an existing policy. For example, processes that have been classified by CSA MC as descendents or as network applications are preserved if the application classes that included them are changed in any way.

On policy changes, process name-based application classes are reevaluated. Old application class memberships are not lost, only new memberships are gained.

Configuring Static Application Classes

Access control rules are application-centric. This means that when you write your rules, you should understand that the application(s) you select are 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 that you specify. So, when you begin creating rules, think in terms of the applications that your enterprise as a whole uses and the manner in which you want to limit an application's ability to perform undesired actions.

This topic explains how to configure static application classes.



Complete the following steps to create an application class:

Step 1 Move the mouse over **Configuration** in the menu bar and select **Applications > Application Classes** (Windows or UNIX) from the drop-down list that appears.

Note The list of existing application classes is displayed. CSA MC ships with several preconfigured applications. Some application classes appear within brackets. These are built-in CSA MC application classes and you cannot edit them.

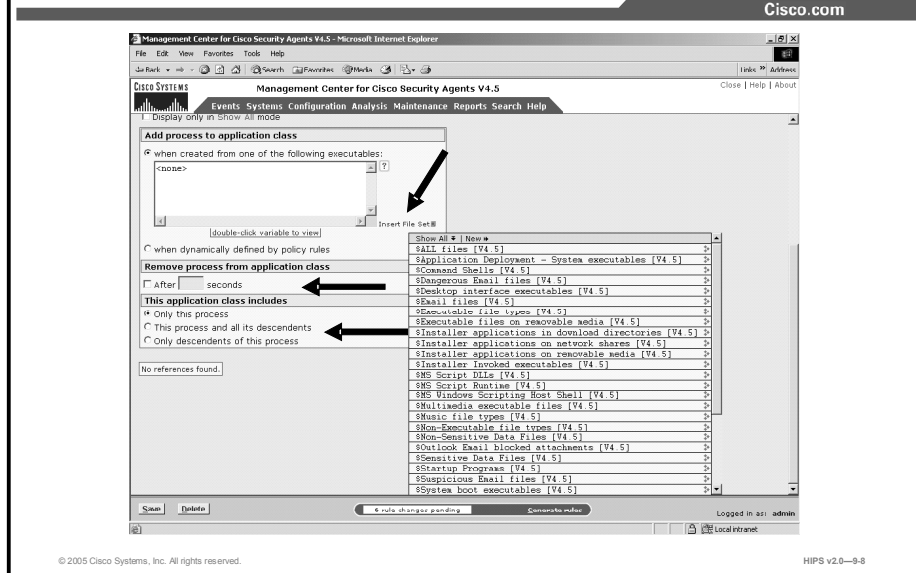
Step 2 Click the **New** button to create a new application class. This takes you to the application class configuration view.

Step 3 Enter a name for the application class that you are creating. 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** Select an operating system. When you create an application class, you must select either a UNIX or a Windows application class. Your application class is then designated for all UNIX or all Windows platforms. Optionally, you can target an operating system more narrowly by selecting a specific UNIX or Windows operating system from the Target drop-down menu.
- Step 6** If your lists of application classes are growing too long in rule pages, click the **Display Only in Show All Mode** check box on an application class page. The item will no longer appear in list pages and selection lists. This feature works in conjunction with Admin Preference settings. You must go to the Admin Preferences page to make the item reappear.

Configuring Static Application Classes (Cont.)



- Step 7** Under Add Process to Application Class, for a static application class, leave the default When Created from One of the Following Executables radio button selected.
- Step 8** Enter the executable filenames (one per line) for the applications you are grouping together in this application class. You can also enter preconfigured file set variables in the executables edit field by clicking the Insert File Set link.
- Step 9** Under Remove Process from Application Class, select the check box beside After and enter a time frame in seconds to configure an application classification that 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.
- Step 10** For UNIX application classes, you have the additional option of selecting the When Session Association Is Voided check box. Selecting this check box causes the application classification to be removed when a process disassociates itself from the current TTY session. When an application class exists for applications descended from "superuser," for example, you might not want the process to continue having the application class of the superuser shell.
- Step 11** When applications are invoked, they often spawn other processes as part of the action that 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

Note Selecting 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 Internet Information Service (IIS) to talk on the network, but create another rule denying descendents of the IIS process from talking on the network.

Step 12 When you are finished, click the **Save** button. This application class name 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 it contains.

Note You can use the Compare button in the application class list view to compare and merge similar application classes.

Dynamic Application Classes

This topic describes dynamic application classes and how to configure them.

Dynamic Application Classes

Cisco.com

Static Application Class:

- **With a static application class, a process is added to the class based on the name of its executable file or the process name.**

Dynamic Application Class:

- **A dynamic application class is defined by process behavior on a system.**

© 2005 Cisco Systems, Inc. All rights reserved.HIPS v2.0—9-10

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 Untrusted Content> application class is a built-in dynamically configured class.

One instance in which you might need a dynamic application class would be if you are writing rules for e-mail clients but you do not know all the different e-mail applications that are being used throughout your corporate network. Any process appearing to act as a client for Simple Mail Transfer Protocol (SMTP) would fall into a dynamic e-mail application class that could be used in rules quarantining dangerous e-mail messages. (You can use your own criteria to define what an e-mail application is.)

Building Classes as Rule Consequences

You can also build a dynamic application class as a consequence of rules triggering. This way, for example, you can configure a query user rule in which a process is added to an application class as a result of a specific user response (yes, no, terminate). You could build a “suspected virus” application class based on a query to the end user when untrusted content arrives on the desktop. If the user clicks the Terminate button on the query box, the process is disallowed. But if the user clicks Yes to allow it, the process would not be added to the suspected virus application class.

Removing Processes from Classes

You can also use a dynamic “remove process” capability in conjunction with dynamically adding a process. For example, you can dynamically add a process to a “suspicious web server descendents” class if a web server spawns a process. Then, if that spawned process attempts to read a script from a normally accessed directory, you can decide this isn’t a dangerous process and have the process removed from the class after the attempt. But if the spawned process attempts to read a script from a directory it should not be accessing, the process should remain in the suspicious web server descendents class.

Defining Dynamic Application Classes

Cisco.com

Define a dynamic application class by doing the following:

- Create a new application class and select the **Processes dynamically defined by selecting the policy rules radio button**.
- Configure an application-builder rule to define your dynamic application class.

© 2005 Cisco Systems, Inc. All rights reserved.

HIPS v2.0—9-11

Define a dynamic application class by doing the following:

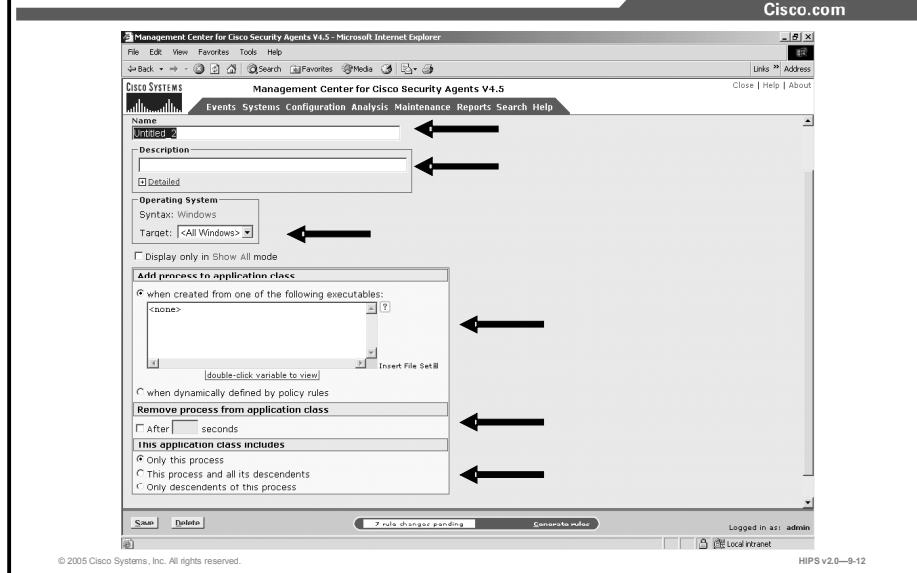
- Create a new application class and select the **When Dynamically Defined by Policy Rules** radio button. (Do not enter any process names in the application class page edit field.)
- Configure an application-builder rule to define your dynamic application class.

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 **Add Process to Application Class** from the drop-down list as the rule action. Then choose the name of the dynamic application class (created in the first bulleted point) from the drop-down 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.

Note A dynamically defined application class can be used in any rule where a static application class can be used.

Configuring Dynamic Application Classes



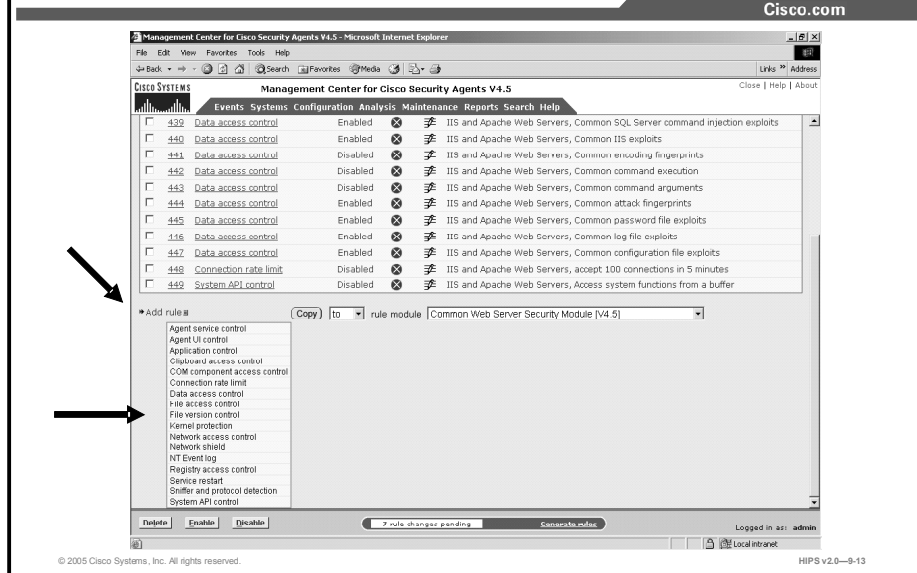
Continuing to use the e-mail client example, we will create an application class that will be dynamically populated by e-mail client applications. You might want to do this if you are writing rules to protect e-mail applications, but you do not know what e-mail applications are being used across your network. Using this dynamic class, rules will restrict e-mail clients based on detected behavior, such as using SMTP to access an e-mail server, rather than by explicitly defining e-mail application executables.

Complete the following steps to create a dynamic application class:

- Step 1** Move the mouse over **Configuration** in the menu bar and select **Applications > Application Classes** (Windows or UNIX) 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.
- Step 3** Enter a name for the dynamic application class that you are creating. Use a descriptive name that you can easily recognize in the application selection lists that appear in the rule views.
- Step 4** Enter a description for your application class.
- Step 5** If your lists of application classes are growing too long in rule pages, click the **Display Only in Show All Mode** check box on an application class page. The item will no longer appear in list pages and selection lists. This feature works in conjunction with Admin Preference settings. You must go to the Admin Preferences page to make the item reappear.
- Step 6** Under Add Process to Application Class, for a dynamic application class, select the **When Dynamically Defined by Policy Rules** radio button. (Do not enter any process names in the edit field.)

- Step 7** Under Remove Process from Application Class, select the check box beside After and enter a time frame in seconds to configure an application classification that 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.
- Step 8** For UNIX application classes, you have the additional option of selecting the When Session Association Is Voided check box. Selecting this check box causes the application classification to be removed when a process disassociates itself from the current TTY session. When an application class exists for applications descended from "superuser," for example, you might not want the process to continue having the application class of the superuser shell.
- Step 9** 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.
- Only This Process
 - This Process and All Its Descendents
 - Only Descendents of This Process
- Step 10** When you are finished, click the **Save** button. This dynamic application class name now appears in the drop-down list beside the Add to Application Class radio button in access control rules and in all application selection fields.

Configuring an Application-BUILDER Rule

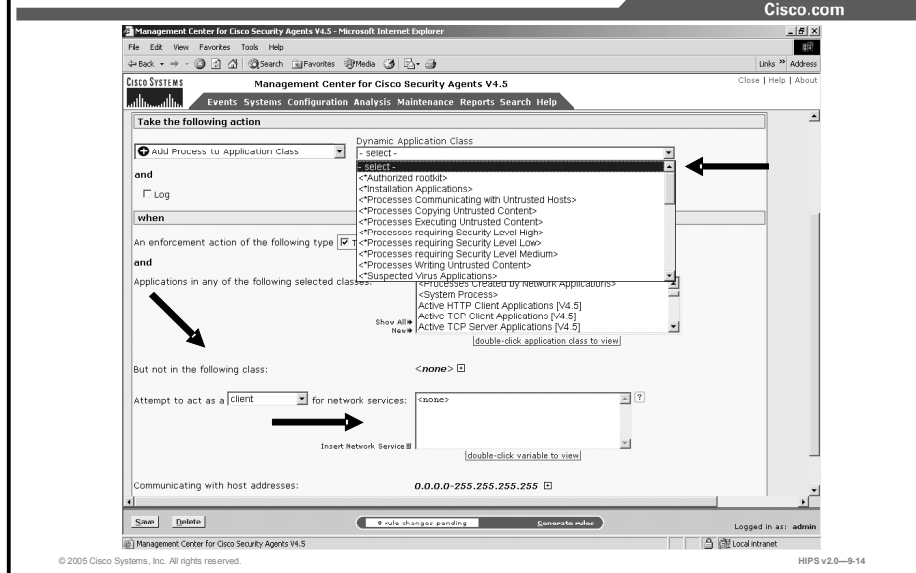


The example shown in the figure uses a network access control rule to define a dynamic application class. You can use any access control rule type as an application-builder rule. Remember, your dynamic application class is not populated with applications until an application-builder rule is triggered by the behavior of the process and added to the class.

Complete the following steps to configure an application-builder rule:

- Step 1** Access the rule module you wish to add the class to, and click the **Modify Rules** link.
- Step 2** Click **Add Rule** and select the access control rule you wish to use.

Configuring an Application-BUILDER Rule (Cont.)



- Step 3** Enter a description in the Description field.
- Step 4** Select the **Add Process to Application Class** option from the Take the Following Action drop-down menu.
- Step 5** Select the dynamic application class from the corresponding drop-down menu.
- Step 6** Optionally, you can select one or more of the available check boxes (Terminate, Deny, Allow) from the When—An Enforcement Action of the Following Type list box. All entries are selected by default, meaning that the tag will apply when the request is made regardless of the action that occurs. All actions apply. If you make a specific selection here, to your dynamic application class will be created based on that action occurring when the request is made (perhaps via another configured rule).

Note All resource requests always result in either an allow, deny, or terminate occurring. Even if there is no rule governing the resource, for example, the implicit action is allow.

- Step 7** 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.
- Step 8** Select **Client** from the drop-down menu under But Not in the Following Class, and select the preconfigured variable that you wish to use for this class.
- Step 9** Leave the default of 0.0.0.0-255.255.255.255 entered in the host addresses field.
- Step 10** Click **Save**.

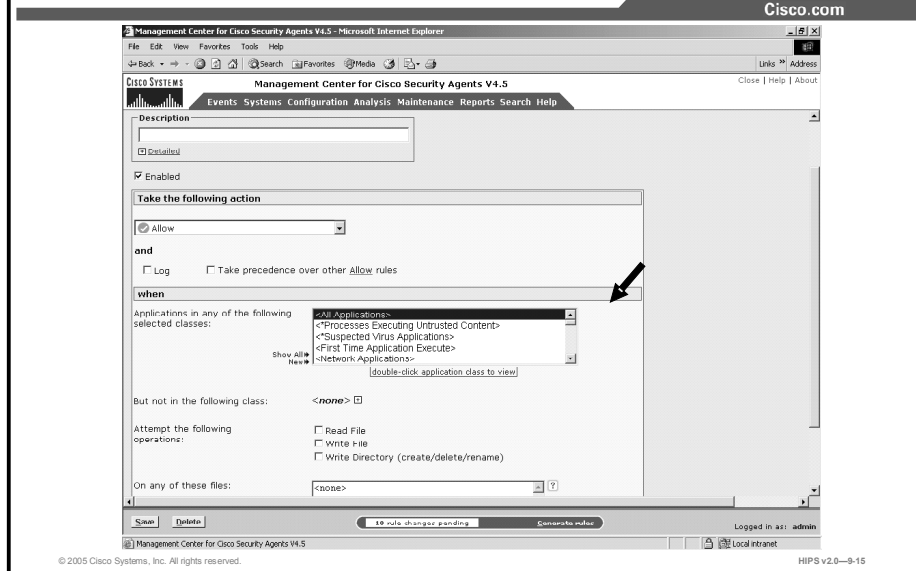
Create New Application Classes from Rule Pages

You can create a new application class from a rule page and have that application class be available to the rule that you are 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 the configuration window. Configure your new application class and click **Save**. It is now available for selection in the 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.

Configuring a Rule Using a Dynamic Application Class

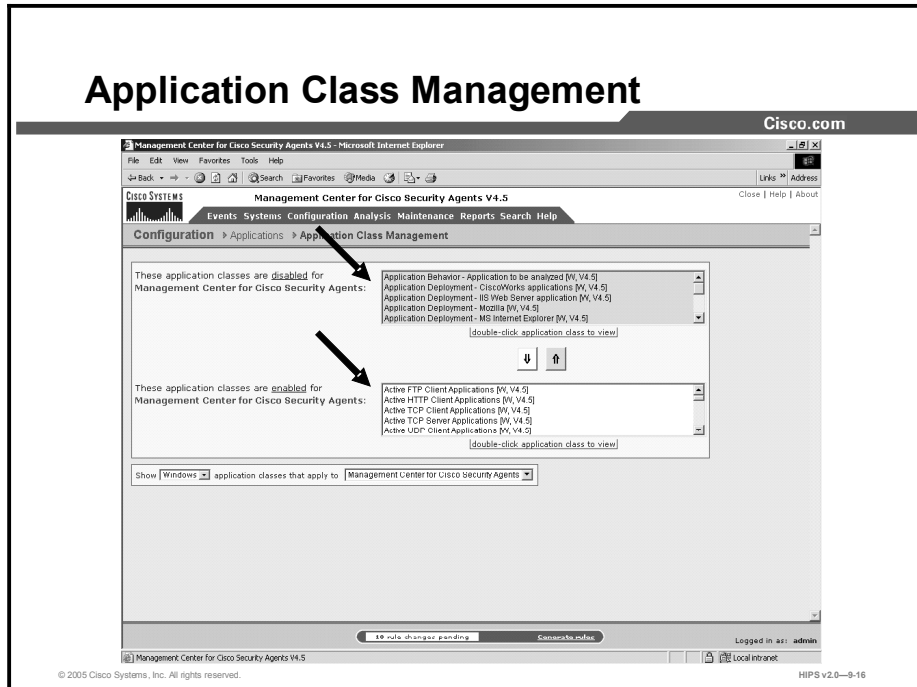


In this example, we are going to use a file access control rule to control the actions of a dynamic application class. Complete the following steps to configure this rule.

- Step 1** Click the **Modify Rules** link.
- Step 2** Click the **Add Rule** link and select **File Access Control**. The file access control configuration window is displayed.
- Step 3** Enter a description in the Description box.
- Step 4** Select **Add Process to Application Class** in the Take the Following Action field.
- Step 5** Choose the class that you wish to use from the Dynamic Application Class drop-down menu.

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.

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 Analysis feature pages. 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 analysis, you can choose to have application classes appear or not appear in features you select.

Choosing not to view certain application classes in certain products does not mean that those application classes are deleted. They will still appear in the main application class list page. They simply will not appear in the application class selection fields in the feature in question. By default, all application classes appear in all application class fields in all feature sets.

Complete the following steps to enable or disable an application for general configuration or for analysis purposes:

Step 1 Select **Configuration > Applications > Application Class Management**. In the Application Class Management page there are swap box fields for CSA MC and for Application Behavior Investigation and Application Deployment Investigation. The application classes appearing in the white swap box(es) (the bottom swap box for each category) are enabled for the feature in question. Those appearing in the gray swap box(es) (the top swap box for each category) will not appear in the feature in question.

Step 2 Select an application class and click the up arrow or down arrow 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.)

Note To narrow the application class categories to specific product components, use the "Show [All, UNIX, Windows] application classes that apply to [<All features>, Management Center for Cisco Security Agents, Application Behavior Investigation, Application Deployment Investigation]."

Summary

This topic summarizes what you have learned in this lesson.

Summary

Cisco.com

- **Rules allow or deny operations by application classes.**
- **Application classes are the key to the rules that you build in your security policies.**
- **Static application classes are defined by the names of the application executables.**
- **Dynamic application classes are defined based on the behavior of an application.**

© 2005 Cisco Systems, Inc. All rights reserved.

HIPS v2.0—9-17

Lesson 10

Working with 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 lesson contains the following topics:

- Objectives
- Variables
- Data Sets
- File Sets
- Network Address Sets
- Network Services Sets
- Registry Sets
- COM Component Sets
- Query Settings
- Summary
- Lab Exercise

Objectives

Upon completion of this lesson, you will be able to perform the following tasks:

- Discuss how events sets are used to ease administration of security policies
- Configure data, file, and network address sets
- Create registry, Component Object Model (COM) component, and network services sets

- Use the COM extraction utility to gather program identifiers (PROGIDs) and class identifiers (CLSIDs) for the software installed on a system
- Configure query setting variables to be used with query rules

Variables

This topic introduces variables.

Variables

Cisco.com

The following are the types of variables that can be used:

- **Event sets**
- **Query settings**
- **File sets**
- **Network address sets**
- **Network services**
- **Registry sets**
- **COM component sets**
- **Data sets**

© 2005 Cisco Systems, Inc. All rights reserved. HIPS v2.0—10-3

We will be discussing the following types of variables in this lesson:

- Event sets
- Query settings
- File sets
- Network address sets
- Network services
- Registry sets
- COM component sets
- Data sets

Using variables is 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.

Display Only in Show All Mode Option

Each individual variable page (including application classes) contains a Display Only in Show All Mode check box. If your lists of variables are growing too long in rule or application configuration pages, you can click this check box on a variable page and that variable will no longer appear in lists for that variable type. To display hidden items, you must go to the Admin Preferences page and choose another preference that always uses show all mode, or change the preference assigned to you.

Data Sets

This topic introduces the configuration of data sets.

Data Sets

Cisco.com

The following are samples of data strings used in HTTP exploits:

///	*(*)*	*[*
]	*^*	*#*	**
* *	*,*	*<?*	*.conf
*%u**.htr*	*.ida*	*.log*	

© 2005 Cisco Systems, Inc. All rights reserved.HIPS v2.0—10-5

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 uniform resource identifier (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 preconfigured data sets that you can use. The preconfigured data sets group patterns to match based upon the following:

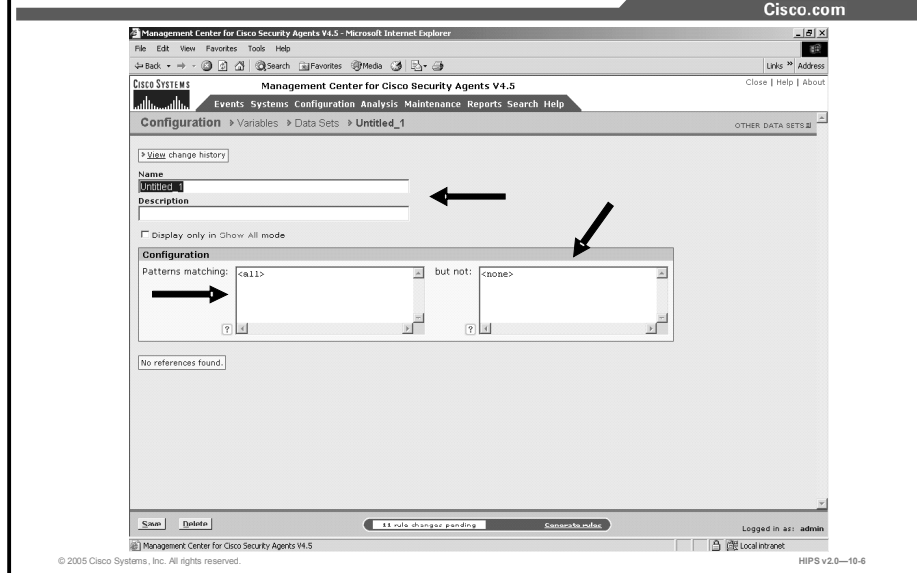
- Functional associations of metacharacters (for example, "(" and ")")
- Examples of known classes of attacks
- Web-server-specific exploits

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/..%255c%255c../winnt/system32/cmd.exe?/c+dir
```

Note Not all preconfigured data sets are used in preconfigured 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, preconfigured data sets used in shipped policies may require modification if legitimate but blocked metacharacters are being used by a web server. Additionally, modifying the preconfigured data sets allows you to block a pattern that specifically matches a new or old exploit or attack.

Configuring a Data Set



Complete the following steps to configure a data set:

- Step 1** Select **Configuration > Variables > Data Sets**. Any existing data set configurations are shown.
- Step 2** Click the **New** button to create a new data set. This takes you to the data set configuration view.
- Step 3** Enter a name in the Name field. This is a unique name for this data set. Generally, it is a good idea to adopt a naming convention that lets you quickly enter data set names in a corresponding rule configuration field.
- Step 4** Enter a description in the Description field. This is a line of text that is displayed in the list view. It will help you to identify this particular data set configuration.
- Step 5** Select the Display Only in Show All Mode check box. If your lists of variables are growing too long in rule or application configuration pages, you can click this check box on a variable page and that variable will no longer appear in selection lists for that variable type. This feature works in conjunction with Admin Preference settings. You must go to the Admin Preferences page to make the item reappear.
- Step 6** Enter a data string in the Patterns Matching pane. Enter the data strings here (one per line) on 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 requested URI to enforce allow or deny data access control rules.

Note When entering data patterns, the * character is a generic wildcard specification.

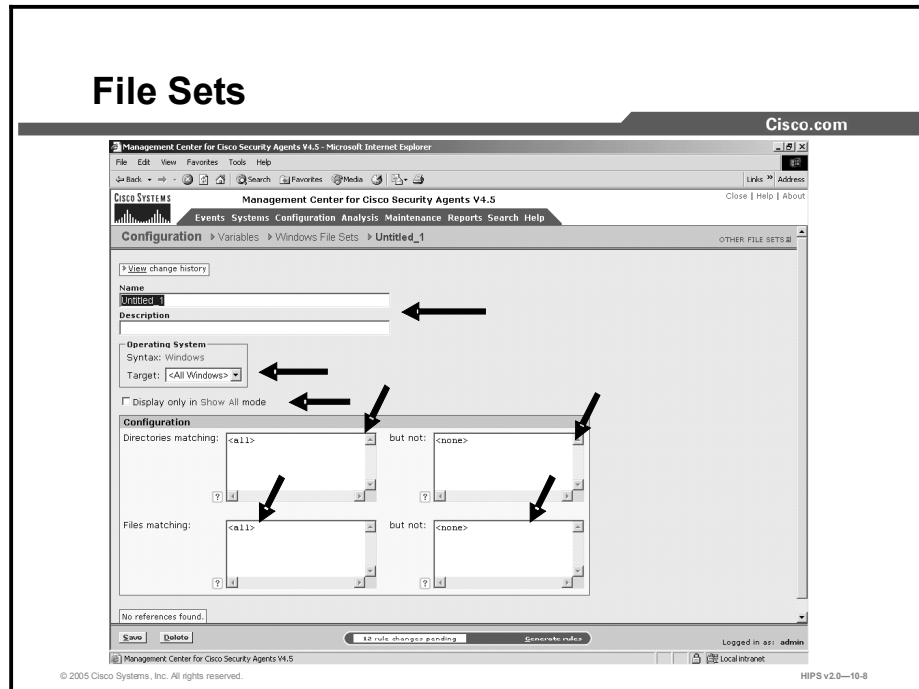
- Step 7** Enter exceptions to the variable in the But Not pane. Make exceptions to the data strings that you have entered in the patterns matching field. By default, this field has

a <none> entry indicating no exceptions. When you click inside this field, the <none> disappears so that you can enter your own exceptions.

- Step 8** 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.

File Sets

This topic discusses file sets.



Configure file sets for use in file access control 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 directory and file permissions and restrictions. All the parameters that exist under that name are then applied to the rule where the name is used. CSA MC ships with several preconfigured file sets that you can use.

Complete the following steps to configure a file set:

- Step 1** Select **Configuration > Variables > File Sets [Unix or Windows]**. Any existing file set configurations are shown.
- Step 2** Click the **New** button to create a new file set. This takes you to the file set configuration view.
- Step 3** Enter a name in the **Name** field. This is a unique name for this file set. Generally, it is 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.
- Step 4** Enter a description in the **Description** field. This is a line of text that is displayed in the list view. It will help you to identify this particular file set configuration.

- Step 5** Select the Display Only in Show All Mode check box. If your lists of variables are growing too long in rule or application configuration pages, you can click this check box on a variable page and that variable will no longer appear in selection lists for that variable type. This feature works in conjunction with Admin Preference settings. You must go the Admin Preferences page to make the item reappear.
- Step 6** Select the operating system from the Target drop-down menu. When you create a file set, you must select to create either a UNIX or a Windows file set. Your file set is then designated for all UNIX or all Windows platforms. Optionally, you select to target an operating system more narrowly by selecting a specific UNIX or Windows operating system from the Target drop-down menu.
- Step 7** Enter the directories and files (one per line) on which you want to impose restrictions in the Directories Matching pane. 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:

Windows example:

```
c:\Program Files\**\*SQL*\bin\**  
\Program Files\**\*SQL*\bin
```

UNIX example:

```
/apache/webroot/**  
/usr/adm/sg
```

- Step 8** Make exceptions to the files and directories that you have entered in the But Not Directories Matching field. For example:

Windows example:

```
c:\Program Files\**\*SQL*\bin\temp
```

Caution The exclusion entry above means that any temp files in the bin folder are ignored by the restrictions that you apply using this file set. This also means that the path that you are protecting in the Directories Matching field is *not* protected when the excluded directory "temp" is being accessed.

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 9** Enter the names of the files to which you are controlling access in the Files Matching field. You can use wildcards here to indicate all of a specific file type, for example:

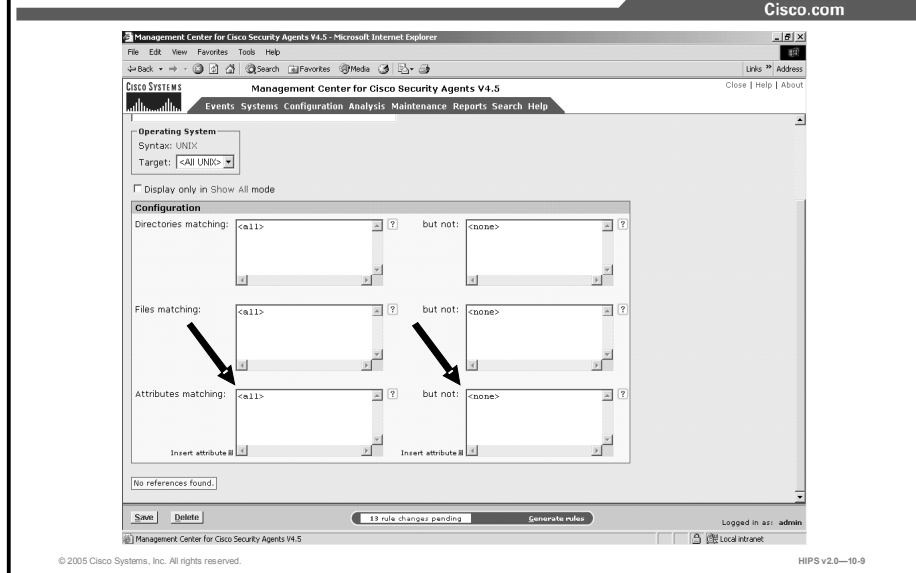
*.exe to specify all executables.

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 10 Make exceptions to the filenames you enter in the But Not 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 in the file set text field to indicate all files that have been quarantined by CSA MC. This list updates automatically (dynamically) as logged quarantined files are received. To view the files that are added to the dynamically quarantined files list and to manually add files to be quarantined, click the **Manage Dynamically Quarantined Files** link on the Global Event Correlation page.

File Sets (Cont.)



Step 11 (UNIX only instruction) File sets created for UNIX have an additional configuration field. In the Attributes Matching edit fields, click the **Insert Attribute** link and optionally select one or more file types to match against. Available file types are as follows:

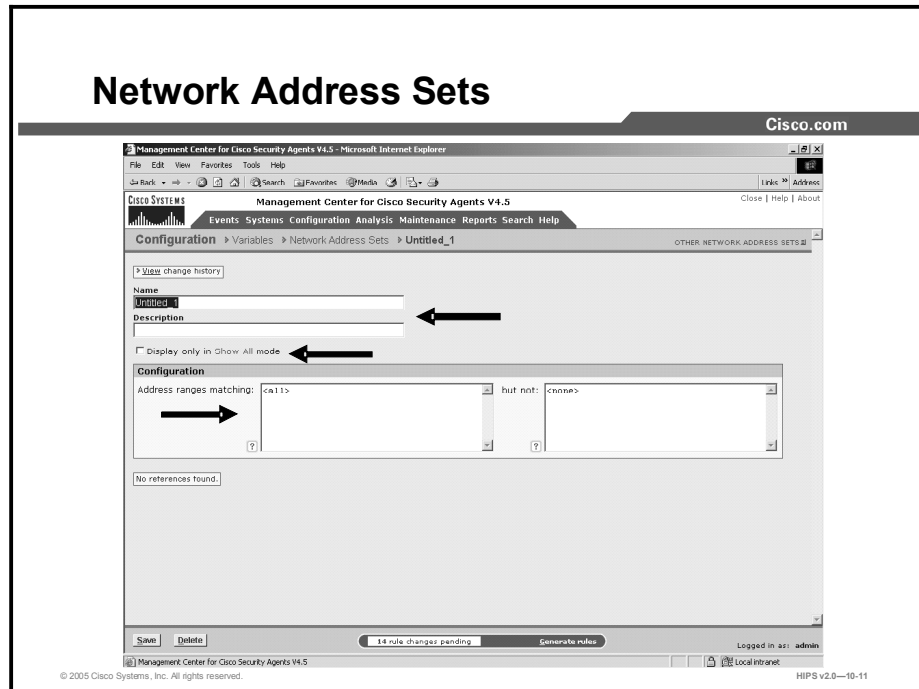
- Block device: A special file used for buffered or block I/O. For example, a disk device.
- Character device: A special file used for unbuffered or character I/O. For example, a tty file.
- Executable file: A file identified in /etc/magic as being executable.
- Interpreter file: A file that contains a script (shell, Perl, etc.) where the first line starts with "#! interpreter [arg]" .
- Java class file: A file identified in /etc/magic as being executable Java byte code.
- Setgid file: A file with the "set group ID on execution" property set in the file mode.
- Setuid file: A file with the "set user ID on execution" property set in the file mode.

Step 12 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 is a View Change History link. Click this link to go to a page that 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.

Network Address Sets

This topic discusses 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 that you create.

Complete the following steps to configure network address sets:

- Step 1** Select **Configuration > Variables > Network Address Sets**. Any existing address set configurations are shown.
- Step 2** Click the **New** button to create a new network address set. This takes you to the configuration view.
- Step 3** Enter a name in the **Name** field. 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.
- Step 4** Enter a description in the Description field. This is a useful line of text that is displayed in the list view. It helps you to identify this particular set of addresses.
- Step 5** Select the Display Only in Show All Mode check box. If your lists of variables are growing too long in rule or application configuration pages, you can click this check box on a variable page and that variable will no longer appear in selection lists for that variable type. This feature works in conjunction with Admin Preference settings. You must go the Admin Preferences page to make the item reappear.

Network Address Sets (Cont.)

Cisco.com

Address range syntax:

- Use one entry per line.
- Use a hyphen to indicate address ranges.
- Address ranges are inclusive.
- Use **@local** to indicate all local addresses on the system.

© 2005 Cisco Systems, Inc. All rights reserved.

HIPS v2.0—10-12

Step 6 Enter the IP address ranges in the Enter Address Ranges pane. In the available edit field, enter a single address or a 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.
- Use a hyphen to indicate the range.
- Address ranges are inclusive, for example:

```
128.66.24.130
```

```
128.67.2.10-20
```

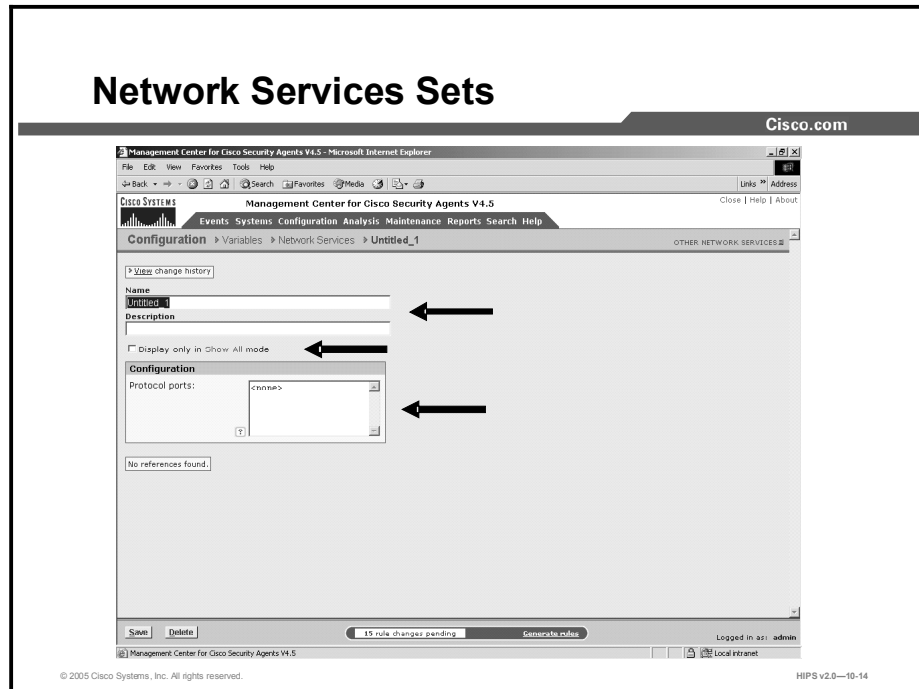
Use **@local** to indicate all local addresses on the Agent system. You would want to use this if you are allowing different applications on a single system to talk to each other without opening these applications to network access.

Note Use **@dynamic** in the addresses set field to indicate untrusted hosts that have been quarantined by CSA MC. Addresses are added to this list when they are seen as an untrusted host. (The built-in "Processes Communicating with Untrusted Hosts" is triggered in a rule.) This list updates automatically (dynamically) as logged quarantined addresses are received.

Caution On UNIX platforms, IP version 6 (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.

Network Services Sets

This topic explains the configuration and deployment of network services sets.



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 preconfigured network services that you can use.

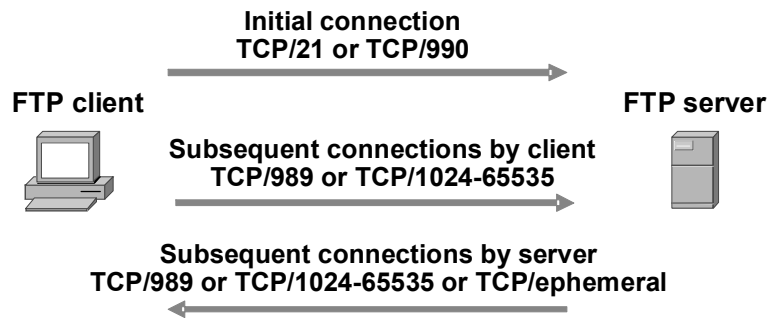
Complete the following steps to configure network services:

- Step 1** Select **Configuration > Variables > Network Services**. Any existing configurations are shown.
- Step 2** Click the **New** button to create a new network service variable. This takes you to the configuration view.
- Step 3** Enter a name in the Name field. This is a unique name for this network service configuration. This name is case insensitive. Generally, it is 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.
- Step 4** Enter a description in the Description field. This is a useful line of text that is displayed in the list view. It helps you to identify this particular configuration.

- Step 5** Select the Display Only in Show All Mode check box. If your lists of variables are growing too long in rule or application configuration pages, you can click this check box on a variable page and that variable will no longer appear in selection lists for that variable type. This feature works in conjunction with Admin Preference settings. You must go the Admin Preferences page to make the item reappear.

Network Services Sets (Cont.)

Cisco.com



Some protocols create additional connections as part of the same session started by the initial connection.

© 2005 Cisco Systems, Inc. All rights reserved.

HIPS v2.0-10-15

- Step 6** Enter the protocol ports in the Protocol Ports pane. Enter a TCP or User Datagram Protocol (UDP) protocol and corresponding port or port range to indicate a restriction according to the system that is initiating the connection. By default, this field has a <none> entry indicating no ports. When you click inside the edit field, the <none> disappears so that you can enter your own port restrictions, for example:

Use the following syntax:

TCP/21

UDP/1025-65535

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 must be defined as another network service and used appropriately in a rule module to consider callback connections. 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:

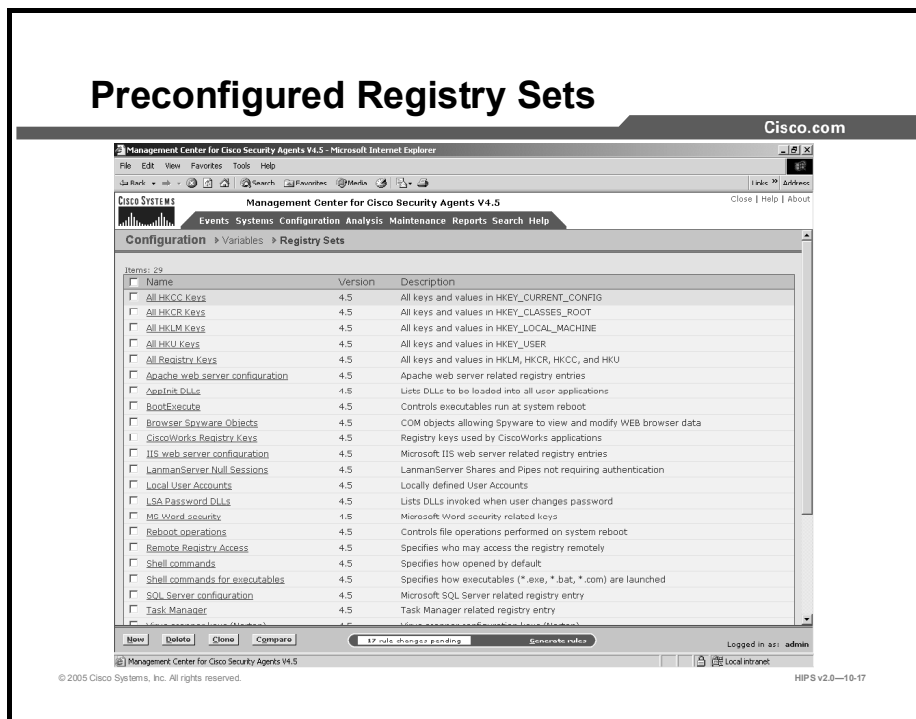
TCP/ephemeral

UDP/ephemeral

- Step 7** 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.

Registry Sets

This topic introduces the configuration of 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 that are 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.

Included Registry Sets

CSA MC ships with several preconfigured registry sets that 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.

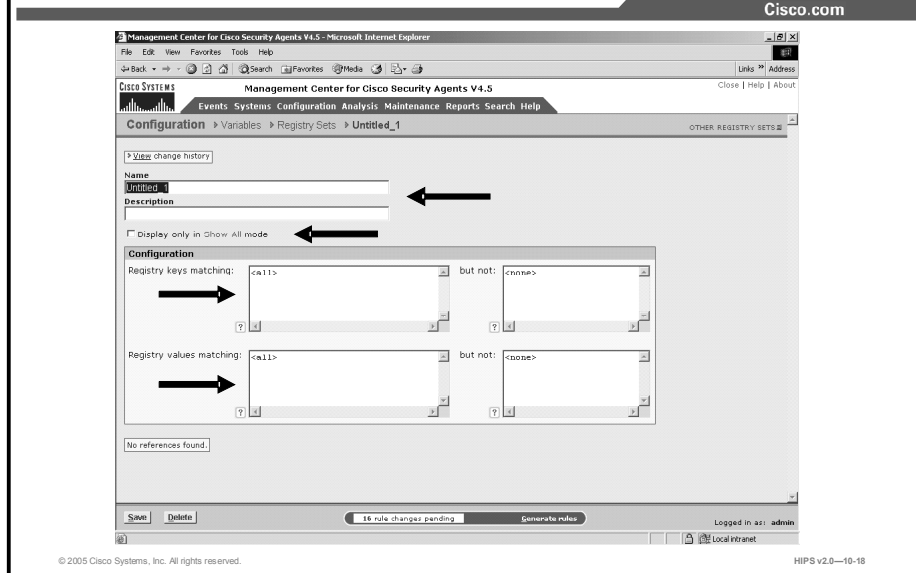
- Run keys are used to register programs so that the system will invoke them as a service. Viruses can make use of this key to become persistent. Protecting this registry value by creating a rule to prevent writing to run keys can prevent certain viruses from invoking and propagating themselves.

Note 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. If users are not installing software, the triggering of this type of query user rule could be treated as a serious issue, and users 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 could go undetected.

Note 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. If users are not uninstalling software, the triggering of this type of query user rule could be treated as a serious issue, and users should answer "No to all" to disallow the action.

Registry Sets



Complete the following steps to configure registry sets:

- Step 1** Select **Configuration > Variables > Registry Sets**. Any existing registry set configurations are shown.
- Step 2** Click the **New** button to create a new registry set variable. This takes you to the configuration view.
- Step 3** Enter a name in the Name field. This is a unique name for this registry set.
- Step 4** Enter a description in the Description field. This is a line of text that is displayed in the list view. It helps you to identify this particular registry set configuration.
- Step 5** Select the Display Only in Show All Mode check box. If your lists of variables are growing too long in rule or application configuration pages, you can clicking this check box on a variable page and that variable will no longer appear in selection lists for that variable type. This feature works in conjunction with Admin Preference settings. You must go the Admin Preferences page to make the item reappear.
- Step 6** Enter a registry key in the Registry Keys Matching pane. 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.

Hives are one of the following strings:

HKLM—refers to the HKEY_LOCAL_MACHINE

HKCR—refers to HKEY_CLASSES_ROOT

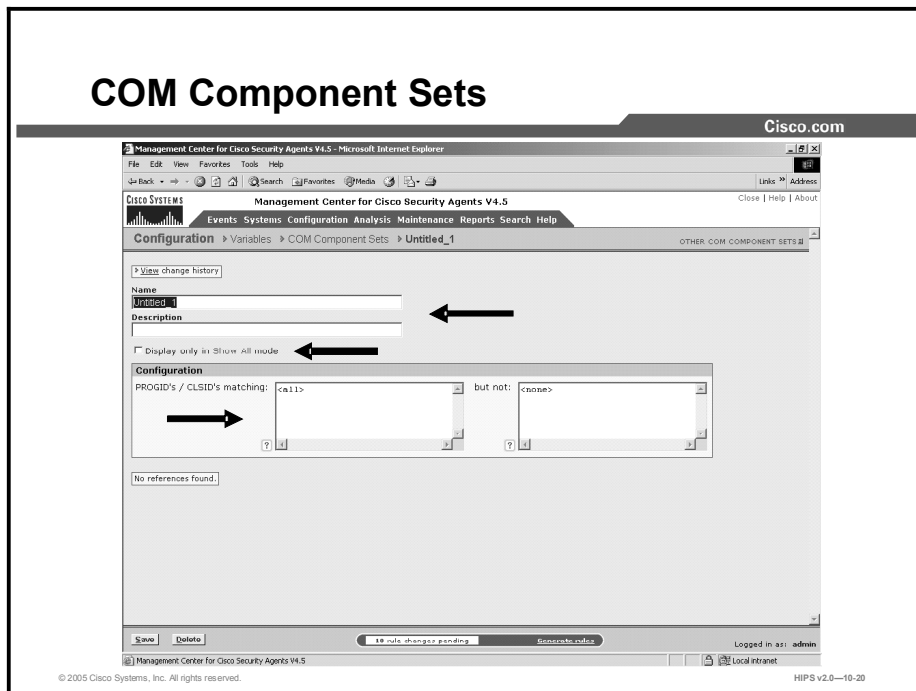
HKCC—refers to HKEY_CURRENT_CONFIG

HKU—refers to HKEY_USERS (HKU* refers to all users)

- Step 7** Enter exceptions to registry keys in the But Not pane.
- Step 8** Enter the registry values that you are controlling access to in the Registry Values Matching pane.
- Step 9** Enter exceptions to registry values in the But Not pane.
- Step 10** 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.

COM Component Sets

This topic discusses COM component sets.



Configure COM component sets for use in COM component access control rules. COM objects are groupings of COM PROGIDs and/or COM CLSIDs 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 preconfigured COM component sets that you can use as well.

This is not available for UNIX configurations.

Complete the following steps to configure a COM component set:

- Step 1** Select **Configuration > Variables > COM Component Sets**. Any existing COM component set configurations are shown.
- Step 2** Click the **New** button to create a new COM component set. This takes you to the configuration view.
- Step 3** Enter a name in the Name field. This is a unique name for this COM component set. Generally, it is a good idea to adopt a naming convention that lets you quickly enter COM component set names in a corresponding rule configuration field.
- Step 4** Enter a description in the Description field. This is a line of text that is displayed in the list view. It helps you to identify this particular COM component set configuration.

Step 5 Select the Display Only in Show All Mode check box. If your lists of variables are growing too long in rule or application configuration pages, you can click this check box on a variable page and that variable will no longer appear in selection lists for that variable type. This feature works in conjunction with Admin Preference settings. You must go the Admin Preferences page to make the item reappear.

Step 6 Enter the COM component PROGIDs or CLSIDs (one per line) on which you want to impose restrictions in the PROGID's/CLSID's Matching pane. By default, this field has an <all> entry indicating all PROGIDs and CLSIDs. When you click inside this field, the <all> disappears so that you can enter your own restrictions.

When entering PROGIDs, use syntax as shown in the following example:

```
Outlook.Application
```

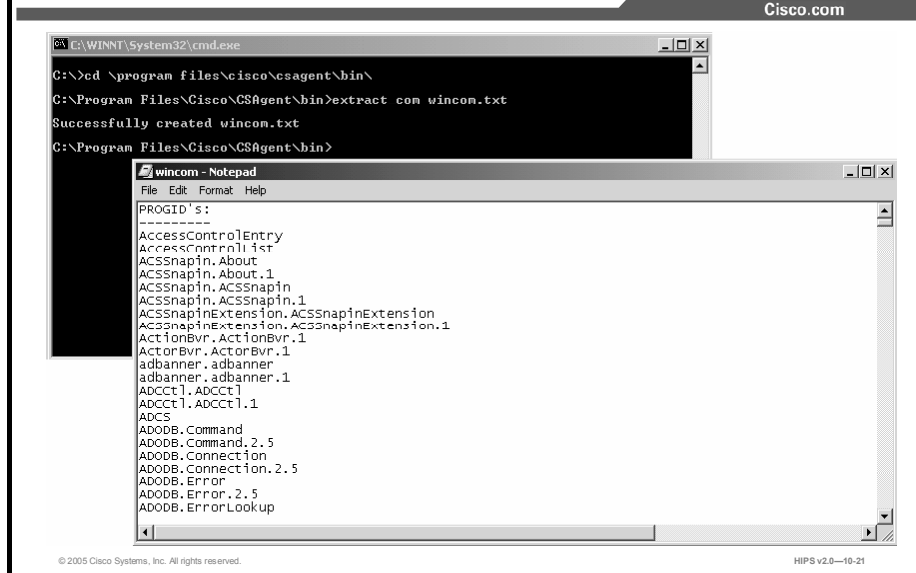
When entering CLSIDs (uppercase hexadecimals), use the following syntax. (You must include the brackets shown here.)

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

Step 7 Make exceptions to PROGIDs or CLSIDs that you want to restrict in the But Not 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.

Step 8 When all required information is entered, click the **Save** button to save your COM component set in the CSA MC database.

COM Component Extraction 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 PROGIDs and CLSIDs for software installed on the system in question and places this data in a text file. You can cut and paste these IDs 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 Systems\CSAgent\bin` directory, type **`extract_com filename`** ("filename" is the name of the text file you want the utility to create). All COM PROGID and CLSID data is placed in this file. For example, enter:

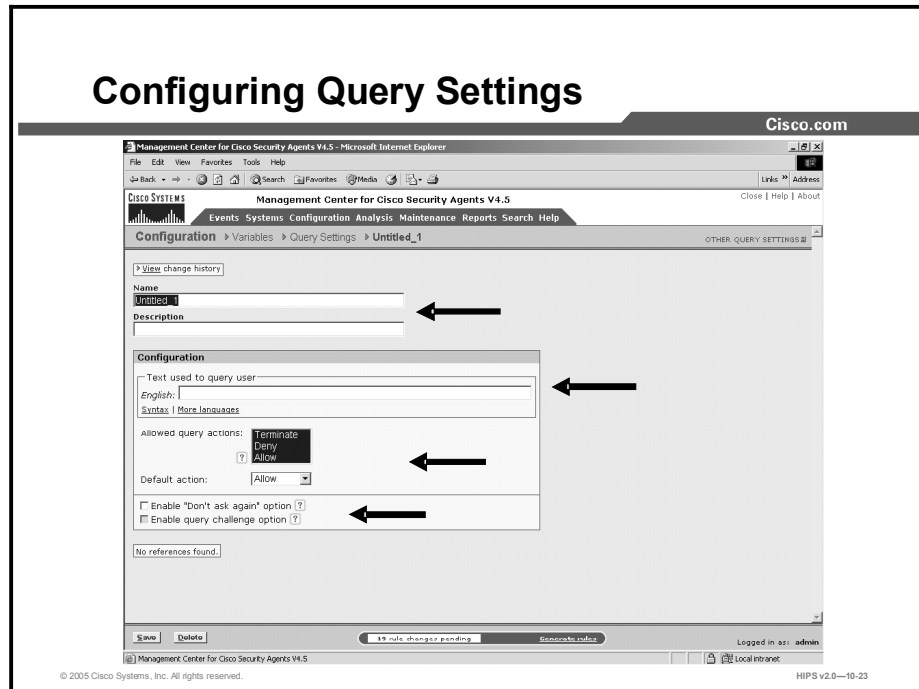
```
\Cisco Systems\CSAgent\bin>extract_com foo.txt
```

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 specific syntax for entering PROGIDs and CLSIDs. The COM component file created by the `extract_com` utility may display PROGIDs and CLSIDs without the proper syntax in the output file. Despite this, when you enter these IDs into text fields for rules or variables you *must* use the correct syntax.

Query Settings

This topic discusses Query settings.



To configure a query popup box for use with a query rule, complete the following steps:

Step 1 Select **Configuration > Variables > Query Settings**. Any existing query settings are displayed.

Note For a query setting, the response to the query is relevant to the question, not to the resource. For example, if a file access control rule queries the user for a response and that identical query is also configured for a network access control rule, the user is not queried again when the network access control rule triggers. The query response from the previous file access control rule is automatically taken.

Step 2 On the Query Settings list page, click the **New** button to create a new query.

Step 3 Enter a unique name for your query in the Name field. 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 rule selection box when you are selecting a specific query setting for a rule.

Step 4 Enter a description of your query in the Description field.

Step 5 In the Text Used to Query User edit field, enter a description of the issue that likely triggered the query. This text field allows you to provide localized query text for Agents using the corresponding language on their desktop. This is the same text that will appear in the query user popup box explaining what is occurring on the system to the user. Therefore, making this information descriptive of the system action that

triggered the popup is important. You can use specially designated tokens to represent the corresponding values presented to the end user who is responding to the query.

Note All Cisco Security Agent kits contain localized support for French, German, and Japanese language desktops. If you do not select a specific language, the default for query text is English. Click the **More Languages** link to enter text to be displayed in a language other than English. This allows you to provide localized query text for Agents using the corresponding language on their desktop.

Step 6 The Allowed Query Actions selection box lets you choose which radio buttons appear on the query popup box. You may not want the user to have a “Terminate” option, for example. Therefore, you would only select the Allow and Deny radio buttons to be displayed. The user reads the information posted on the query and is given the option of selecting one of the following possible choices and clicking Apply:

- **Allow (Yes)**—Allows the application access to the resource in question.
- **Deny (No)**—Denies the application access to the resource in question.
- **Terminate**—Denies the application access to the resource in question and also attempts to terminate the application process. (Some processes cannot be safely terminated, such as winlogon.)

Step 7 Of the radio buttons you decide to display, you also choose one of those buttons in the Default Action drop-down box. If the query is not answered by the user within 5 minutes, or if the user is not logged in to the system, the default action is taken immediately.

Step 8 You can also decide to display a “Don't ask again” check box so that the user's query response is remembered. If the user selects that check box when he or she responds to the query, and then attempts the same action on the same resource, the remembered response is automatically taken and the user is not queried again.

Step 9 For added security, you can issue a query challenge on the query popup box. If the default answer is not selected by the user and the selected answer is weaker than the default, a challenge will appear. This ensures that the user sitting in front of the system is answering the query rather than a malicious remote user or program attempting to respond. To pass the challenge, the user enters the information displayed in a graphic on the popup box itself.

Step 10 Click the **Save** button to save your changes.

Note When you phrase the question that will appear to users and select the radio button options to be displayed, make sure that the logic you use is in sync with the response that the user should select. For example, you probably should not phrase a question in the following way: “Do you want to prevent this action from occurring?” In this case, if the response is Yes, this is counterintuitive to how queries should be used. The user is selecting Yes to indicate No. Instead, phrase the question as follows: “Select No to prevent this action from occurring.”

Query Tokens

Cisco.com

When entering query text in the edit field, you can use the following tokens to represent the values presented to the end user who is responding to the query.

- @parent
- @ActiveXname
- @appname
- @child
- @progid
- @clsid
- @dataname
- @filename
- @fileop
- @funcname
- @hostaddr
- @localaddr
- @netop
- @netservice
- @regname
- @targetapp

© 2005 Cisco Systems, Inc. All rights reserved.

HIPS v2.0—10-24

When entering query text into the edit field, you can use the following tokens to represent the values presented to the end user who is responding to the query.

- @parent—The path of the parent process. Use in application control rules only.
- @ActiveXname—The name of the ActiveX control being downloaded. Use in system API access control rules only.
- @appname—The path of the process triggering the action. Use in all access control rule types, except application control rules.
- @child—The path of the process being invoked. Use in application control rules only.
- @progid—The PROGID of the COM object. Use in COM component access control rules only.
- @clsid—The GUID of the COM object. Use in COM component access control rules only.
- @dataname—The name of data being filtered. Use in data access control rules only.
- @filename—The full file path of the file being accessed. Use in file access control rules only.
- @fileop—The type of file operation (file/directory, read/write). Use in file access control rules only.
- @funcname—The system API function being called. Use in system API access control rules only.
- @hostaddr—The remote address of a connection. Use in network access control rules only.
- @localaddr—The local address of a connection. Use in network access control rules only.
- @netop—The type of network operation (client/server). Use in network access control rules only.
- @netservice—The service/destination port used by the remote connection end. Use in network access control rules only.

- @regname—The registry entry being accessed. Use in registry access control rules only.
- @targetapp—The path of the application being targeted for code injection or modification. Use in system API access control rules only.

Localized Language Version Support

On systems running multiple locales (for example, Multilingual User Interface installations or Terminal Services), queries are displayed in the supported language used for the Windows desktop on which the query is shown. Events appear in the Windows Event Log in the default system's language. On a Windows 2000 Multilingual User Interface (MUI) installation, for example, if a user is running a Japanese language version desktop, queries will appear in Japanese. But the Windows Event Log on this system will store events formatted in U.S. English because the system language on a Windows MUI system is English. On a localized Japanese system, both the queries and the events appearing in the Windows Event Log appear in Japanese.

Summary

This topic summarizes what you learned in this lesson.

Summary

Cisco.com

- **Variables are configuration data items, such as network addresses or services, registry or data values, or COM components.**
- **Variables allow rules to be built more easily and changes to be deployed faster.**
- **Data sets name text strings and metacharacters.**
- **File sets group files and directories.**
- **Network address sets specify IP addresses.**
- **Network services sets name protocol/port pairs.**
- **Registry sets group registry keys and values.**
- **COM component sets specify PROGIDs and CLSIDs.**
- **Query settings can be configured to work with query rules.**

© 2005 Cisco Systems, Inc. All rights reserved.

HIPS v2.0—10-26

Lesson 11

Using Cisco Security Agent Analysis

Overview

This lesson introduces the Cisco Security Agent (CSA) Analysis software and explains how to configure and use the Analysis feature. This lesson includes the following topics:

- Objectives
- Application Deployment Investigation
- Group Settings
- Product Associations
- Data Management
- Application Deployment Reports
- Application Behavior Investigation
- Behavior Analysis Reports
- Behavior Analysis Rule Modules
- Summary
- Lab Exercise

Objectives

Upon completion of this lesson, you will be able to perform the following tasks:

- Understand and configure application deployment investigation
- Understand and configure product associations for application deployment investigation
- Understand how to archive and purge data collected by analysis
- Configure and run application deployment reports
- Understand and configure application behavior investigation
- Understand and use behavior analysis reports

- Import and use behavior analysis rule modules

Application Deployment Investigation

This topic introduces application deployment investigation.

Application Deployment Investigation

Cisco.com

With application deployment investigation, administrators can perform the following tasks:

- **See what applications are running on systems and determine what their usage patterns are.**
- **See what applications are installed but remain largely unused on systems.**
- **See what applications are accessing critical network resources.**
- **Use collected data to accurately deploy policies or to generate new policies for unprotected applications using the Cisco Security Agent.**

© 2005 Cisco Systems, Inc. All rights reserved. HIPS v2.0—11-3

Cisco Security Agent Analysis functionality works with Cisco Security Agent Management Center (CSA MC) and the Agent, serving as a data collection and behavior analysis tool for administrators who are deploying policies across systems and networks.

Because the rules that make up policies are aimed at protecting your enterprise resources, knowing exactly what those resources are and how they are used is essential to deploying effective policies.

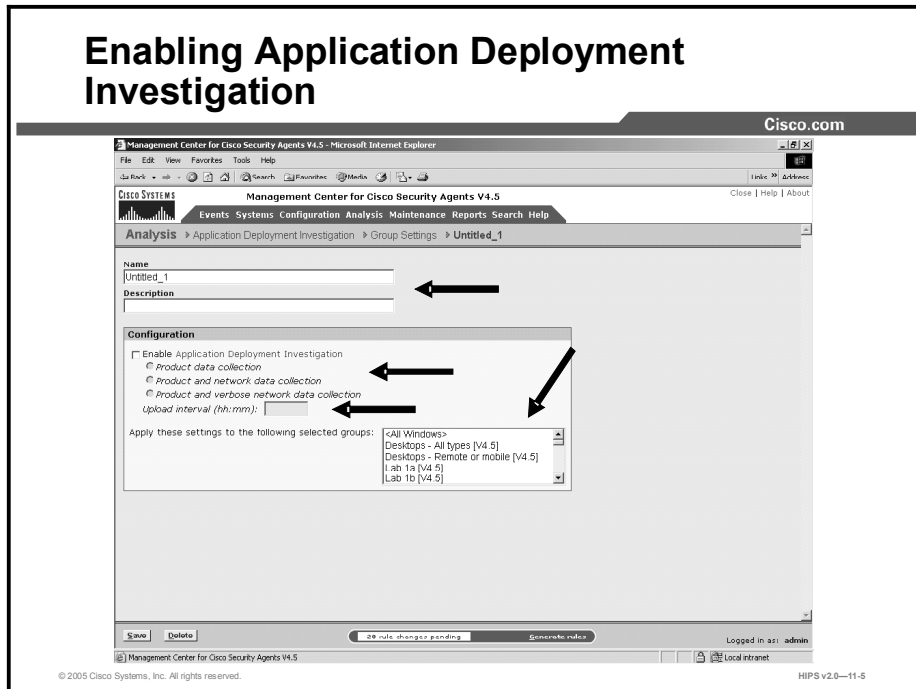
Application deployment investigation allows you to perform the following tasks:

- You can see what applications are running on systems and determine what their usage patterns are.
- You can see what applications are installed but remain largely unused on systems.
- You can see what applications are accessing critical network resources.
- You can use collected data to accurately deploy policies or to generate new policies for unprotected applications using the Cisco Security Agent.

Application deployment investigation is only supported on Windows platforms. By default, application deployment investigation is disabled for all Windows groups until you enable it.

Group Settings

This topic explains how to configure group settings for analysis.



Deployment investigation is controlled on a per-group basis, and it is enabled or disabled using the **Analysis > Application Deployment Investigation > Group Settings** page.

If deployment investigation is enabled for the group, it begins on all hosts in the group in question after you generate rules and the hosts next poll in to CSA MC. If you want to enable application deployment investigation for only one host, you must create a new group with application deployment investigation enabled and add the host to that group. If a host belongs to multiple groups, having application deployment investigation enabled—if present in any group of which the host is a member—takes precedence over not having it enabled. Once application deployment investigation is enabled for a group, it continues to collect data until you disable it and generate rules.

Complete the following steps to configure group settings for application deployment investigation:

- Step 1** Select **Analysis > Application Deployment Investigation > Group Settings**. Any existing group settings are displayed.
- Step 2** Click the **New** button to create a new group setting.
- Step 3** Enter a name in the Name field. This is a unique name for this group setting. 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.
- Step 4** Enter a description in the Description field. This is a useful line of text that is displayed in the list view and helps you to identify this particular group setting.

- Step 5** Click the **Enable Application Deployment Investigation** check box and select one of the following radio button options:
- **Product Data Collection:** This applies to the following reports: AntiVirus Installations, Installed Products, Unprotected Products, Product Usage.
 - **Product and Network Data Collection:** This applies to the following additional report: Network Server Applications.
 - **Product and Verbose Network Data Collection:** This applies to the following additional reports: Unprotected Hosts, Network Data Flows.

Note It is recommended that you choose the lowest verbosity level available in reports whenever possible to keep the volume of network data collection manageable.

Step 6 Enter an Upload Interval time for the Agent to send collected data to the CSA MC. The default and minimum interval is 24 hours. Note that uploads occur at the end of an interval. Therefore, it may take longer than one interval to receive the initial data.

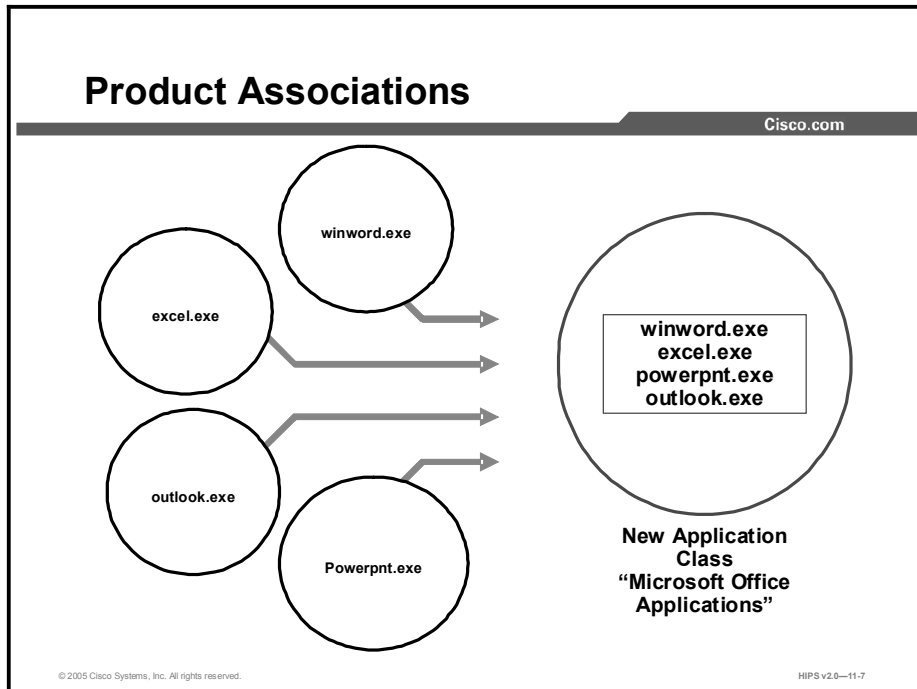
Step 7 In the Apply These Settings to the Following Selected Groups list box, select one or more groups for data collection.

Step 8 Click the **Save** button when your group settings configuration is finished.

Deployment investigation begins on all hosts in the group in question after you generate rules and the hosts next poll in to CSA MC.

Product Associations

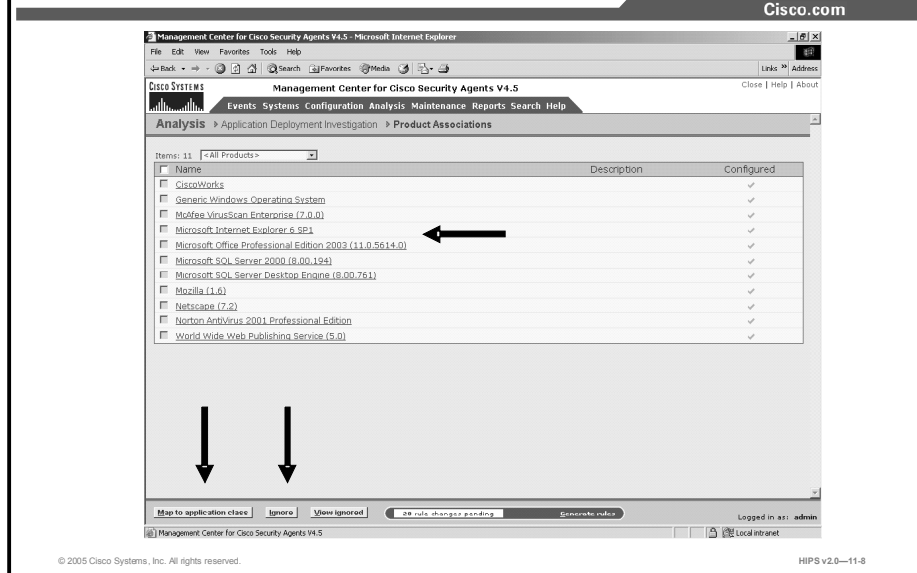
This topic discusses product associations and how to configure them.



You can use application deployment investigation to generate reports that use installed software products as part of the report criteria. In order to do this, there is some prerequisite configuration. This is necessary because the deployment investigation process, in part, gathers data on systems according to the application name that it finds. That is the application executable itself and not the product with which the application is associated. Application deployment investigation does gather information on the products that are installed on systems, but it does not then map the applications back to the installed products. For example, application deployment investigation may find that excel.exe is running on a system. It may also find that Microsoft Office is installed on that same system. But it will not know that excel.exe is part of Microsoft Office. You must tell it so. Therefore, in order to generate certain report types using installed product information, you must first associate the installed products found by application deployment investigation with the application(s) that make up the product. (This could entail creating new application classes for this purpose.) You must make this application class/product association to use product criteria to generate the Product Usage report type.

Caution Preconfigured application classes that ship with CSA MC are not available to application deployment investigation functionality. It is recommended that you configure application classes that are separate and solely for the purpose of analysis reports and investigation. This way, you are not compromising existing application classes that are used in CSA MC security policies.

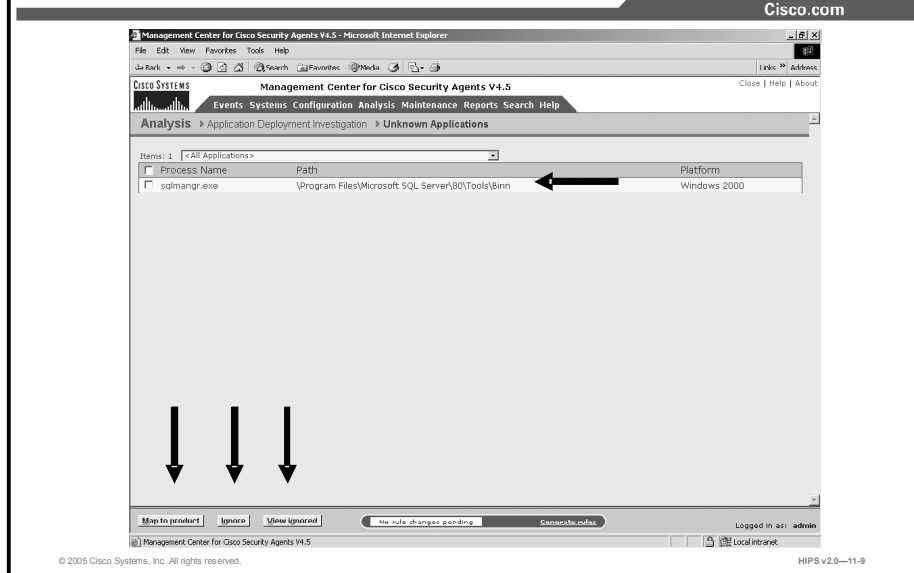
Configuring Product Association



Complete the following steps to create application class/product associations, after application deployment investigation has collected data:

- Step 1** Select **Analysis > Application Deployment Investigation > Product Associations**. The deployment investigation Products page contains a list of all the installed products (not applications) found on systems that were investigated. These are the product names that would be viewable through the Microsoft Add/Remove Programs window.
- Step 2** To associate a product with an application, click the product in the Product Associations window. This takes you to a window where you can select an application class or classes that will define the product. You can also associate a product with an application by selecting the check box beside the product name link and clicking the Map to Application button. This opens a new window, which allows you to select an application class that will define the product. You can map the product to an existing or new application class.
- Step 3** Select a product and click the **Ignore** button to have that product be “ignored” and not appear in reports. Undo an ignore setting by clicking the **View Ignored** button to launch a new window that allows you to “un-ignore” the product in question.
- Step 4** Click the **Save** button once you have selected the application class(es).

Configuring Unknown Product Association



This window displays a list of applications (processes) that have run on systems but have no product associated with them. (This is the inverse of the application deployment investigation Product Associations page.)

You can use application deployment investigation to generate reports that use installed software products as part of the report criteria. In order to do this, there is some prerequisite configuration.

This is necessary because the investigation process, in part, gathers data on systems according to the application name it finds. That is the application executable itself and not the product with which the application is associated. Application deployment investigation does gather information on the products that are installed on systems, but it does not then map the applications back to the installed products. For example, it may find that excel.exe is running on a system. It may also find that Microsoft Office is installed on that same system. But the analysis process will not know that excel.exe is part of Microsoft Office. You must tell it so.

Therefore, to generate certain report types using installed product information, you must first associate the installed products found by the investigation process with the application(s) that make up the product. (You can also associate a product with an existing application class from the Product Associations page.)

You must make this application/product association to use product criteria to generate the Product Usage report type.

Complete the following steps to create application/product associations after the data has been collected.

- Step 1** Select **Analysis > Deployment Investigation > Unknown Applications**. The Unknown Applications window contains a list of all the processes found on systems that were tracked but have no association with an installed product.

Step 2 To associate an application with a process, select the check box beside the Process Name link and click the **Map to Product** button. This opens a new window where you can select a product that will define the application process. You can also map the application to an existing or new application class.

Note You can only map and/or ignore products that have not yet been mapped.

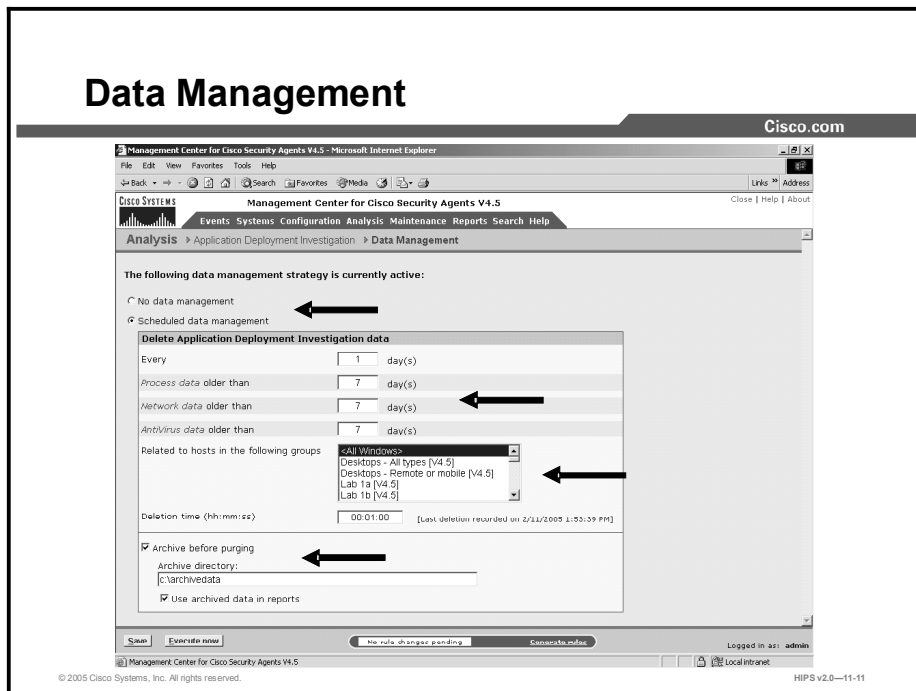
Step 3 Select a process and click the **Ignore** button to have that process be “ignored” and not appear in reports.

Step 4 You can undo an ignore setting by clicking the **View Ignored** button to launch a new window that allows you to “un-ignore” the process in question.

Step 5 Click **Save** once you select a product. The process will then disappear from the Unknown Applications list, as it is no longer unknown.

Data Management

This topic discusses data management and how to configure it.



Accessible from the Analysis > Application Deployment Investigation > Data Management menu, the Data Management window allows you to archive and purge the data collected by the deployment investigation.

Use the Data Management page to purge deployment investigation data at scheduled intervals and, optionally, to archive the data that you are deleting from the active database. This page gives you the option of having no scheduled data management (No Data Management radio button) or to set parameters for a scheduled purging of data (Scheduled Data Management radio button).

You can configure your data management to purge certain types of data at different time intervals as you choose. Process data, network data, and antivirus data can be purged according to the “day” interval that you set. Note that AntiVirus Data has been added as a separate category due to the large volume of this data type that can accumulate.

If you click the Execute Now button, you can trigger data management to occur immediately based on the current configuration, regardless of the data management type you have configured using the available radio buttons.

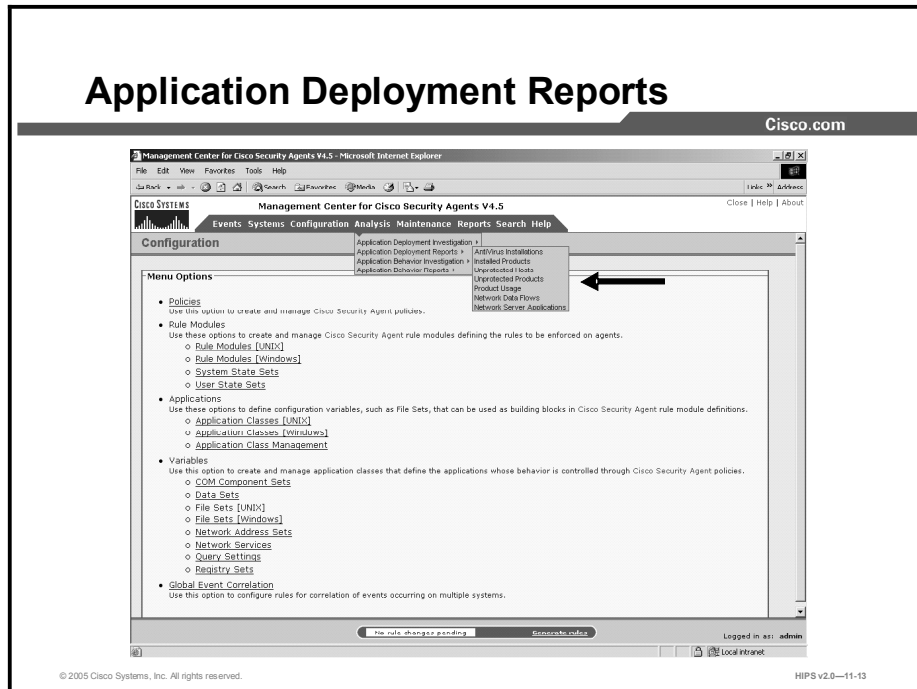
Select the Archive Before Purging check box and enter a directory in which to store archived data if you do not want to lose the report data that you are purging. You can continue to use this archived data in your reports.

Note If you change the archive directory after you have already archived data, that data is automatically moved to the new directory, and new archived data will be stored in the newly specified directory as well.

Note You can click the Archive History link at the top of this page to view an informational list of data purges that have taken place on the system.

Application Deployment Reports

This topic discusses application deployment reports and how to run them.



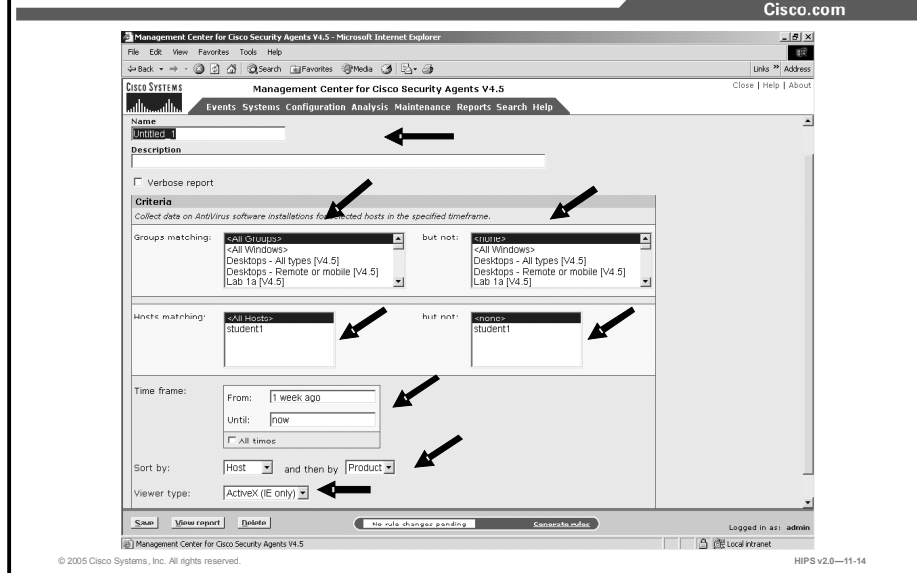
You can generate several different application deployment report types using the data gathered during the tracking process. The following sections describe each of these reports.

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.

Complete the following steps to generate an application deployment report:

- Step 1** Select **Analysis > Application Deployment Reports**. A drop-down list and a cascading menu of report types appears.
- Step 2** Select a report type from the cascading menu to enter parameters and generate that report.

AntiVirus Installations Report



Use AntiVirus Installations report type to view software version and signature version information for detected Norton and McAfee antivirus installations. (Note that for McAfee AntiVirus software, you will also see the engine version in the report.)

Complete the following steps to configure an AntiVirus Installations report:

- Step 1** Select **Analysis > Application Deployment Reports > AntiVirus Installations**. The Antivirus Installations Report window appears.
- Step 2** In the report window, click the **New** button to create a new report. This takes you to the report configuration view.
- Step 3** Enter a name in the Name field.
- Step 4** Enter a description in the Description field.
- Step 5** Optionally, enable the Verbose Report check box. If you do not enable this check box and you have selected to generate a report for <All Groups> and <All Hosts>, you will only see summary information (number of overall installation copies found) for installed antivirus products. If you enable Verbose, you will see a much longer report containing details of installed antivirus products on each host by hostname.

AntiVirus Installations nonverbose reports contain the following data:

- Antivirus product name, product version, engine and signature version, and the number of hosts running this combination.

AntiVirus Installations verbose reports contain the following data:

- Host name, product version, engine version, signature version, and time this information was obtained.

- Step 6** From the Groups Matching field, you can select a specific group for which to generate antivirus installation information. You can view information for <All Groups> or only for those you select.
- Step 7** Optionally, use the But Not field to exclude certain groups from those selected in the Groups Matching field.
- Step 8** From the Hosts Matching field, you can select hosts within the selected group(s) for which to generate antivirus installation information.
- Step 9** You can view information for <All Hosts> or only for those you select by using the But Not field to exclude certain hosts from those selected in the Hosts Matching field. Using exclusions, you can generate a report for a specific host within a selected group.

Note Individual hosts do not appear in the Hosts report field until they have uploaded data at least once.

- Step 10** Enter a time frame by which to view the collected data. This time indicates the last or most recent time the antivirus product was used on the system(s) in question.

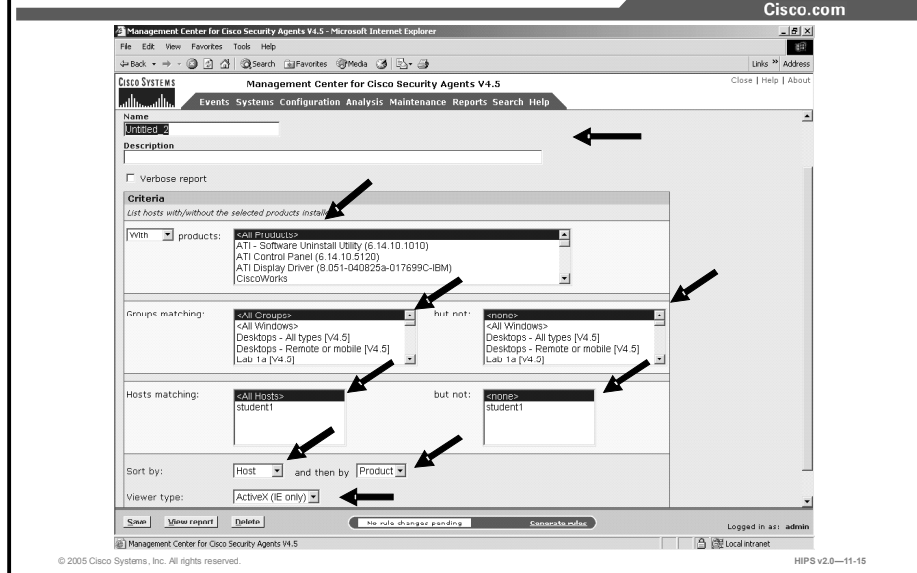
You can enter From and Until time parameters using the syntax described next, or you can check the All Times check box for all time frames.

Time syntax:

You can specify a relative time 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.

- Step 11** Select the criteria by which to sort the report. You can first sort by host and then by product or vice versa.
- Step 12** Select a viewer type. By default, ActiveX is selected. This is the recommended viewer. Click the **Save** button to save the parameters that you have just configured for generating this report.
- Step 13** Click the **View Report** button and the report is automatically displayed in a new window.

Installed Products Report



Use the Installed Products report type to view a list of products that are installed or not installed on various selected host machines. The products listed alphabetically in the report page are the software programs found to be installed (or not installed) on the systems that were analyzed. These are software programs that are visible in the Add/Remove Programs window.

Note This report provides only the latest installed product information. It does not provide any historic data on installed products. Therefore, there is no time range available in this report.

Complete the following steps to configure and run an Installed Products report.

- Step 1** Select **Analysis > Application Deployment Reports > Installed Products**. Any preconfigured reports will show.
- Step 2** In the report window, click the **New** button to create a new report. This takes you to the report configuration view.
- Step 3** Enter a name in the Name field.
- Step 4** Enter a description in the Description field.
- Step 5** Optionally, enable the Verbose Report check box. If you do not enable this check box and you have selected to generate a report for <All Groups> and <All Hosts>, you will only see summary information (number of overall installation copies found) for installed products. If you enable Verbose, you will see a much longer report containing details for installed products on each host by hostname.

Installed Products nonverbose reports contain the following data:

- Distinct product name and the overall number of hosts that have this product installed.

Installed Products verbose reports contain the following data:

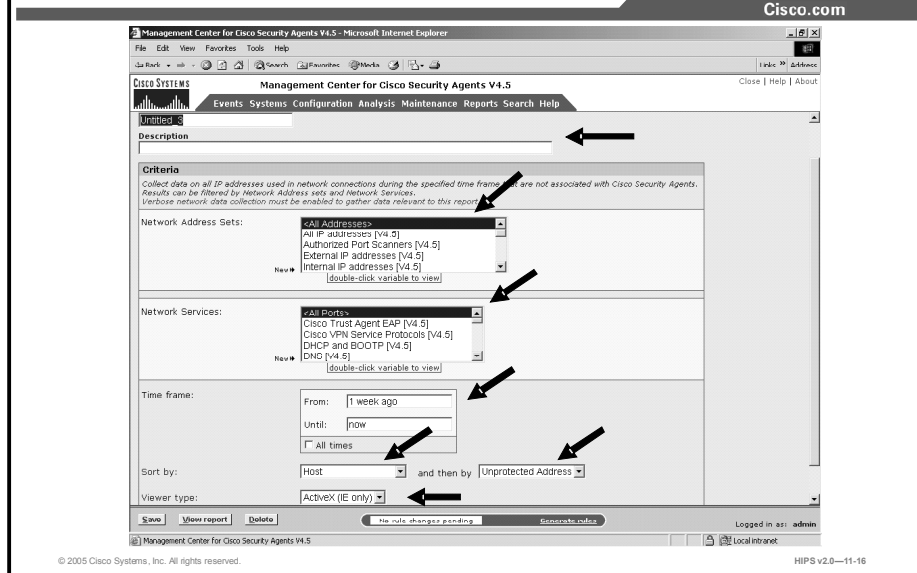
- Distinct product name and the individual hosts that have this product installed.

- Step 6** If you are creating a report of products not installed on the system(s) in question, select “List hosts **without** the selected product installed.” If this is a report on products installed on selected hosts, leave the default choice of **with** in the drop-down view.
- Step 7** From the Products list field, you can select one or more products and view which hosts and/or groups have that product installed (or not installed) on their system (verbose). You can also select <All Products> depending on the type of report you wish to generate.

Note You do not have to associate products with application classes to run this report type.

- Step 8** From the Groups Matching field, you can select a specific group for which to generate product installation information. You can view information for <All Groups> or only for those you select. Optionally, use the But Not field to exclude certain groups from those selected in the Groups Matching field.
- Step 9** From the Hosts Matching field, you can select hosts within the selected group(s) for which to generate product installation information. You can view information for <All Hosts> or only for those you select by using the But Not field to exclude certain hosts from those selected in the Hosts Matching field. Using exclusions, you can generate a report for a specific host within a selected group.
- Step 10** Select criteria by which to sort the report. You can first sort by host and then by product or vice versa.
- Step 11** Select a viewer type. By default, ActiveX is selected. This is the recommended viewer. Click the **Save** button to save the parameters that you have just configured for generating this report.
- Step 12** Click the **View Report** button and the report is automatically displayed in a new window.

Unprotected Hosts Report



Use the Unprotected Hosts report type to view hosts that are being used in network connections, but are not protected by Cisco Security Agents.

This report type uses network address sets and network services for filtering criteria.

Complete the following steps to configure and run an Unprotected Hosts report:

- Step 1** Select **Analysis > Application Deployment Reports > Unprotected Hosts**. The report window appears with any preconfigured reports showing.
- Step 2** In the report window, click the **New** button to create a new report. This takes you to the report configuration view.
- Step 3** Enter a name and description for the report.
- Step 4** From the Network Address Sets field, select a preconfigured network address set. You can view information for <All Addresses> or only for those you select.

Note You can create a new network address set or edit an existing one from this page by clicking the New link or by double-clicking an item in the selection field.

Note From the Network Services list field, select a preconfigured network service. You can view information for <All Ports> or only for those you select.

Note You can create a new network service or edit an existing one from this page by clicking the New link or by double-clicking an item in the selection field.

Step 5 Enter a time frame by which to view the collected data. You can enter From and Until time parameters using the syntax described here, or you can check the All Times check box for all time frames.

Time syntax:

- You can specify a relative time 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 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.

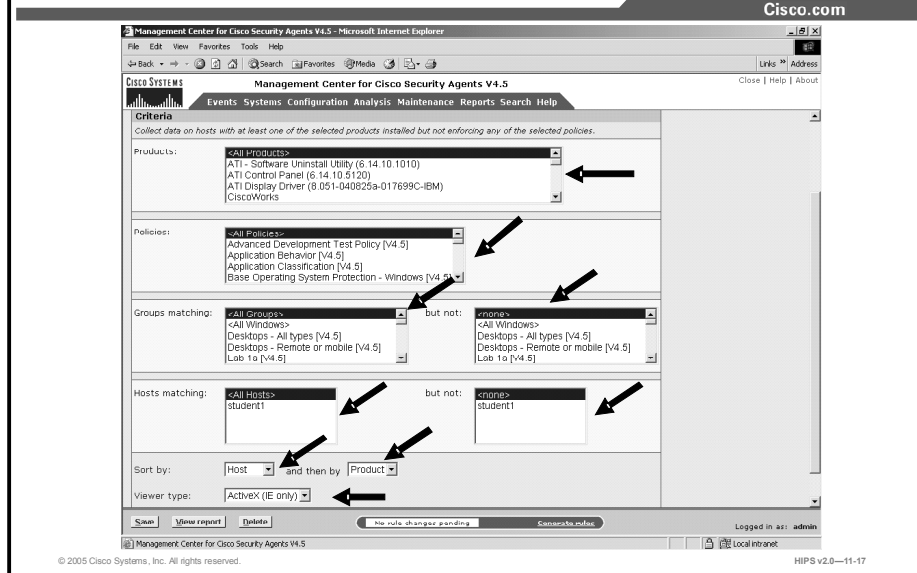
Step 6 Select criteria by which to sort the report. You can sort by operation, host, unprotected, address, or protocol.

Step 7 Select a viewer type. By default, ActiveX is selected. This is the recommended viewer.

Step 8 Click the **Save** button to save the parameters that you have just configured for generating this report.

Step 9 Click the **View Report** button and the report is automatically displayed in a new window.

Unprotected Products Report



Use the Unprotected Products report type to view hosts with products installed that have no associated Cisco Security Agent policies (that is, hosts running products for which there is no deployed policy).

Complete the following steps to configure and run an Unprotected Products report:

- Step 1** Select **Analysis > Application Deployment Reports > Unprotected Products**. The report window appears with any preconfigured reports showing.
- Step 2** In the report window, click the **New** button to create a new report. This takes you to the report configuration view.
- Step 3** Enter a name in the Name field.
- Step 4** Enter a description in the Description field.
- Step 5** From the Groups Matching field, you can select a specific group for which to generate unprotected product information. You can view information for <All Groups> or only for those you select.
- Step 6** Optionally, use the But Not field to exclude certain groups from those selected in the Groups Matching field.
- Step 7** From the Hosts Matching field, you can select hosts within the selected group(s) for which to generate report information. You can view information for <All Hosts>.
- Step 8** Optionally, you can view information only for those you select by using the But Not field to exclude certain hosts from those selected in the Hosts Matching field. Using exclusions, you can generate a report for a specific host within a selected group.
- Step 9** From the Products list field, you can select one or more products and view which hosts and/or groups have used that product on their system (verbose) but have no policy for that product enforced. You can also select <All Products> depending on the type of report that you wish to generate.

Note You must first associate products with application classes to run this report type.

Step 10 From the Policies list field, you can select one or more policies (preferably a policy that you know enforces rules for the product also selected for this report) and view which hosts and/or groups have that policy enforced on their system. You can also select <All Policies> depending on the type of report you wish to generate.

Note You must have some policies defined to run this report type.

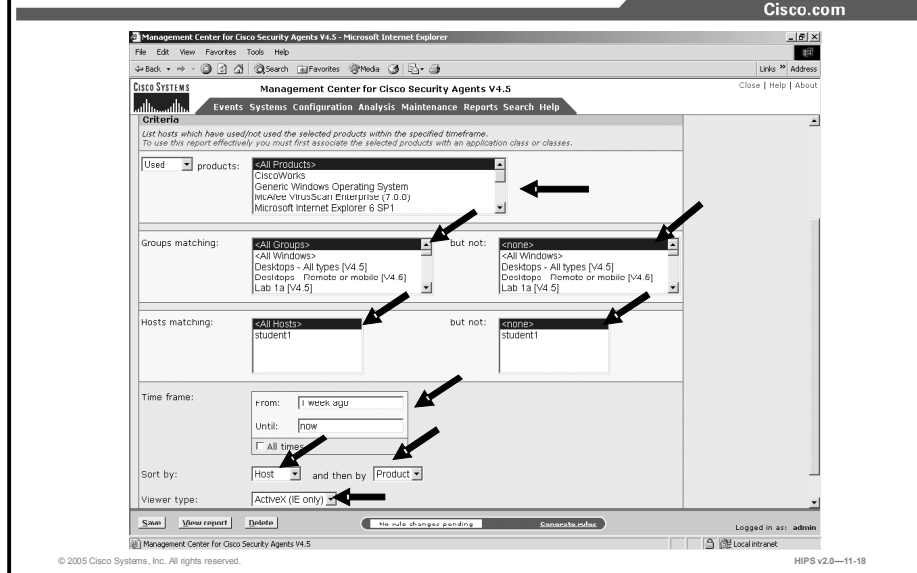
Step 11 Select criteria by which to sort the report. You can first sort by host and then by product or vice versa.

Step 12 Select a viewer type. By default, ActiveX is selected. This is the recommended viewer.

Step 13 Click the **Save** button to save the parameters that you have just configured for generating this report.

Step 14 Click the **View Report** button and the report is automatically displayed in a new window.

Product Usage Report



Use the Product Usage report type to view the number of systems on which installed products are used or not used.

Note In order to generate this report type, you must first associate products (all or just the particular ones that you are interested in) with an application class or classes.

Complete the following steps to configure and run a Product Usage report:

- Step 1** Select **Analysis > Application Deployment Reports > Product Usage**. The Product Usage Report page appears with any existing reports shown.
- Step 2** In the report window, click the **New** button to create a new report. This takes you to the report configuration view.
- Step 3** Enter a name in the Name field.
- Step 4** Enter a description in the Description field.
- Step 5** Optionally, enable the Verbose Report check box. If you do not enable this check box and you have selected to generate a report for <All Groups> and <All Hosts>, you will only see summary information displaying the overall number of systems each product is used on. If you enable Verbose Report, you will see a much longer report containing details for product usage on each host by hostname.

Product Usage nonverbose reports contain the following data:

- Product name and the overall number of hosts that have used the product.

Product Usage verbose reports contain the following data:

- Product name and the individual name of the host(s) that have used the product.

- Step 6** If you are running a report to determine which products are not used on systems, select “List hosts which have **not used** the selected products within the specified time” in the options drop-down list. Otherwise, leave the default choice of **used** selected.
- Step 7** From the Products list field, you can select one or more products and view which hosts and/or groups have used that product on their system (verbose). You can also select <All Products> depending on the type of report you wish to generate.

Note You must first associate products with application classes to run this report type.

- Step 8** From the Groups Matching field, you can select a specific group for which to generate product usage information. You can view information for <All Groups> or only for those you select.
- Step 9** Optionally, use the But Not field to exclude certain groups from those selected in the Groups Matching field.
- Step 10** From the Hosts Matching field, you can select hosts within the selected group(s) for which to generate report information. You can view information for <All Hosts>.
- Step 11** You can optionally view information only for those that you select by using the But Not field to exclude certain hosts from those selected in the Hosts Matching field. Using exclusions, you can generate a report for a specific host within a selected group.
- Step 12** Enter a time frame by which to view the collected data. This time indicates when the product was used on the system(s) in question. You can enter From and Until time parameters using the syntax described here, or you can check the All Times check box for all time frames.

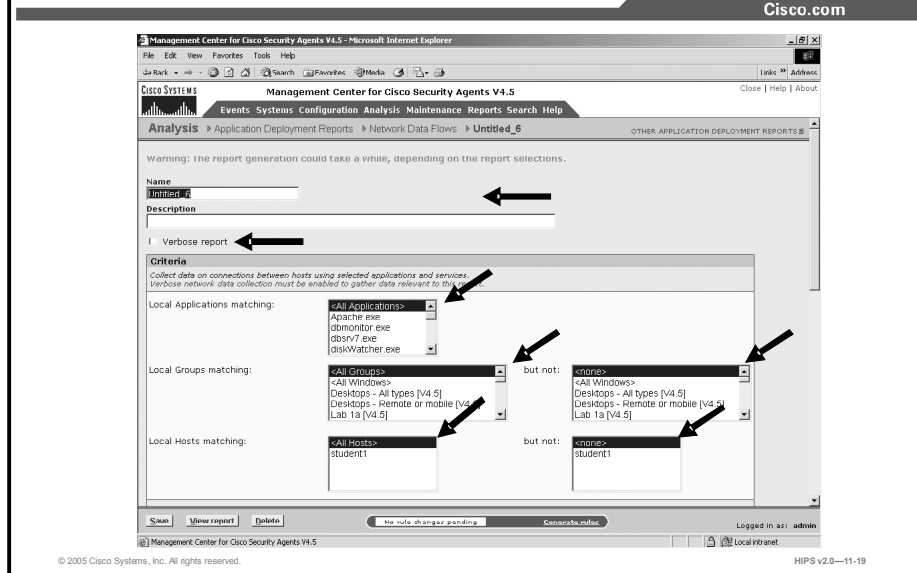
Time syntax:

You can specify a relative time 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). 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? (The question marks here indicate that the information inside the question marks is optional.) The default year is the current year.

- Step 13** Select criteria by which to sort the report. You can first sort by host and then by product or vice versa.
- Step 14** Select a viewer type. By default, ActiveX is selected. This is the recommended viewer.
- Step 15** Click the **Save** button to save the parameters that you have just configured for generating this report.
- Step 16** Click the **View Report** button and the report is automatically displayed in a new window.

Network Data Flows Report



Use the Network Data Flows report type to view, by network service, the number of data flows (unique source/destination address combinations), the number of hosts acting as clients, and the number of hosts acting as servers. This data can be filtered by protocol, source address set, and destination address set. You could use the results of this report to constrain a host's communication to only those hosts that it typically talks to.

Note Verbose network data collection must be enabled to gather data relevant to this report.

Complete the following steps to configure and run a Network Data Flows report:

- Step 1** Select **Analysis > Application Deployment Investigation Reports > Network Data Flows**. The Network Data Flows report page is displayed with any existing reports showing.
- Step 2** In the report window, click the **New** button to create a new report. This takes you to the report configuration view.
- Step 3** Enter a name in the Name field.
- Step 4** Enter a description in the Description field.
- Step 5** Optionally, enable the Verbose Report check box. If you do not enable this check box and you have selected to generate a report for <All Groups> and <All Hosts>, you will only see summary information displaying the overall number of data flows rather than data flows per host. If you enable Verbose Report, you will see a much longer report containing details for hosts, source and destination addresses, protocols, and client/server connections.

Network Data Flows nonverbose reports contain the following data:

- Unique protocol/port combinations, unique combination of source IP address, destination IP address (including address resolved to hostname whenever

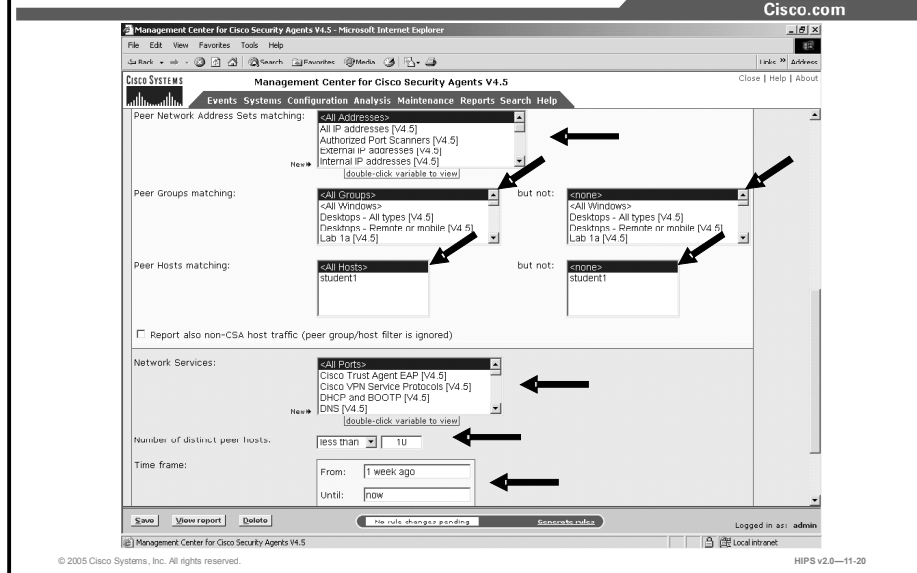
possible), and number of incoming and outgoing connections between the source/destination combination in the specified time frame.

Network Data Flows verbose reports contain the following data:

- Local host, local IP address, local process name, network operation, peer host, peer IP address, and number of network requests with the distinct combination of all items mentioned.

- Step 6** From the Applications list field, you can select one or more applications with which to filter this report. You can also select <All Applications> depending on the type of report you wish to generate.
- Step 7** From the Local Groups Matching field, you can select a specific group for which to generate network data flow information. You can view information for <All Groups> or only for those you select.
- Step 8** Optionally, use the But Not field to exclude certain groups from those selected in the Groups Matching field.
- Step 9** From the Local Hosts Matching field, you can select hosts within the selected group(s) for which to generate report information. You can view information for <All Hosts>.
- Step 10** You can view information only for those hosts you select by using the But Not field to exclude certain hosts from those selected in the Hosts Matching field. Using exclusions, you can generate a report for a specific host within a selected group.

Network Data Flows Report (Cont.)



- Step 11** From the Peer Network Address Sets Matching list field, select a preconfigured address. You can view information for <All Addresses> or only for those you select.
- Step 12** You can create a new network address set or edit an existing one from this page by clicking the New link or by double-clicking an item in the selection field.
- Step 13** From the Peer Groups Matching field, you can select a specific peer group for which to generate network data flow information. You can view information for <All Groups> or only for those you select.
- Step 14** Optionally, use the But Not field to exclude certain groups from those selected in the Groups Matching field.
- Step 15** From the Peer Hosts Matching field, you can select a specific peer host for which to generate network data flow information. You can view information for <All Hosts> or only for those you select.
- Step 16** Optionally, use the But Not field to exclude certain hosts from those selected in the Hosts Matching field.
- Step 17** Optionally, enable the Report Also Non-CSA Host Traffic (Peer Group/Host Filter Is Ignored) check box. This will produce a much longer report and will ignore any peer settings that you may have configured.
- Step 18** From the Network Service list field, select a preconfigured network service. You can view information for <All Ports> or only for those you select.
- Step 19** You can create a new network service or edit an existing one from this page by clicking the New link or by double-clicking an item in the selection field.
- Step 20** For Number of Distinct Peer Hosts, enter a number by which to filter this report.

- Step 21** Enter a time frame by which to view the collected data. You can enter From and Until time parameters using the syntax described here, or you can check the All Times check box for all time frames.

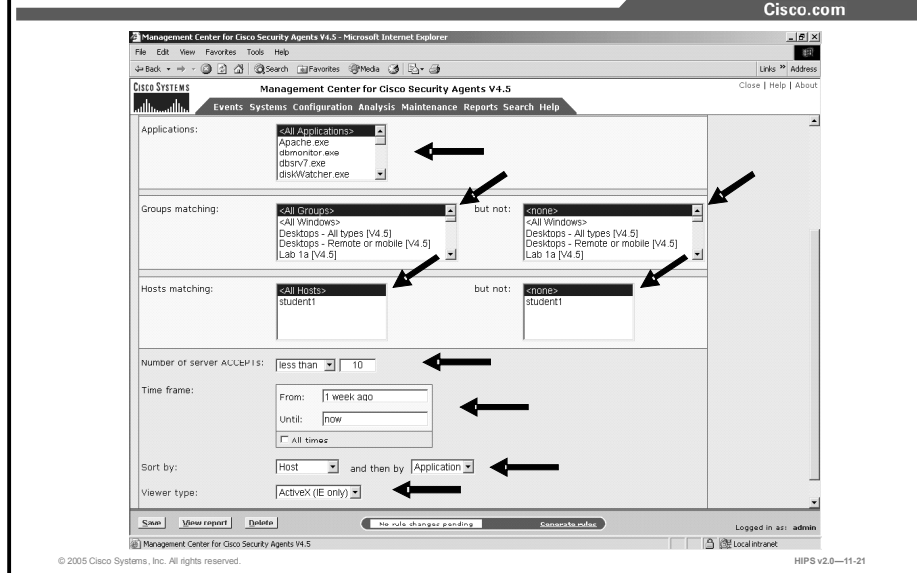
Time syntax:

You can specify a relative time 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? (The question marks here indicate that the information inside the question marks is optional.) The default year is the current year.

- Step 22** Select criteria by which to sort the report. You can first sort by host and then by application or vice versa.
- Step 23** Select a viewer type. By default, ActiveX is selected. This is the recommended viewer.
- Step 24** Click the **Save** button to save the parameters that you have just configured for generating this report.
- Step 25** Click the **View Report** button and the report is automatically displayed in a new window.

Network Server Applications Report



The Network Server Applications report is intended to break down network server application activity on a given set of hosts. You could use this report type to view which network server applications are listening on ports but not accepting any (or very few) connections. You could also use this to determine which are the most active web servers or database servers on your network.

Complete the following steps to configure and run a Network Server Applications report:

- Step 1** Select **Analysis > Application Deployment Reports > Network Server Applications**. The Network Server Applications report page appears with any existing reports showing.
- Step 2** In the report window, click the **New** button to create a new report. This takes you to the report configuration view.
- Step 3** Enter a name in the Name field.
- Step 4** Enter a description in the Description field.
- Step 5** From the Applications list field, you can select one or more applications that hosts and/or groups use to listen on the network. You can also select <All Applications> depending on the type of report that you wish to generate.
- Step 6** From the Groups Matching field, you can select a specific group for which to generate unused network application information. You can view information for <All Groups> or only for those you select.
- Step 7** Optionally, use the But Not field to exclude certain groups from those selected in the Groups Matching field.
- Step 8** From the Hosts Matching field, you can select hosts within the selected group(s) for which to generate report information. You can view information for <All Hosts>.

- Step 9** You can view information only for those you select by using the But Not field to exclude certain hosts from those selected in the Hosts Matching field. Using exclusions, you can generate a report for a specific host within a selected group.
- Step 10** For the Number of Server ACCEPTs, enter the maximum number. By default, this field has 0 entered. Use this number to find network listens with no or very few subsequent network connections.
- Step 11** Enter a time frame by which to view the collected data. This time indicates when the network listen/connection was seen on the system(s) in question. You can enter From and Until time parameters using the syntax described here, or you can check the All Times check box for all time frames.

Time syntax:

You can specify a relative time 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.

- Step 12** Select criteria by which to sort the report. You can first sort by port, host, and then by application or vice versa.
- Step 13** Select a viewer type. By default, ActiveX is selected. This is the recommended viewer.
- Step 14** Click the **Save** button to save the parameters that you have just configured for generating this report.
- Step 15** Click the **View Report** button and the report is automatically displayed in a new window.

Viewing Reports

When you generate your reports, you are given the option of selecting the type of viewer through which to display the report. From the Viewer Type drop-down 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).
- **HTML Frame:** 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 viewer type you have selected and the printer settings on the printer.

Exporting Reports

When you export analysis reports, you are only exporting the report itself and the names of objects referenced in the report, such as a group or policy. The group and policy objects

themselves are not automatically exported with the report. There is no cascading inheritance when you export either reports or event sets as occurs when other items are exported. You must select groups and policies separately if you want to export them for the purpose of the report.

Application Behavior Investigation

This topic introduces application behavior investigation.

Application Behavior Investigation

Cisco.com

Cisco Security Agent application behavior investigation 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.

© 2005 Cisco Systems, Inc. All rights reserved.HIPS v2.0—11-23

Cisco Security Agent application behavior investigation 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 make up CSA MC policies are application-centric, understanding the resources that applications require for normal operations is integral to building effective policies. Behavior investigation does that by analyzing applications as they operate in a normal environment and generating useful reports and rule modules (rule module creation is a separately licensed feature) based on that analysis.

When deployed on a system running a Cisco Security Agent, application behavior investigation monitors the actions of designated applications on that system, logging all resource access attempts made by the application. It then analyzes the logging data that it collects and develops detailed reports for the application in question. It also, optionally, generates a rule module. The generated rule module 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.

Application Behavior Investigation Process

Cisco.com

The application behavior investigation is performed by three different contributing components:

- **CSA MC**
- **The agent (logging agent)**
- **The behavior investigation functionality**

© 2005 Cisco Systems, Inc. All rights reserved.

HIPS v2.0—11-24

The application behavior investigation is performed by three different contributing components: CSA MC, the agent (logging agent), and the behavior investigation functionality.

- Through *CSA MC*, you designate which application you want to investigate. You also select an agent host on which the investigation is to take place and a time frame within which the investigation will be completed. This investigation configuration is then sent to the agent on the selected host in the same way that policies are sent to agents.
- *Application behavior investigation* examines all the logged data that 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 that are 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 that are 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 behavior investigation function for processing.

Optionally, CSA MC imports the rule module created by the behavior investigation.

Behavior Analysis

Cisco.com

When you are ready to configure a behavior analysis for an application, you must have the following information:

- **What application you want to analyze**
- **Which host you want to select for application analysis**

© 2005 Cisco Systems, Inc. All rights reserved.

HIPS v2.0—11-25

When you are ready to configure a behavior analysis for an application, you must have the following information:

- **What application you want to analyze:** 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 a time.)
- **Which host you want to select for application analysis:** You should have an appropriate host chosen for the behavior analysis.

Creating, Saving, and Canceling Behavior Analysis

Cisco.com

Available buttons and links are as follows:

- **New**
- **Delete**
- **Clone**
- **Save**
- **Stop Logging**
- **Start Analysis**
- **Optional Import**

© 2005 Cisco Systems, Inc. All rights reserved.

HIPS v2.0—11-26

Similar to most CSA MC windows, behavior analysis action items appear in a frame at the bottom of CSA MC.

The available buttons in the bottom frame change in accordance with the actions that are available for the page that you are viewing. With a behavior analysis, several actions are performed from the same page as the behavior analysis progresses. You may have to refresh the behavior analysis page for the buttons to change appropriately.

Available buttons and links are as follows:

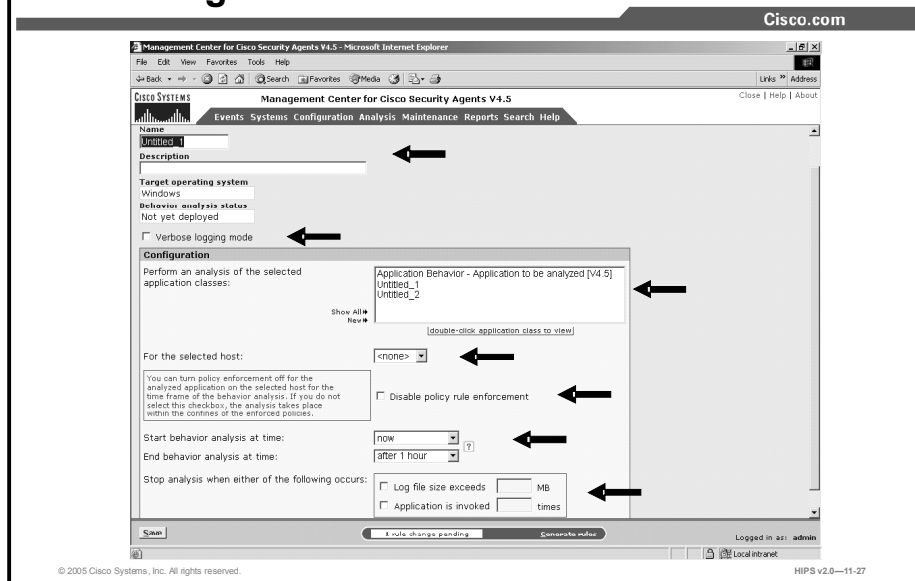
- **New:** Use the New button to create a new configuration item within the list view that you have selected. Click the New button and a new item appears in the list view. Click the new item link to access the configuration view for that item.
- **Delete:** Use the Delete button in conjunction with the check boxes beside each list view item. To delete a configuration, select its check box (you can select several at once) and click the Delete button. All checked items are deleted. To quickly select all check boxes, click the very top check box in the list view heading bar. Clicking the Delete button then deletes all items.
- **Clone:** Use the Clone button in conjunction with the check boxes beside each list view item. To clone a particular configuration, select its check box 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 analysis early and send collected data to the workstation for analysis, click this Stop Logging button.
- **Start Analysis:** When the logging for the analysis is complete, a Start Analysis button appears in the bottom frame of the behavior analysis page. Click this button to have the analysis workstation begin to analyze the logging data.
- **Optional Import:** When the analysis of the logging data is complete, the behavior analysis creates a policy that you can import into CSA MC. The Import button appears when the policy creation is complete if you have a license for analysis policy creation and import.

Configuring a Behavior Analysis Investigation



Complete the following steps to configure a behavior analysis investigation:

Note In some cases, you can configure a behavior analysis investigation using the Event Management Wizard, which accessible from particular Event Log entries.

Step 1 Select **Analysis > Application Behavior Investigation > Behavior Analyses** (Windows or UNIX). The list of existing analyses (if any) is displayed.

Step 2 Click the **New** button to create a new behavior analysis. This takes you to the behavior analysis configuration page.

Step 3 Enter a name in the Name field for the behavior analysis you are creating.

Step 4 Enter a description in the Description field for your behavior analysis. This description becomes visible in the behavior analysis list view.

Step 5 Select the **Verbose Logging Mode** check box: By default, behavior analysis filters its logging process so that duplicate events are not logged. You can turn this feature off by selecting this check box. 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 behavior analysis reports.

Note The target operating system that you select is displayed in a read-only field. The Behavior Analysis Status field is also a read-only field. It displays text, informing you of each stage of the analysis. When you first configure your behavior analysis, it displays "Not yet deployed."

Step 6 In the Perform An Analysis of the Selected Application Classes list box, select the application class or classes that you want to analyze. To select multiple items in a list box, hold down the Control key as you select each item. To deselect a single

item, hold down the Control 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.

Note You can select an application class that contains more than one application for the analysis. But in that case, the reports created would apply equally to all applications included in the analyzed application class. For example, if the application class that you are analyzing contains both Microsoft Word and Microsoft Outlook, the reports created by the behavior analysis would be a combination of the resources required by both applications.

Step 7 Select the host that you are assigning the behavior analysis to in the For the Selected Host list box. You cannot have more than one behavior analysis running on a host at one time.

Note Once the behavior analysis begins, you can click the Stop Logging button that appears in the bottom frame. The behavior analysis stops automatically according to the parameters that you enter on this behavior analysis page. But if you want to stop the analysis early and send collected data to the workstation for analysis, click this Stop Logging button.

Step 8 Optionally, you can select the check box for Disable Policy Rule Enforcement for the time frame of the analysis. 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 check box, when the logging agent receives a behavior analysis investigation, any policies relevant to the application being analyzed are disabled on the selected host until the analysis is completed. This may be undesired if the application in question is unknown or is in any way suspicious.

Step 9 From the Start Behavior Analysis at Time drop-down menu, select a time for the behavior analysis to start once the host polls in and receives the behavior analysis. If you specify no time here, "now" is automatically entered. This means that the behavior analysis will start immediately when the host receives it.

Step 10 From the End Behavior Analysis at Time drop-down list, you must select a time for the behavior analysis to end. The behavior analysis process will not allow you to save the analysis until you do.

Step 11 To stop behavior analysis according to specified criteria, select either or both of the following:

- **Log File Size Exceeds __ MB:** You can enter a size restriction on the log file. When it reaches the size that you indicate, the analysis is finished. The maximum log file size that 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 that you indicate, the analysis is finished.

Note Using an invocation number limit is not always appropriate. For example, for server applications, time frame parameters might be a more appropriate criteria for ending a behavior analysis.

Note If you enter analysis completion parameters in more than one field, the parameter that is reached first is the one that applies.

Step 12 Click the **Save Behavior Analysis** button in the bottom frame of CSA MC to save it.

Step 13 Once your behavior analysis 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 behavior analysis to the group hosts that you have selected. Depending on the behavior analysis parameters that you have configured, the selected host will begin the behavior analysis after it polls in to CSA MC and receives the new rules.

Note Keep in mind that if you have configured your behavior analysis to begin immediately and your agents are configured to poll in to CSA MC once every hour, the behavior analysis 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 behavior analysis. In this case, the analysis will not run at all.

Monitoring the Behavior Analysis

You can check your CSA MC Event Log to view the behavior analysis progression. An event is sent when the behavior analysis begins and again when it finishes.

You can also monitor Progress Status fields in the Behavior Analysis configuration page. These fields appear when the analysis is in progress. You can monitor the size of the log file, and if you have set an application invocation limit, you can monitor the number of application innovations as well. These progress fields update each time the logging agent polls in to the MC.

When reports and the policy are ready to be imported to CSA MC, an Event Log message appears indicating this.

Start Behavior Analysis

When the Event Log in CSA MC displays "Log files for behavior analysis were sent to the analysis workstation," you can begin the data analysis of the logging information.

Begin this analysis by accessing the behavior analysis window for this particular analysis and clicking the Start Analysis button in the bottom frame. This begins the analysis. An Event Log message appears, informing you that "Data analysis has started."

When the analysis is complete, you can view reports.

If you have a license for rule module creation, when the analysis is complete, the Event Log file displays the message "Rule module creation for behavior analysis completed successfully." Once rule module creation is complete, you can import the module.

Importing the Rule Module

Note If you do not have a separate license for importing behavior analysis rule modules, the behavior analysis results in creation of a report without the added step of creating a rule module.

When the behavior analysis has completed its analysis of the logging data, the rule module that it created is ready to be imported into CSA MC.

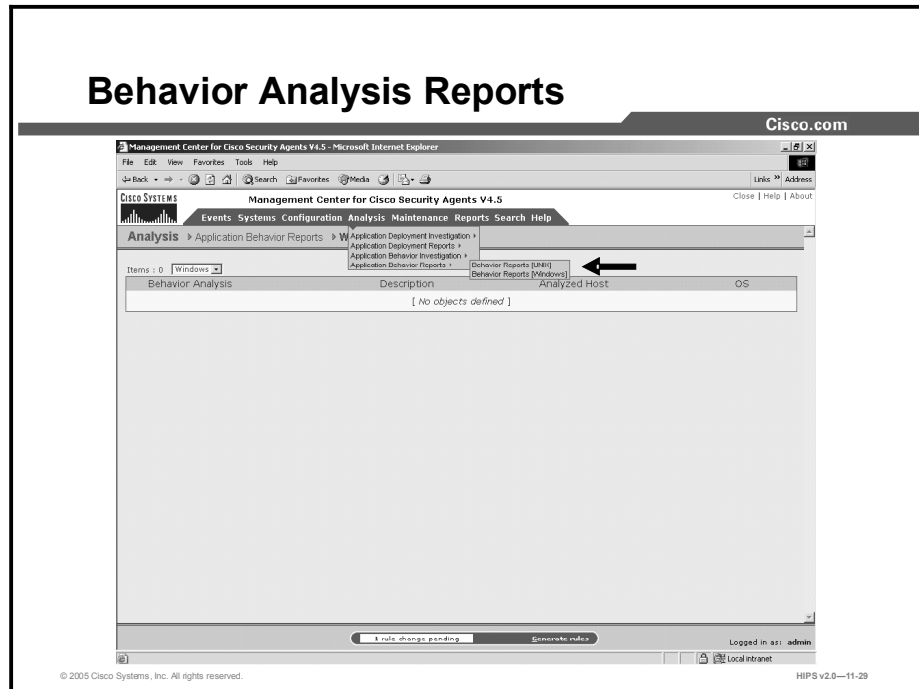
Import the rule module by once again accessing the behavior analysis window for this particular analysis. Click the **Import** button in the bottom frame. (This button only appears when the rule module is ready for importing.)

Note The rule module and its accompanying "variables" are imported into CSA MC. The behavior analysis creates its own variables for use in the rules that it also creates.

Note In order to deploy the rule module that the behavior analysis has created, you must associate it with an existing policy or with a new policy that you create. This policy must be attached to a group for the rules to be deployed to hosts.

Behavior Analysis Reports

This topic describes behavior analysis reports.



During the analysis process, the behavior analysis sorts the logging data that it receives from the logging agent into categorized reports. You can view these reports on the CSA MC system by accessing the Analysis > Application Behavior Reports > Behavior Report (Windows or UNIX).

Reports on specific analyses only become available once the behavior analysis has successfully completed. The CSA MC Event Log displays a message to inform you that reports have been created.

Report Components

When you access the application behavior reports window, you can view individual reports for all completed analysis from the same window by selecting a particular behavior analysis from the Reports for Behavior Investigation drop-down 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 behavior analysis logging session. Each report category has several subtopics that 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 that was accessed during the analysis time frame. If you select to view Behavior Summary for a report category, additional views further sort the information that the behavior analysis has collected by time frame, individual resource (for example, 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.

File Event Reports

Cisco.com

File event 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 the following information:

- **Time**
- **Directory**
- **File type**
- **Operation**
- **Process name**
- **Number of events**

© 2005 Cisco Systems, Inc. All rights reserved.

HIPS v2.0—11-30

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 the following information:

- **Time:** This is 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 that the event in question occurred during the logging period.

Registry Event Reports (Windows Only)

Cisco.com

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 the following information:

- Time
- Key name
- Value name
- PID
- Process name
- Number of events

© 2005 Cisco Systems, Inc. All rights reserved.

HIPS v2.0—11-31

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 the following information:

- **Time:** This is 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)

Cisco.com

COM reports display information on the COM class ID that was accessed and the process that made the request. More specifically, they provide the following information:

- **Time**
- **Object name**
- **PID**
- **Process name**
- **Number of events**

© 2005 Cisco Systems, Inc. All rights reserved.

HIPS v2.0—11-32

COM reports display information on the COM class ID that was accessed and the process that made the request. More specifically, they provide the following information:

- **Time:** This is 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

Cisco.com

Network reports display details such as the protocol that is accessing the network, the source and destination addresses of the connection, and the source and destination ports. More specifically, they provide the following information:

- Time
- Role
- Protocol
- Source address
- Source port
- Destination address
- Destination port
- PID
- Process name
- Number of events

© 2005 Cisco Systems, Inc. All rights reserved.

HIPS v2.0—11-33

Network reports display details such as the protocol that is accessing the network, the source and destination addresses of the connection, and the source and destination ports. More specifically, they provide the following information:

- **Time:** This is 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 the address from which the connection originated 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

Cisco.com

Summary reports display the number of times that each resource type was accessed during the logging time frame.

© 2005 Cisco Systems, Inc. All rights reserved.

HIPS v2.0—11-34

Summary reports display the number of times that each resource type was accessed during the logging time frame.

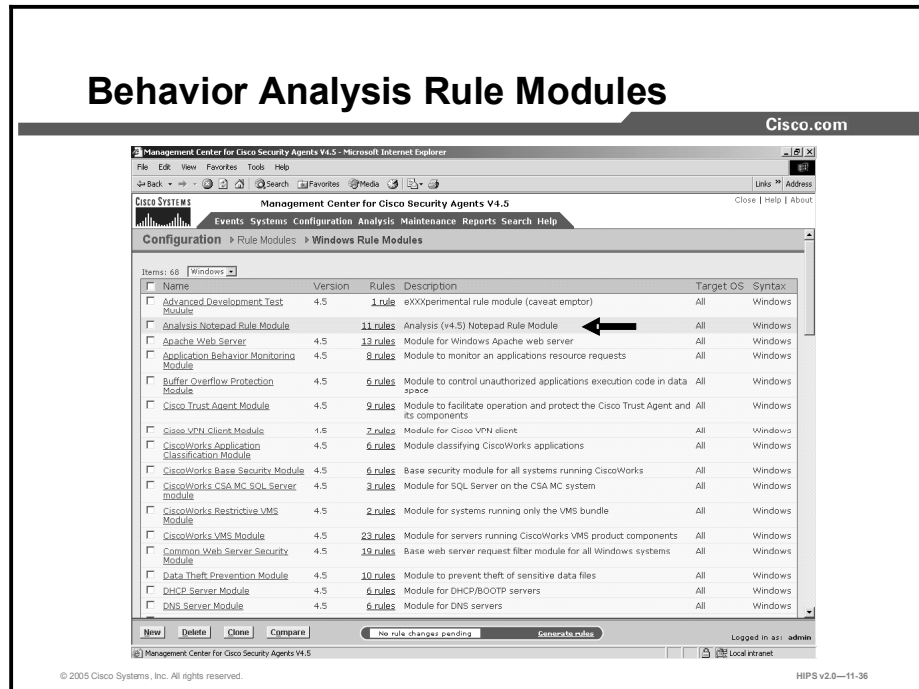
Working with Reports

Behavior analysis 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 **Control** key 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.

Behavior Analysis Rule Modules

This topic describes the behavior analysis rule modules.



Once imported, the behavior analysis rule module is added to your list of rule modules (Windows or UNIX) with the word "Analysis" appended to the original behavior analysis name. For example, if the analysis name is "Notepad," the name of the policy would be "Analysis Notepad Rule module."

Reviewing the Rule Module

The rule modules created by the behavior analysis process enforce normal application behavior and maintain application and system integrity. To achieve this, the general strategy behind the creation of behavior analysis rule modules is to protect the application from the system and to protect the system from the application.

As with all new rule modules that you create, you should review the rules that are generated by the behavior analysis and run the module in test mode for some period of time to ensure that it works as intended. You should also review the reports that are generated during the analysis, as they are valuable resources for understanding the application as well as the rule module.

Note Behavior analysis does not add system hardening or global correlation "built-in" rules to the policy. For example, you can add system API control to the policy.

Behavior Analysis Methodology

Cisco.com

- **Protecting the application from the system**
- **Protecting the system from the application**

© 2005 Cisco Systems, Inc. All rights reserved.

HIPS v2.0—11-37

Protecting the application from the system:

As part of the rule module, the behavior analysis 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 module.

Protecting the system from the application:

Resources that are accessed by the application are broken down into file, network, registry, and COM categories, and then rules for each category are created by the behavior analysis. Allow rules permit what was seen as normal application behavior, while deny rules prevent access to all resources that were not used by the application during the logging period.

Because security requirements may vary from site to site, the behavior analysis generates several rules that are disabled by default. The disabled rules are generally network and registry restrictions. The behavior analysis 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 the behavior analysis in combination with the Sample Network (Permissive, Selective, and Restrictive) policies shipped with the CSA MC.

If you decide to edit behavior analysis rule modules based on your site's requirements, the reports generated during the logging analysis process contain information on all the resources that are 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.

The general methodology behind the creation of rules for each resource type is as follows:

- **File access control rules:** The behavior analysis 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. It 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):** The behavior analysis 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 that were not used during the logging period. Use COM Object Summary reports to help refine these rules as needed.
- **Registry access control rules (Windows only):** The behavior analysis 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. The behavior analysis creates Registry Set variables based on the registry resources that were 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 that you are imposing. Use Registry Key Summary reports to help refine these rules, if needed.
- **Network access control rules:** The behavior analysis 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 the behavior analysis creates the rules for the rule module, it also creates all the registry and COM component variables that are required by the rules. All Windows files are entered as literals. (UNIX files are grouped into sets.)

Additionally, the behavior analysis creates a new application class for the analyzed application and uses this new application class in all rules that make up the rule module. You should note that if you select more than one application class for the analysis, the application class created for the rule module is an aggregate of all the analyzed applications. If you decide that the application is not dangerous and it can run without any rule module restrictions, you can begin to configure the behavior analysis.

Summary

This topic summarizes what you have learned in this lesson.

Summary

Cisco.com

- **Application deployment investigation analysis provides administrators with application deployment information that is vital to protecting hosts.**
- **Application deployment reports are used to design optimal protection policies.**
- **Application behavior analysis can make CSA policies more effective, with an analysis of the resources that are accessed by an application in normal operation.**
- **A behavior analysis job can evaluate the operation of an application, or a number of applications, in an application class.**
- **Behavior analysis will create a policy to protect the application from the system and to protect the system from the application.**
- **The new policy will be specific to the analyzed application. System-hardening rules would need to be added.**

© 2005 Cisco Systems, Inc. All rights reserved. HIPS v2.0—11-39

Lesson 12

Using Event Logs and Generating Reports

Overview

Events and messages logged by the Cisco Security Agent (CSA) can be viewed from the CSA Management Center (CSA MC). You can configure alerts to be sent based on the severity level of the logged event, the specific event, and the host that generated the alert. You can configure the CSA MC to send e-mail, issue Simple Network Management Protocol (SNMP) traps, beep pagers, log to a text file, and execute custom programs. You can generate reports on events logged by the CSA MC.

This lesson contains the following topics:

- Objectives
- How Logging Works
- The Event Log and Event Monitor
- Event Log Management
- Event Sets
- Configuring Alerts
- Generating Reports
- Summary
- Lab Exercise

Objectives

Upon completion of this lesson, you will be able to perform the following tasks:

- Explain the features of the Event Log and Event Monitor
- Manage the size of the Event Log
- Configure filtering of events for logging, reports, and alerts
- Create event-based alerts
- Generate reports on events selected by sorting criteria

How Logging Works

This topic explains how logging works for CSA.

What Is Logged

Cisco.com

The following information is logged for each rule type:

- **File access control logging:** Process path and filenames and file operation are logged.
- **Network access control logging:** Process path, network address, port, and direction are logged.
- **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.

© 2005 Cisco Systems, Inc. All rights reserved.HIPS v2.0—12-3

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.

Caution In some cases, when an event is logging continuously, the Agent 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 filenames and file operation are logged.
- **Network access control logging:** Process path, network address, port, and direction are logged.

Note No network access control rule denial events are logged for any TCP or (User Datagram Protocol (UDP) port resulting from multicast packet signals.

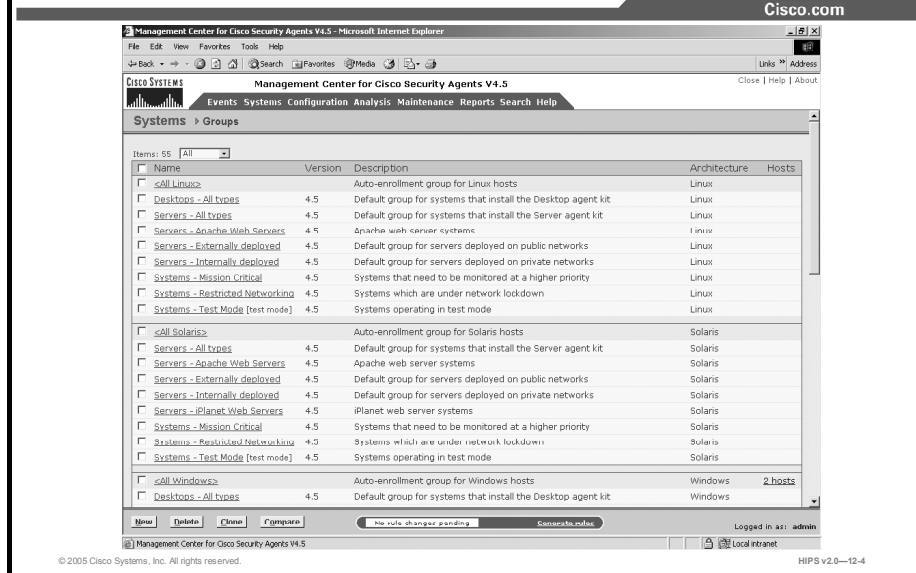
- **Registry access control logging:** Process path and registry key are logged.
- **Component Object Model (COM) component access control logging:** Process path and COM component program identifiers (PROGIDs) or class identifiers (CLSIDs) are logged.

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



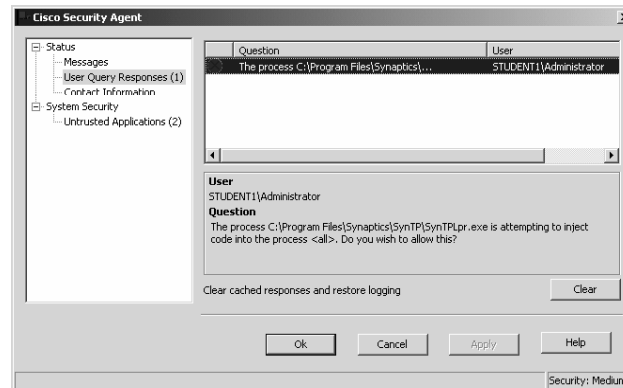
Enable verbose logging mode in the Group configuration view to change the Event Log timer to log *all* recurring events rather than only logging recurring events once every hour. Verbose logging applies to all policies that are attached to the group that have logging turned on.

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—that is, 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.

Note Verbose logging is enabled on a host if any group in which the host is a member has verbose logging turned on.

Logging and Query User Rules

Cisco.com



When a user responds to a Query User box (by selecting Yes, No, or Terminate), 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 popup query box appearing again.

When the user responds to a triggered Query User popup box, the system action that triggered the popup, 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.

The Event Log and Event Monitor

This topic explains the use of the Event Log and Event Monitor.



The Event Log view, available from the **Events** 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 tells you the following:

- **Filter by Eventset:** This displays the name of the event set, if any, used to filter the Event Log view.
- **Optionally, you can define a filter with the following parameters:**
 - **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 Module:** From the drop-down list, select a rule module to search for events generated by that module.
 - **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.
 - **Filter Out Duplicates:** Use this radio button to pare down events and remove all duplicate events from the display.

Start Date and End Date

To search events, click the **Change Filter** link to access a popup 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 the following time format: hh:mm:ss. If no meridian (AM or PM) is specified, hh is interpreted on a 24-hour clock (0-23). 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. The default year is the current year.

Minimum and Maximum Severity Settings

From the Minimum and Maximum Severity drop-down 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 have 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 can enter a specific hostname to search for that host. Click the **Change** link beside the Host field for a host selection box.

Events per Page

Enter the number of events per page that 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.

Filter Out Duplicates

You can select to filter out duplicate events. This will cause the Event Log to display only the most recent event for duplicated event entries.

Event Log (Cont.)

Management Center for Cisco Security Agents V4.5 - Microsoft Internet Explorer

Management Center for Cisco Security Agents V4.5

Events Systems Configuration Analysis Maintenance Reports Search Help

Events > Event Log

Viewing 77 - 28 of 77 events [Change filter](#)

Event log generation time: 2/16/2005 7:54:43 AM
Severity: Information - Emergency
Hosts: All
Rule Module: All
Events per page: 50
Filter out duplicates: No

[Latest](#) [Earliest](#)

#	Date	Host	Severity	Event
77	2/16/2005 7:16:36 AM	student1	Alert	The process 'System' (as user NT AUTHORITY\SYSTEM) attempted to communicate with 10.0.1.50 on TCP port 139. The attempted access was to accept a connection as a server (operation = ACCEPT). The operation was denied. Details Rule Wizard Find Similar
76	2/16/2005 6:15:52 AM	student1	Alert	The process 'System' (as user NT AUTHORITY\SYSTEM) attempted to communicate with 10.0.1.50 on TCP port 139. The attempted access was to accept a connection as a server (operation = ACCEPT). The operation was denied. Details Rule Wizard Find Similar
75	2/16/2005 5:15:07 AM	student1	Alert	The process 'System' (as user NT AUTHORITY\SYSTEM) attempted to communicate with 10.0.1.50 on TCP port 139. The attempted access was to accept a connection as a server (operation = ACCEPT). The operation was denied. Details Rule Wizard Find Similar
74	2/16/2005 4:14:22 AM	student1	Alert	The process 'System' (as user NT AUTHORITY\SYSTEM) attempted to communicate with 10.0.1.50 on TCP port 139. The attempted access was to accept a connection as a server (operation = ACCEPT). The operation was denied. Details Rule Wizard Find Similar
73	2/16/2005 3:13:37 AM	student1	Alert	The process 'System' (as user NT AUTHORITY\SYSTEM) attempted to communicate with 10.0.1.50 on TCP port 139. The attempted access was to accept a connection as a server (operation = ACCEPT). The operation was denied. Details Rule Wizard Find Similar

© 2005 Cisco Systems, Inc. All rights reserved. HIPS v2.0—12-8

The Event Log screen 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. The screen also provides the exact time that 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 that 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 that you wish to search by and select a time frame greater or less than the time that the event in question was logged.

Use the Wizard link, where available, to edit the rule that caused the event.

Event Monitor

The screenshot displays the Cisco Management Center for Cisco Security Agents V4.5 interface. The main content area is titled "Event Monitor" and shows a monitoring filter and a table of events.

Monitoring filter

- Displaying: last 50 events
- Severity: Information - Emergency
- Host: All
- Rule Module: All
- Filter out duplicates: No
- Next refresh: in 8 seconds
- Refresh interval: 15 seconds

Event List

#	Date	Host	Severity	Event
50	2/16/2005 8:17:21 AM	student1	Alert	The process 'System' (as user NT AUTHORITY\SYSTEM) attempted to communicate with 10.0.1.50 on TCP port 139. The attempted access was to accept a connection as a server (operation = ACCEPT). The operation was denied. Details: Rule 160 Wizard
49	2/16/2005 7:16:36 AM	student1	Alert	The process 'System' (as user NT AUTHORITY\SYSTEM) attempted to communicate with 10.0.1.50 on TCP port 139. The attempted access was to accept a connection as a server (operation = ACCEPT). The operation was denied. Details: Rule 160 Wizard
48	2/16/2005 6:15:52 AM	student1	Alert	The process 'System' (as user NT AUTHORITY\SYSTEM) attempted to communicate with 10.0.1.50 on TCP port 139. The attempted access was to accept a connection as a server (operation = ACCEPT). The operation was denied. Details: Rule 160 Wizard
47	2/16/2005 5:15:07 AM	student1	Alert	The process 'System' (as user NT AUTHORITY\SYSTEM) attempted to communicate with 10.0.1.50 on TCP port 139. The attempted access was to accept a connection as a server (operation = ACCEPT). The operation was denied. Details: Rule 160 Wizard
46	2/16/2005 4:14:22 AM	student1	Alert	The process 'System' (as user NT AUTHORITY\SYSTEM) attempted to communicate with 10.0.1.50 on TCP port 139. The attempted access was to accept a connection as a server (operation = ACCEPT). The operation was denied. Details: Rule 160 Wizard

© 2005 Cisco Systems, Inc. All rights reserved. HIPS v2.0-12-9

Similar to the Event Log, the Event Monitor, available from the Events 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 popup window from which you can edit these values and change the event filter. Refer back to the earlier discussion for more information on these fields.

Unlike the Event Log page, the Event Monitor page automatically refreshes itself at set intervals. The event list is updated with the latest events each time the page refreshes.

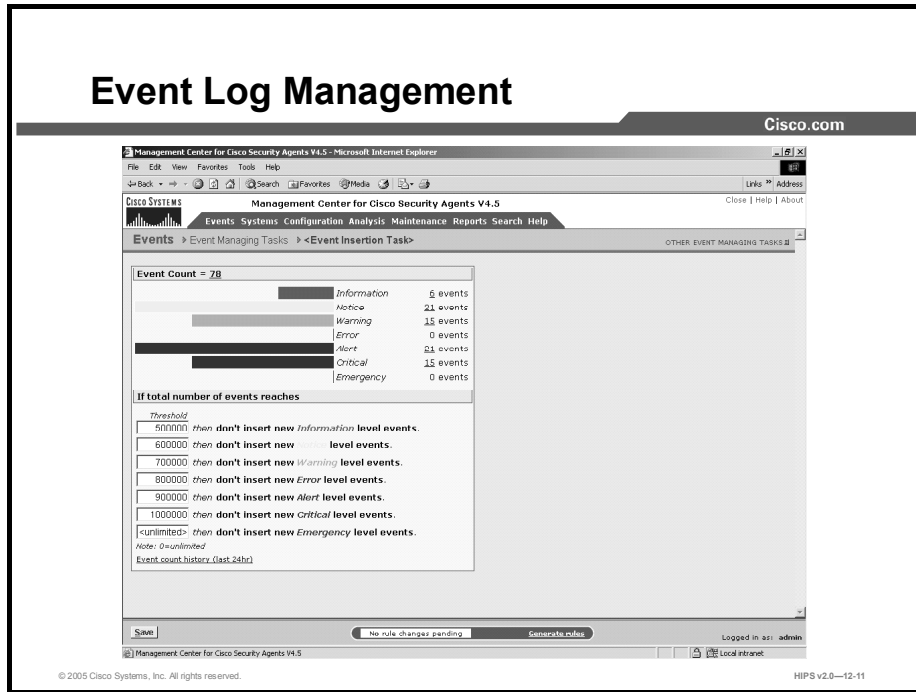
The footer of this page provides a Refresh button and a Pause button. Use the Refresh button to refresh the page immediately without waiting for the set refresh interval to occur. Use the Pause button to immediately stop the page from refreshing. The set refresh interval will then stop wherever it is in the countdown. This pause feature is useful when you are testing policies and you want to mark a certain place as a starting point for receiving new events. When you click it, the Pause button becomes a Resume button.

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

This topic explains the configuration and use of Event Log management.



The Event Log Management feature, available from the **Events** 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 it from growing too large and from maintaining stale information.

You can configure global event insertion threshold parameters from the global Event Insertion Tasks page. This page already contains default settings for stopping the insertion of additional events for each event level when the specified threshold setting is reached. You can change these settings, if necessary. The thresholds on this page only trigger if the Event Log Management parameters that you configure (described in the second set of instructions here) do not adequately keep events pruned below configured levels. If there is a sudden flurry of events and configured pruning parameters do not trigger immediately, for example, the global thresholds will kick in.

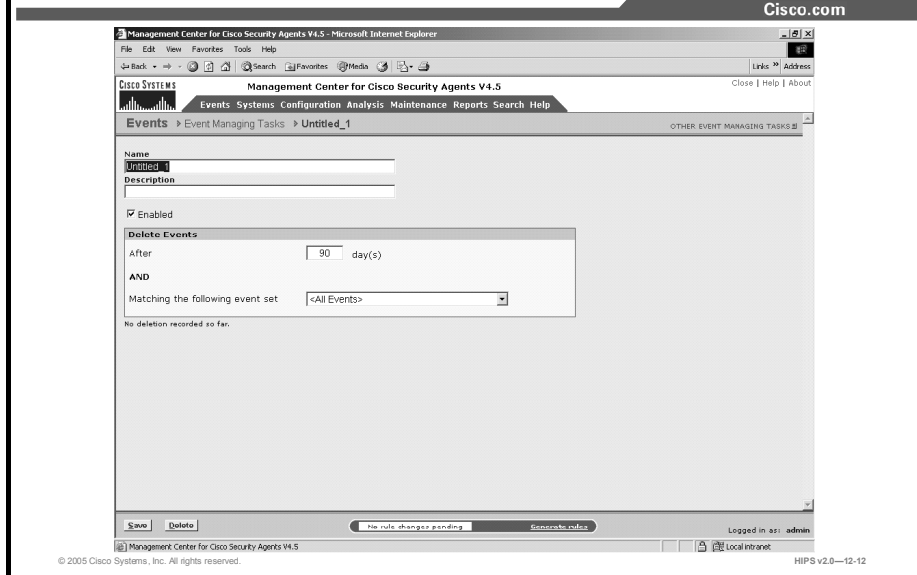
Complete the following steps to access the global Event Insertion Tasks page:

Step 1 Select **Events > Event Log Management**. The Event Log Management window is displayed.

Step 2 Click the top bracketed link **<Event Insertion Tasks>** to access the page.

This page displays the total number of events in the Event Log. It also breaks events out to the number of events that exist for each severity level. Beneath this graphical event display are the default threshold settings for each event level. These thresholds represent the upper limit of events that must be reached for each severity level before no more events of this type will log. Event pruning must occur in order for these event types once again to be written to the Event Log.

Event Log Management (Cont.)

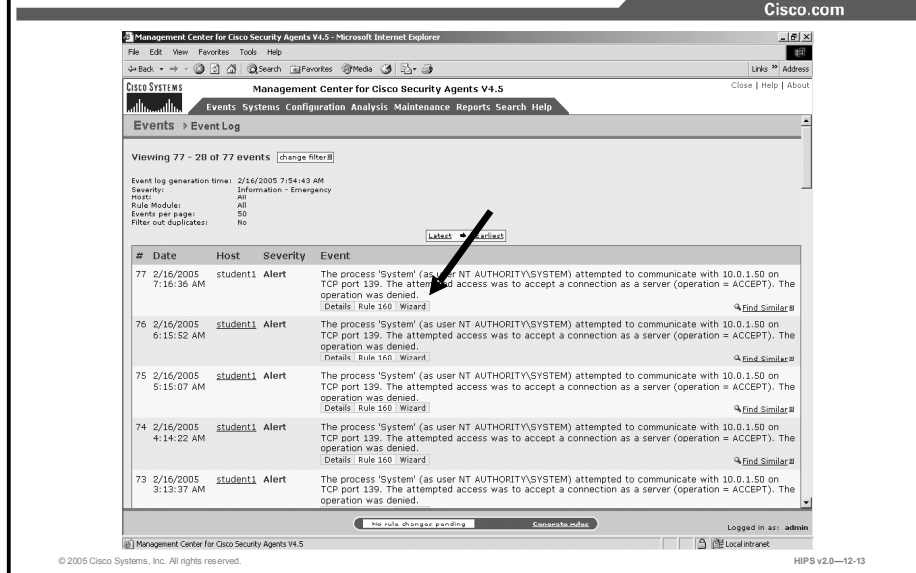


To configure an event autopruning task, complete the following tasks:

- Step 1** Select **Maintenance > Event Log Management**. The Event Log Management Window is displayed.
- Step 2** Click the **New** button to create a new entry. This takes you to the autopruning configuration view.
- Step 3** Enter a name for the autopruning task.
- 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** Use the Enabled check box to enable this event autopruning configuration. (It is enabled by default.) By not selecting this check box, you can save this item, but it will not be active.
- Step 6** Enter a value in the Delete Events - After field. This is the value for which events, once having been in the log for this number of days, are deleted. Before these events are removed, they must also match the parameters of the event set selected on this page.
- Step 7** Select the preconfigured event set for the event type you want to prune from the Event Log in the drop-down box next to Matching the Following Event Set. Configuring event sets provides flexibility in selecting the events for autopruning.
- Step 8** Click the **Save** button.

Note This purging of events will occur periodically based upon the configured autopruning items. Generally, this pruning will take place at a time when the least activity is registered on the MC. When event autopruning occurs, a message appears in the Event Log notifying you of this action.

Event Management Wizard



Use the Event Management Wizard to accomplish the following:

- 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.
- Create an exception rule that stops a specific event from logging. The wizard makes use of the Take Precedence Over Other <Action Type> Rules feature to manipulate rule precedence and prevent logging of an event.
- Perform a Behavior Analysis Investigation for the application that caused the event.

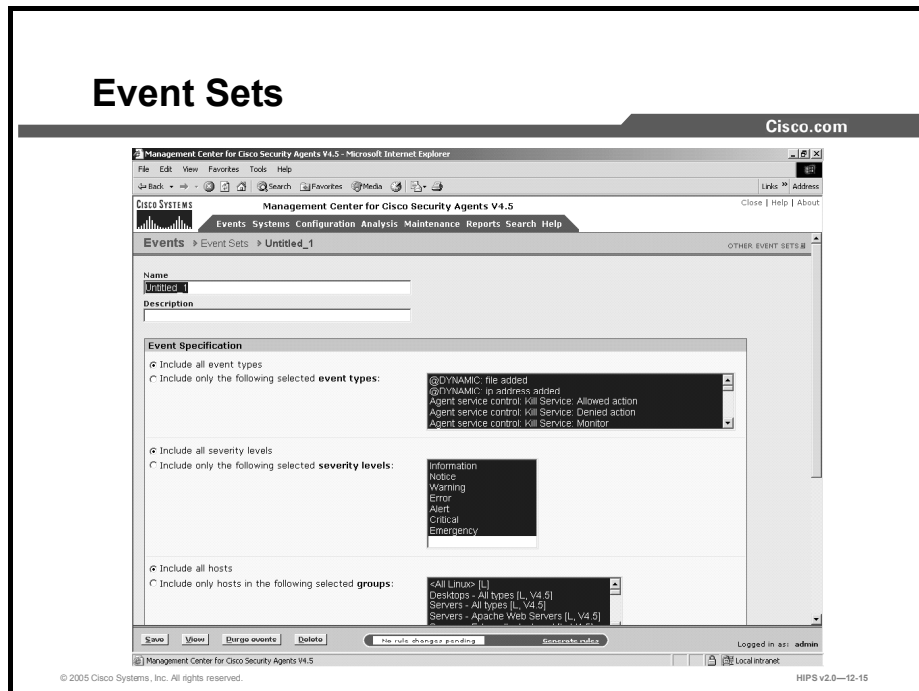
The Event Management Wizard is available for events triggered by deny rules and query user rules of the following types:

- Application control
- Buffer overflow
- COM component access control
- File access control
- Network access control
- Registry access control
- Rootkit/kernel protection
- System API control

You launch the wizard from a Wizard link, which appears with certain Event Log messages.

Event Sets

This topic discusses event sets and how to configure them.



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.

CSA MC ships with several preconfigured event sets that 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.

Complete the following steps to configure event sets:

- Step 1** Select **Events > Event Sets**. 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** Enter a name in the Name field. This is a unique name for this event set. Generally, you will want to adopt a naming convention that lets you quickly recognize event sets in alert configuration fields.
- Step 4** Enter a description in the Description field. 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 Control key as you select each item. To deselect a single item, hold down the Control key when you click on the item in question. Press the Shift key to select multiple successive items.

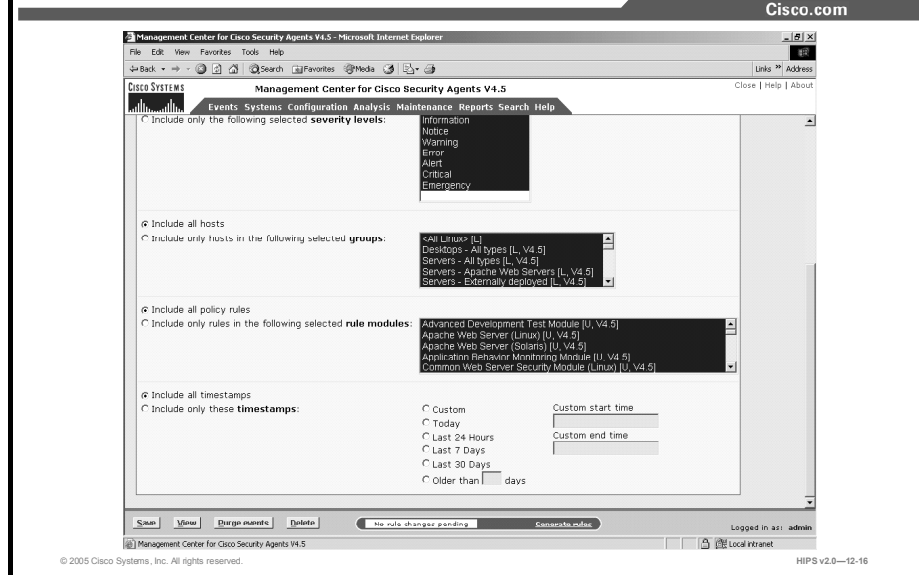
Step 5 Select filter by event specifications.

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, 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 6 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, 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.

Event Sets (Cont.)



Step 7 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, 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 8 Select filter by policy specifications.

Leave the Include All Policy Rules radio button selected to have events generated by all rule modules included, or select the Include Only Rules in the Following Selected Rule Modules radio button. If you select the second radio button, you must select the rule module(s) that trigger an alert for this event set. Any rule modules selected here that log the event in question will trigger an alert.

Step 9 Select filter by time specifications.

Note If you do not have Include All Timestamps selected, the event set is not available for use in alerts.

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, you can create a custom time here or select from available times: Today, Last 24 hours, Last 7 days, Last 30 days, and Events older than <you specify> days to trigger an alert when an event occurs with the specified time range.

Step 10 You can also enter custom start and custom end times in the following manner:

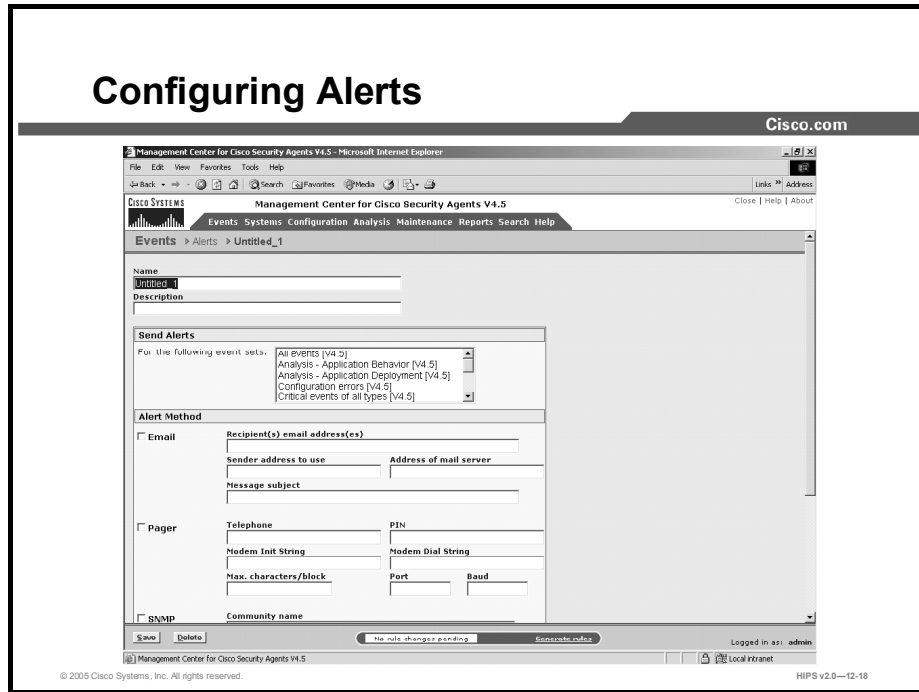
- Specify a relative time using any of the following terms: tomorrow, yesterday, today, now, last, this, next, ago, year, month, week, day, hour, minute, and second.
- Enter a specific time using the following time format: hh:mm:ss. If no meridian (AM or PM) is specified, hh is interpreted on a 24-hour 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. The default year is the current year.

Note When you select multiple categories to filter by, all selections have to match.

- Step 11** When all required information is entered, click the **Save** button to enter and save your event set in the CSA MC database.
- Step 12** 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.
 - 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.

Configuring Alerts

This topic discusses alerts and how to configure them.



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.

Complete the following steps to configure CSA MC to issue alerts when specified system events occur:

- Step 1** Select **Events > Alerts**. 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.

Configuring Alerts (Cont.)

The screenshot shows the 'Alert Method' configuration page in the Cisco Management Center. The page is divided into several sections, each with a checkbox and associated input fields:

- Email:** Includes fields for 'Recipient(s) email address(es)', 'Sender address to use', 'Address of mail server', and 'Message subject'.
- Pager:** Includes fields for 'Telephone', 'PIN', 'Modem Init String', 'Modem Dial String', 'Max. characters/block', 'Port', and 'Baud'.
- SNMP:** Includes fields for 'Community name' and 'Manager IP address'.
- Log:** Includes a field for 'Log file'.
- Custom:** Includes a field for 'Custom program'.
- Named Pipe:** Includes a field for 'Named Pipe' with a dropdown menu.

At the bottom of the page, 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'.

- Step 3** In the alert configuration view, 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 Sets list box, select the event set(s) to trigger the alert that you are 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 Control key as you select each item. To deselect a single item, hold down the Control 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 that 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 each alert type you want to send, select the corresponding check box 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 check box, the alert in question is not enabled; however, the information you enter 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.

Generating Reports

This topic discusses generating event reports.

Viewing Reports

Cisco.com

When you generate your reports, you are given the option of selecting the type of viewer through which to display the report. From the Viewer Type drop-down menu, you can select either of the following:

- **ActiveX**
- **HTML Frame**

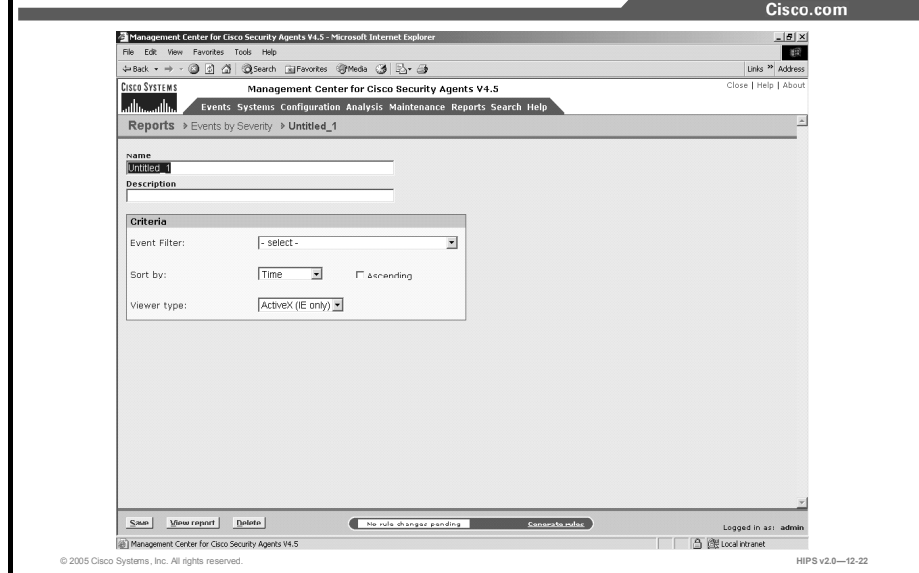
© 2005 Cisco Systems, Inc. All rights reserved. HIPS v2.0—12-21

When you generate your reports, you are given the option of selecting the type of viewer through which to display the report. From the Viewer Type drop-down 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, but not supported by Netscape).
- **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 viewer type you have selected and the printer settings on the printer.

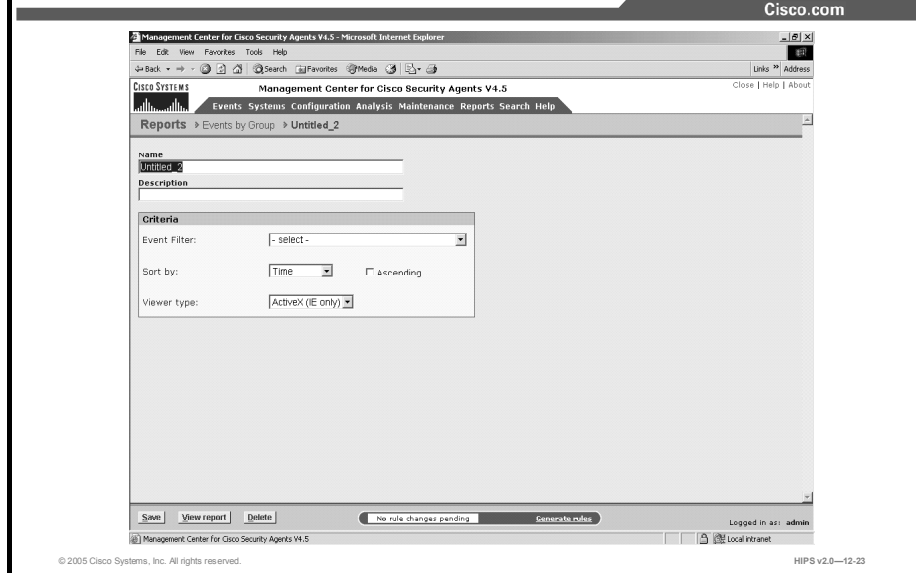
Events by Severity Reports



Complete the following steps to generate an Events by Severity report:

- 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 drop-down list, select an Event Filter. This is an event set that you create from the Monitor > Event Sets configuration view.
- Step 5** From the Sort By drop-down list, select a parameter for sorting the contents of this report.
- Step 6** Enable or disable the Ascending check box 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.
- Step 8** Click the **Save** button to save the parameters that you have just configured for generating this report.
- Step 9** Click the **View Report** button and the report is automatically displayed in a new browser window.

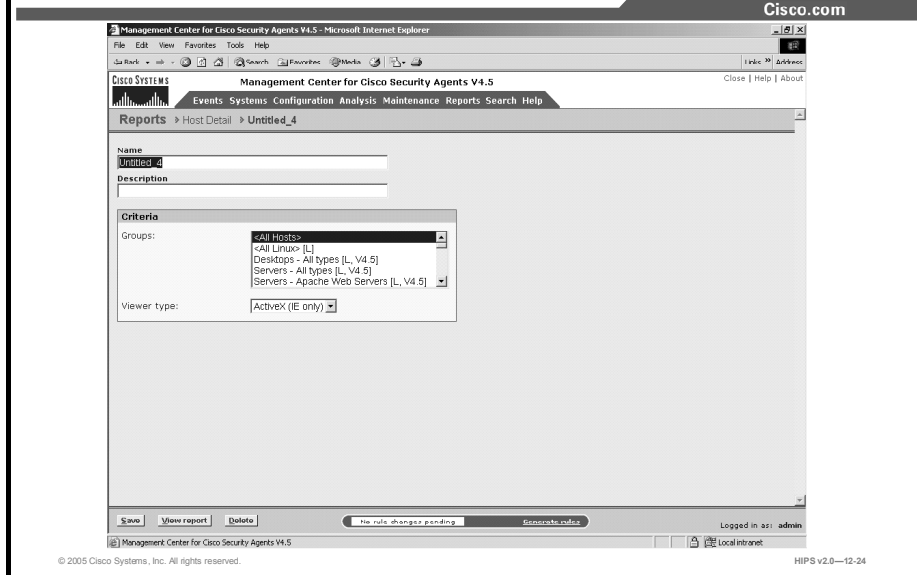
Events by Group Reports



Complete the following steps to generate an Events by Group report:

- 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 drop-down list, select an Event Filter. This is an event set that you create from the Monitor > Event Sets configuration view.
- Step 5** From the Sort By drop-down list, select a parameter for sorting the contents of this report.
- Step 6** Enable or disable the Ascending check box 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.
- Step 8** Click the **Save** button to save the parameters that you have 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 Reports



Complete the following steps to generate a Host Detail report.

- 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, 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 Control key as you select each item. To deselect a single item, hold down the Control 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.
- Step 5** By default, ActiveX is selected as the viewer type. This is the recommended viewer.
- Step 6** Click the **Save** button to save the parameters that you have just configured for generating this report.
- Step 7** Click the **View Report** button and the report is automatically displayed in a new browser window.

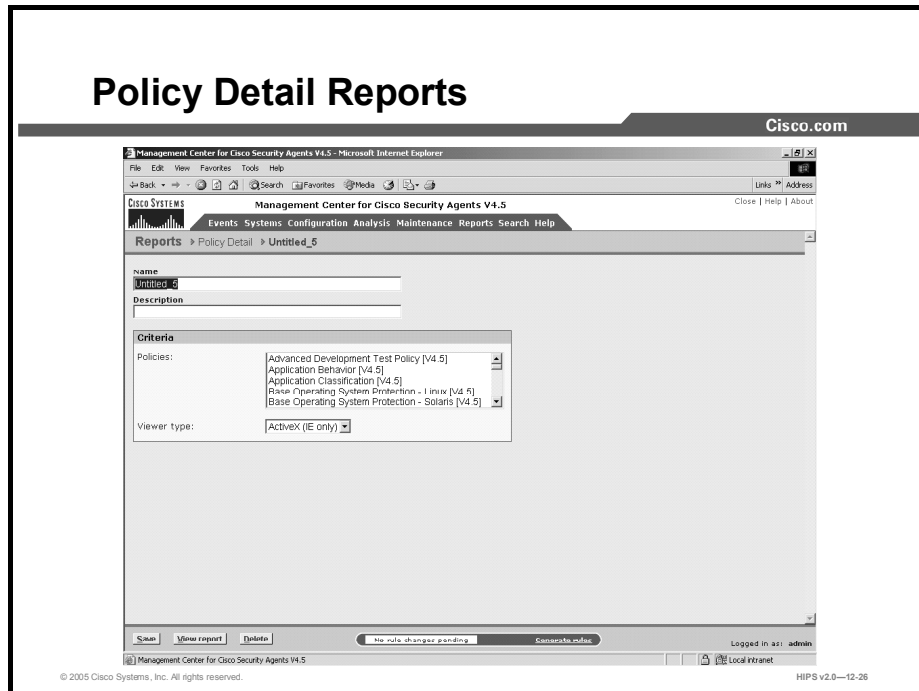
Group Detail Reports



Complete the following steps to generate a Group Detail report:

- 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 Control key as you select each item. To deselect a single item, hold down the Control 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.
- Step 6** Click the **Save** button to save the parameters that you have 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.

Policy Detail Reports



Complete the following steps to generate a Policy Detail report:

- 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 Control key as you select each item. To deselect a single item, hold down the Control 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.
- Step 6** Click the **Save** button to save the parameters that you have 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.

Summary

This topic summarizes the information that you learned in this lesson.

Summary

Cisco.com

- The Event Log provides access to in-depth information on events and tools to filter events by many criteria.
- The Event Monitor refreshes frequently and can provide an open monitoring screen that secures the CSA MC from nonadministrative changes.
- You can configure Event Log Management to selectively prune the size of the Event Log.
- Event sets can be used for logging, reports, and alerts to increase granularity of the output.
- Event-based alerts can be configured to be sent by e-mail, pager, SNMP, log, or a custom application.
- Reports on events can be generated based on criteria that you have selected.

© 2005 Cisco Systems, Inc. All rights reserved. HIPS v2.0—12-27

HIPS

Securing Hosts Using Cisco Security Agent

Version 2.0

Lab Guide

CLS Production Services: 06.07.05

Copyright © 2005, Cisco Systems, Inc. All rights reserved.

Cisco Systems has more than 200 offices in the following countries and regions. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at www.cisco.com/go/offices.

Argentina • Australia • Austria • Belgium • Brazil • Bulgaria • Canada • Chile • China PRC • Colombia • Costa Rica
Croatia • Cyprus • Czech Republic • Denmark • Dubai, UAE • Finland • France • Germany • Greece
Hong Kong SAR • Hungary • India • Indonesia • Ireland • Israel • Italy • Japan • Korea • Luxembourg • Malaysia
Mexico • The Netherlands • New Zealand • Norway • Peru • Philippines • Poland • Portugal • Puerto Rico • Romania
Russia • Saudi Arabia • Scotland • Singapore • Slovakia • Slovenia • South Africa • Spain • Sweden • Switzerland
Taiwan • Thailand • Turkey • Ukraine • United Kingdom • United States • Venezuela • Vietnam • Zimbabwe



Copyright © 2005 Cisco Systems, Inc. All rights reserved. CCSP, the Cisco Square Bridge logo, Follow Me Browsing, 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 Access Registrar, Aironet, ASIST, BPX, Catalyst, CCDA, CCDP, CCIE, CCIP, CCNA, CCNP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Cisco Unity, Empowering the Internet Generation, Enterprise/Solver, EtherChannel, EtherFast, EtherSwitch, Fast Step, FormShare, GigaDrive, GigaStack, HomeLink, Internet Quotient, IOS, IP/TV, iQ Expertise, the iQ logo, iQ Net Readiness Scorecard, LightStream, Linksys, MeetingPlace, MGX, the Networkers logo, Networking Academy, Network Registrar, Packet, PIX, Post-Routing, Pre-Routing, ProConnect, RateMUX, ScriptShare, SlideCast, SMARTnet, StrataView Plus, 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 United States and certain other countries.

All other trademarks mentioned in this document or Website 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. (0501R)

DISCLAIMER WARRANTY: THIS CONTENT IS BEING PROVIDED "AS IS." CISCO MAKES AND YOU RECEIVE NO WARRANTIES IN CONNECTION WITH THE CONTENT PROVIDED HEREUNDER, EXPRESS, IMPLIED, STATUTORY OR IN ANY OTHER PROVISION OF THIS CONTENT OR COMMUNICATION BETWEEN CISCO AND YOU. CISCO SPECIFICALLY DISCLAIMS ALL IMPLIED WARRANTIES, INCLUDING WARRANTIES OF MERCHANTABILITY, NON-INFRINGEMENT AND FITNESS FOR A PARTICULAR PURPOSE, OR ARISING FROM A COURSE OF DEALING, USAGE OR TRADE PRACTICE. This learning product may contain early release content, and while Cisco believes it to be accurate, it falls subject to the disclaimer above.

Lab Guide

Overview

This guide presents the instructions and other information concerning the activities for this course.

Outline

This guide includes these activities:

- Lab 3-1: Cisco Security Agent Quick Start Installation
- Lab 5-1: Configuring Groups and Managing Hosts
- Lab 6-1: Building Policies
- Lab 7-1: Rule Basics
- Lab 9-1: Defining Application Classes
- Lab 10-1: Working with Variables
- Lab 11-1: Using Cisco Security Agent Analysis
- Lab 12-1: Using Event Logs and Generating Reports

Lab 3-1: Cisco Security Agent Quick Start Installation

Complete the following lab exercise to practice what you learned in this lesson.

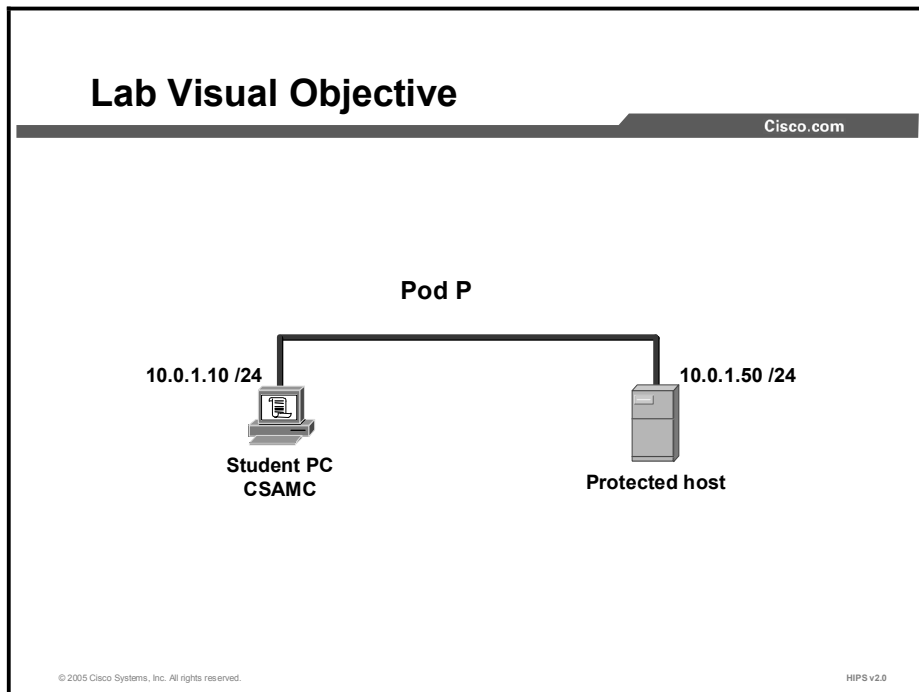
Objectives

In this lab exercise you will complete the following tasks:

- Install the CSA MC module
- Launch CiscoWorks and access the CSA MC interface
- Verify the Web Servers group for Windows and obtain the Agent kit URL for that group
- Install the CSA software on the host

Visual Objectives

The following figure displays the configuration that you will complete in this lab exercise.



Setup

Before starting this lab exercise, verify that CiscoWorks is installed on your system. Verify that the host can ping your system.

Task 1: Install the CSA MC

Complete the following steps to install the CSA MC on your system:

- Step 1** Log in to your system as local administrator with username **administrator** and password **attack**.

- Step 2** Choose **My Computer > C: > Apps > CSAMC_4.5** to open the CSA MC folder.
- Step 3** Double-click the **setup.exe** file to begin the installation. When the Management Center for Cisco Security Agents V4.5 installation screen appears, click **Next** to begin the installation.
- Step 4** Select the **Local Database** radio button to maintain all configuration data on the local machine. Click the **Next** button.
- Step 5** A window appears stating that Cisco Security Agents V4.5 requires version 8.0 or higher of Microsoft SQL Server 2000 desktop engine (MSDE) and notifies you that setup will install the necessary components. Click the **Yes** button. The Specified the Installation Path for MSDE window appears.
- Step 6** Click **Next** to accept the default destination folder. A window appears stating that you must restart the system for the configuration changes made to Microsoft SQL Server desktop engine to take effect.
- Step 7** Click the **Yes** button to restart the system.
- Step 8** Log in to your system as local administrator with username **administrator** and password **attack**.
- Step 9** Choose **My Computer > C: > Apps > CSAMC_4.5** to open the CSA MC folder.
- Step 10** Double-click the **setup.exe** file to begin the installation. When the Management Center for Cisco Security Agents V4.5 installation screen appears, click **Next** to begin the installation. The Welcome to the Installation Wizard for Management Center for Cisco Security Agents V4.5 window appears.
- Step 11** Select the **Local Database** radio button to maintain all configuration data on the local machine.
- Step 12** Click the **Next** button.
- Step 13** A window appears stating that the Management Center for Cisco Security Agents V4.5 requires a license to be fully functional. Click the **Yes** button to install the license. The select license file window appears.
- Step 14** Select **local disk (C:)** from the Look in: drop-down menu.
- Step 15** Double-click the **Apps** folder.
- Step 16** Double-click the **CSAMC_4.5** folder.
- Step 17** Double-click the **license** folder.
- Step 18** Select the **CSAMC.lic** file and click the **Open** button. The Ready to Install the Program window appears.
- Step 19** Click the **Install** button to begin the installation.
- Step 20** Once the installation is complete, a window appears stating that the system will restart in 5 minutes. Click the **OK** button to restart immediately.

Task 2: Launch CiscoWorks and Access the CSA MC Interface

Complete the following steps to establish secure communication with CiscoWorks and access the CSA MC interface:

- Step 1** Log in to your system as local administrator with the username **administrator** and password **attack**.
- Step 2** Launch CiscoWorks by choosing **Start > Programs > CiscoWorks > CiscoWorks**.
- Step 3** The CiscoWorks login manager appears.
- Step 4** Log in to CiscoWorks with username **admin** and password **cisco**.
- Step 5** Choose **VPN/Security Management Solution > Management Center > Security Agents V4.5**.

Task 3: Verify the Default Servers Group and Obtain the Agent Kit URL

Complete the following steps to use the CSA MC interface to examine the Web Servers group for Windows and access the CSA MC to obtain the Agent kit URL for the Web Servers group:

- Step 1** At the CSA MC interface, click **Systems** on the main toolbar, and then click **Groups** in the drop-down menu.
- Step 2** Scroll down to the **Servers IIS Web Servers** group link in the lower box of groups. The correct group will show “Windows” in the Operating System column to the right.
- Step 3** Click the **Servers IIS Web Servers (for Windows)** link. This page also displays the policies attached to this preconfigured group.
- Step 4** Click **Systems** on the main menu bar, and then click **Agent Kits** in the drop-down menu.
- Step 5** Click the **New** button. The Select Target Architecture window appears.
- Step 6** Click the **Windows** button. The New Agent Kit window appears.
- Step 7** Enter **Web_Servers** in the name text box.
- Step 8** Enter **IIS Web Servers** in the description text box.
- Step 9** Select **Servers-IIS Web Servers [V4.5 r369]** from the Select the groups with which this kit should be associated.
- Step 10** Select the Force reboot after install check box.
- Step 11** Click the **Make kit** button. A window appears stating that the Agent kit was successfully created.
- Step 12** Click the **Generate rules** link. The Generate Rules Program window appears.
- Step 13** Click the **Generate** button. After the rules are generated, a window appears stating that the rule program generation was successful.
- Step 14** Click **Systems** on the main menu bar, and then click **Agent Kits** in the drop-down menu. The Systems > Agent Kits window appears.
- Step 15** Select the **Web_Servers** link. The Web_Servers Agent kit window appears.
- Step 16** Copy the URL as shown in the download URL portion of the Agent Kits window.

Task 4: Install CSA Software on the Host

Complete the following steps to use the Agent kit URL to download the CSA software from the CSA MC to the host and install the CSA software:

- Step 1** Log in to the host as the local administrator.
- Step 2** Launch the Internet browser on the host. Type or paste the Agent kit URL into the Address field and press **Enter**.
- Step 3** Click the **Yes** button when the security alert window appears. The file download window appears.
- Step 4** Click the **Open** button.
- Step 5** The CSA software begins downloading to the host.
- Step 6** When the software download is complete, CSA automatically installs and then announces that the system will reboot in 3 minutes. Click **OK**.
- Step 7** Log in to the host as the local administrator.
- Step 8** To verify that the Protected host has registered with the CSA MC, right-click the **red CSA flag** in the system tray and click **Open Agent Panel**. The CSA control panel is displayed.
- Step 9** Click **Status** from the main menu. If the Protected host has successfully registered with the CSA MC, a registration date and time are indicated.
- Step 10** Also verify registration by the Protected host on the CSA MC by choosing **Systems > Hosts > <name of the Protected host>**. If the Protected host has successfully registered with the CSA MC, a registration date and time are indicated.

Lab 5-1: Configuring Groups and Managing Hosts

Complete the following lab exercise to practice what you learned in this lesson.

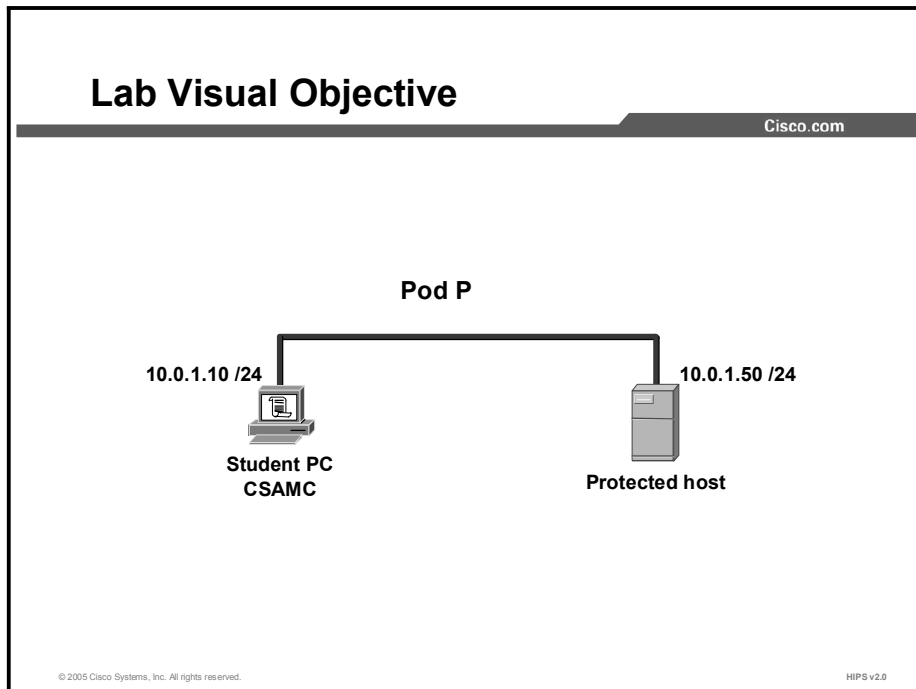
Objectives

In this lab exercise you will complete the following tasks:

- Create a new group
- Build a new Agent kit
- Change the group membership of a host

Visual Objectives

The figure displays the configuration that you will complete in this lab exercise.



Task 1: Create a New Group

Complete the following steps to create a new group for Windows hosts and configure it for test mode and verbose logging:

- Step 1** Choose **Systems > Groups**. A list of existing groups will be displayed in the left column.
- Step 2** Click the **New** button to create a new group entry.
- Step 3** Click the **Windows** button.

- Step 4** Enter **Test Servers** in the Name field and **Group for testing new server policies** in the Description field in the appropriate areas for the new group.
- Step 5** Enter **2 minutes** in the Polling interval field.
- Step 6** Click the **Save** button.
- Step 7** Verify creation of the new Test Servers group by choosing **Systems > Groups**.

Task 2: Build a New Agent Kit

Complete the following steps to build an Agent kit:

- Step 1** Choose **Systems > Agent Kits**. Agent kits that are preconfigured or that have been added are displayed.
- Step 2** Click the **New** button to create a new Agent kit. A popup window asks which operating system is the target of the Agent kit.
- Step 3** Click the **Windows** button.
- Step 4** Enter **Test_kit_1** in the Name field and **Kit for testing new policies** in the Description field.
- Step 5** Choose **Test Servers** from the Select the groups with which this kit should be associated field.
- Step 6** Click the **Make kit** button.
- Step 7** Click the **Generate rules** link.
- Step 8** Click **Generate** at the Generate Rules Program window.
- Step 9** Verify the new Agent kit by choosing **Systems > Agent Kits**.

Task 3: Change the Group Membership of a Host

Complete the following steps to change the group membership of a host and join a host to an additional group:

- Step 1** Choose **Systems > Hosts**. Hosts that have registered with the CSA MC will be displayed.
- Step 2** Click the hostname of the Protected host that registered with the CSA MC in the previous lab exercise.
- Step 3** Click the **Modify group membership** link located in the upper-left corner of the host detail window.
- Step 4** Select **Test Servers** in the Does not belong to the following groups pane. Click the **Add** button. Do not remove the current group Servers-IIS Web Servers.
- Step 5** Click the **Generate rules** link.
- Step 6** Click the **Generate** button.
- Step 7** Verify the changed group membership of the host by choosing **Systems > Hosts**, selecting the hostname, and scrolling down to **Group Membership and Policy Inheritance**.

Lab 6-1: Building Policies

Complete the following lab exercise to practice what you learned in this lesson.


Objectives

In this lab exercise you will complete the following tasks:

- Create a policy
- Add a rule module to a policy
- Copy a rule between modules
- Compare rule modules
- View change history
- Filter the rules display
- Attach a policy to a group

Visual Objective

The following figure displays the configuration that you will complete in this lab exercise.



Task 1: Create a Policy

Complete the following steps to create a new policy:

- Step 1** Choose **Configuration > Policies**. The list of existing policies is displayed.
- Step 2** Click the **New** button to create a new policy.
- Step 3** Enter **Windows_Test_Policy** in the Name field of the policy configuration window.

- Step 4** Enter **Policy to be tested on all Windows web server systems** in the Description field.
- Step 5** Select the **Windows** check box from the Target Architectures box.
- Step 6** Click the **Save** button.

Task 2: Adding a Rule Module to a Policy

This task involves adding a rule module to the newly created policy.

- Step 1** Choose **Configuration > Policies**. The existing policies are displayed.
- Step 2** Select the **Windows_Test_Policy** link. The Windows Test Policy window appears.
- Step 3** Select the **Modified Rule Module Associations** link. The Rule Module Associations window appears.
- Step 4** Select **Common Web Server Security Module [All Windows, V4.5 R369]** from the Unattached Windows rule modules pane.
- Step 5** Click the **Add** button. The Common Web Server Security module now appears in the Attached Windows rule modules pane.
- Step 6** Click the **Generate rules** link. The Generate Rules Program window appears.
- Step 7** Click the **Generate** button to generate rules.

Task 3: Copying Rules Between Rule Modules

This task involves copying a file access control rule to the Common Web Server Security module.

- Step 1** Choose **Configuration > Rule Modules > Rule Module Windows**. The existing rule modules appear.
- Step 2** Select the **Apache Web Server** module link.
- Step 3** Select the **Modify Rules** link. The existing rules for the Apache Web Server rules module appear.
- Step 4** Select the **File Access Control** rule with ID 235 check box.
- Step 5** Select **Common Web Server Security Module [v4.5 r369]** from the Copy to rule module drop-down menu.
- Step 6** Click the **Copy** button. A window appears asking if you are sure you want to copy this rule.
- Step 7** Click the **OK** button.
- Step 8** Click the **Generate rules** link. The Generate Rules Program window appears.
- Step 9** Click the **Generate** button to generate rules.
- Step 10** Click **Configuration > Rule Modules > Rule Module Windows**.
- Step 11** Select the **Common Web Server Security Module**.
- Step 12** Click **Modify Rules** link.

Step 13 Verify that you now have a new file access control rule.

Task 4: Compare Two Rule Modules

This task involves comparing the Apache Web Server rule module to the Common Web Server Security rule module.

- Step 1** Choose **Configuration > Rule Modules > Rule Module Windows**. Any existing rule modules will appear.
- Step 2** Select the **Apache Web Server** check box.
- Step 3** Select the **Common Web Server Security Module** check box.
- Step 4** Click the **Compare** button. The compare output window appears.
- Step 5** View the output to compare the rules.

Task 5: View Change History

This task involves viewing the change history of a rule module.

- Step 1** Choose **Configuration > Rule Modules > Rule Module Windows**. Any existing rule modules will appear.
- Step 2** Select the **Common Web Server Security Module** link.
- Step 3** Click the **View Change History** link in the upper-left corner.
- Step 4** The **Audit Trail** window appears.
- Step 5** View the changes that have been made to this rule module.

Task 6: Filter the Rules Display

This task involves filtering the rules that are shown when viewing a rule module.

- Step 1** Choose **Configuration > Rule Modules > Rule Module Windows**. Any existing rule modules will appear.
- Step 2** Select the **Common Web Server Security Module** link.
- Step 3** Select the **Modify Rules** link in the upper-left corner. All existing rules are shown.
- Step 4** Select the **All** link in the upper-left corner.
- Step 5** Choose **Data Access Control** from the popup menu.
- Step 6** Verify that only data access control rules are shown.

Task 7: Attach a Policy to a Group

This task involves attaching the newly created policy to a group.

- Step 1** Choose **Configuration > Policies**.
- Step 2** Select the **Windows_Test_Policy** link. The Windows test policy window appears.

- Step 3** Select the **Modify Group Associations** link. The group association window appears.
- Step 4** Select **Test_Servers** from the Windows_Test_Policy is not attached to the following groups pane.
- Step 5** Click the **Add** button. The Test_Servers policy now appears in the Windows_Test_Policy is attached to the following groups pane.
- Step 6** Click the **Generate rules** link. The Generate Rules Program window appears.
- Step 7** Click the **Generate** button to generate rules.
- Step 8** Choose **Systems > Groups**. All existing groups are shown.
- Step 9** Select the **Test_Servers** group. The Test_Servers configuration window appears.
- Step 10** Verify that the **Windows_Test_Policy** appears in the Attached Policies box.

Lab 7-1: Rule Basics

Complete the following lab exercise to practice what you learned in this lesson.

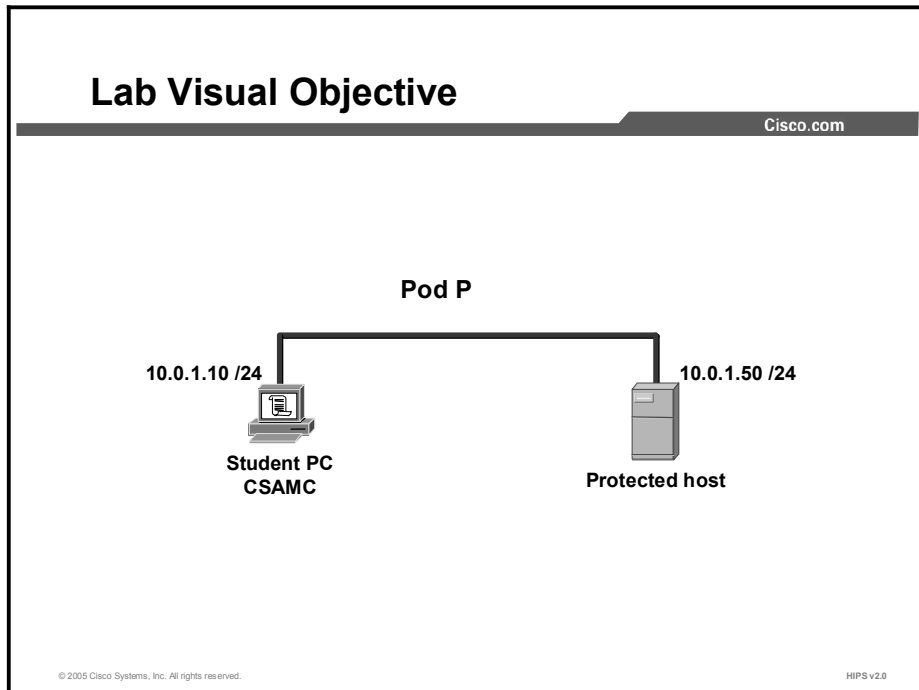
Objectives

In this lab exercise you will complete the following tasks:

- Create a new rule module
- Add an Agent service control rule to a rule module
- Add an application control rule to a rule module
- Add a connection rate limit rule to a rule module
- Add a data access control rule to a rule module
- Add a file access control rule to a rule module
- Add a network access control rule to a rule module

Visual Objective

The following figure displays the configuration that you will complete in this lab exercise.



Task 1: Create a Rule Module

Complete the following steps to create a new rule module:

- Step 1** Choose **Configuration > Rule Modules > Rule Module Windows**. The list of existing rule modules appears.

- Step 2** Click the **New** button to create a new rule module. The new rule module window appears.
- Step 3** Enter **Test_Rule_Module** in the Name field.
- Step 4** Enter **Windows Web Server Rule Module** in the Description field.
- Step 5** Verify that **All Windows** is selected in the target system drop-down menu.
- Step 6** Click the **Save** button.

Task 2: Add an Agent Service Control Rule to a Rule Module

This task involves adding an Agent service control rule to a rule module that will prevent users, including administrators, from suspending CSA security or stopping the CSA service. Complete the following steps to add an Agent service control rule to a rule module:

- Step 1** Choose **Configuration > Rule Modules > Rule Module Windows**.
- Step 2** Select the **Test_Rule_Module** link. The test rule module window appears.
- Step 3** Select the **Modify Rules** link in the upper-left corner of the window.
- Step 4** Select the **Add Rule** link and choose **Agent Service Control**.
- Step 5** Enter **Prevent users or administrators from stopping the agent service** in the Description field.
- Step 6** Check the **Enabled** check box to enable this rule within the policy.
- Step 7** Choose **High Priority Deny** from the Take the following action drop-down menu.
- Step 8** Verify that the **Log** check box is selected to enable logging.
- Step 9** Do not check the **Take precedence over other High Priority Deny rules** check box.
- Step 10** Verify that **<All applications>** is selected from the application in any of the following classes pane.
- Step 11** Verify that the **attempt to disable the agent security** check box is selected. Do not check the **attempt to modify the local agent configuration** check box.
- Step 12** Click the **Save** button.
- Step 13** Choose **Configuration > Policies**. Any existing policies will appear.
- Step 14** Select the **Windows_Test_Policy** link. The Windows_Test_Policy configuration window appears.
- Step 15** Click the **Modify rule module association's** link.
- Step 16** Select **Test Rule Module [All Windows]** from the Unattached Windows Rule Modules pane.
- Step 17** Click the **Add** button. The Test Rule Module [All Windows] now appears in the attached Windows Rule Modules pane.
- Step 18** Click **Common Web Server Security Module [All Windows]** from the attached Windows Rule Modules pane.

- Step 19** Click the **Remove** button.
- Step 20** Click the **Generate rules** link at the bottom of the window. The Generate Rules Program window is displayed.
- Step 21** Click the **Generate** button to generate rules.
- Step 22** Access the protected host. Log on as **administrator** with password **attack**.
- Step 23** On the protected host desktop, right-click the red CSA flag in the system tray and select **Open Agent Panel**. The CSA applet is displayed.
- Step 24** Click the **Poll** button. Verify that the last poll time matches the current time.
- Step 25** Right-click the red system flag in the system tray. Select **Security Level > Off**. Note that the request is prevented, which shows that the policy has been enforced.

Task 3: Add an Application Control Rule to a Rule Module

This task involves adding an application control rule to a rule module that will prevent the Windows Task Scheduler from running editor programs. Complete the following steps to add an application control rule to a rule module:

- Step 1** Choose **Configuration > Rule Modules > Rule Module Windows**. The existing rule modules appear.
- Step 2** Select the **Test_Rule_Module** link. The test rule modules configuration window appears.
- Step 3** Click the **Modify rules** link in the upper-left corner.
- Step 4** Click the **Add rule** link at the bottom of the rule list. The Add rule drop-down menu is displayed.
- Step 5** Select **Application control**. The application control rule configuration window is displayed.
- Step 6** Enter **Stop applications from running editors** in the Description field.
- Step 7** Verify that the **Enabled** check box is checked to enable this rule within the policy.
- Step 8** Select **Query User** from the Take the following action drop-down menu.
- Step 9** Click the **New Link** located next to Query Settings.
- Step 10** Enter **Query User Default Deny** in the Name field.
- Step 11** Enter **Test Application Control Rule Query** in the Description field.
- Step 12** Enter **An application is attempting to run an editor. Do you want to allow this?** in the Text used to query user field.
- Step 13** Select **Deny, Terminate, and Allow** from the Allowed query actions pane.
- Step 14** Select **Deny** from the Default action drop-down menu.
- Step 15** Click the **Save** button. Close the Query Settings window.
- Step 16** Verify that the **Log** check box is checked and the Take Precedence over other Query User (Default Deny) rules check box is not checked.

- Step 17** Select **<All Applications>** from the when current application in any of the following selected classes pane.
- Step 18** Select **Editor Applications** from the New applications in any of the selected classes pane.
- Step 19** Click the **Save** button. A “saved changes” message is displayed.
- Step 20** Click the **Generate rules** link at the bottom of the window. The Generate Rules Program window is displayed.
- Step 21** In the Generate Rules Program window, click the **Generate** button at the bottom of the window. A message will announce that rule program generation was successful.
- Step 22** Access the protected host and log in as **administrator** with password **attack**.
- Step 23** On the protected host desktop, right-click the red CSA flag in the system tray and select **Open Agent Panel**. The CSA applet is displayed.
- Step 24** Click the **Poll** button.
- Step 25** Click the **Status** option and verify that the last poll time is within the last minute or so.
- Step 26** Choose **Start > Run**. The run window is displayed.
- Step 27** Type **Notepad** in the open text box.
- Step 28** Click the **OK** button.
- Step 29** When the user query window appears note the query and select the **Yes** radio button to allow the application to launch.
- Step 30** Close the Notepad window.
- Step 31** In the CSA MC window on the Student PC, choose **Events > Event Log**. Note the new Event Logs for the application control rule.

Task 4: Add a Connection Rate Limit Rule to a Policy

This task involves adding a connection rate limit rule to the rule module that will protect the web server on the protected host from being overloaded by accepting too many connections in a specified period. Complete the following steps to add a connection rate limit rule to a rule module:

- Step 1** Choose **Configuration > Rule Modules > Rule Module Windows**. Existing rule modules are displayed.
- Step 2** Select the **Test_Rule_Module** link. The Test Rule Module configuration window appears.
- Step 3** Click the **Modify rules** link in the upper-left corner.
- Step 4** Click the **Add rule** link at the bottom of the rule list.
- Step 5** Select **Connection rate limit** to open the configuration window for this rule type.
- Step 6** Enter **Protect web server from too many connections** in the Description field.
- Step 7** Check the **Enabled** check box to enable this rule within the policy.

- Step 8** Select **Deny** from the Take the following action drop-down menu.
- Step 9** Check the **Log** check box to enable CSA MC logging for this rule.
- Step 10** Select **<All applications>** from the Applications in any of the selected classes pane.
- Step 11** Choose **server** in the Attempt to act as a drop-down menu.
- Step 12** Choose **all** in the Communicating with <value> hosts drop-down menu.
- Step 13** Enter **2** in the Over Limit text box and **5** in the in <number> minutes field.
- Step 14** Click the **Save** button. A “saved changes” message is displayed.
- Step 15** Click the **Generate rules** link at the bottom of the window. The Generate Rules Program window is displayed.
- Step 16** In the Generate Rules Program window, click the **Generate** button at the bottom of the window. A message announces that rule program generation was successful.
- Step 17** Access the protected host and log on as **administrator** with the password **attack**.
- Step 18** On the protected host desktop, right-click the red CSA flag in the system tray and select **Open Agent Panel**. The CSA applet is displayed.
- Step 19** Click the **Status** option and verify that the last poll time is within the last minute or so.
- Step 20** Open Internet Explorer on the Student PC, enter **10.0.1.50** in the URL field, and press **Enter**. An Under Construction web page will open.
- Step 21** Without closing the first Internet Explorer window, open two more Internet Explorer windows, enter **10.0.1.50** in the address field of each window, and press **Enter**. Note the different result in the third Internet Explorer window.
- Step 22** In the CSA MC window, choose **Events > Event Log**. Note the event logged regarding the connection rate limit rule.

Task 5: Add a Data Access Control Rule to a Rule Module

This task involves adding a data access control rule to a rule module that will protect the web server on the protected host from the dreaded ***]*** attack (fictitious example), which would be delivered in the URI portion of an HTTP request. Complete the following steps to add a data access control rule to a rule module:

- Step 1** Choose **Configuration > Rule Modules > Rule Module Windows**.
- Step 2** Select the **Test_Rule_Module** link. The Test Rule Module configuration window appears.
- Step 3** Click the **Modify rules** link. The Test Rules Module window is displayed.
- Step 4** Click the **Add rule** link at the bottom of the rule list. The Add rule drop-down menu is displayed.
- Step 5** Select **Data access control**. The data access control rule configuration window is displayed.
- Step 6** Enter **Block *]* attack** in the Description field.

- Step 7** Check the **Enabled** check box to enable this rule within the policy.
- Step 8** Choose **High Priority Deny** from the Take the following action drop-down menu.
- Step 9** Check the **Log** check box, but do not check the Take precedence over other High Priority Deny rules check box.
- Step 10** Select **Network Applications** and **Processes created by Network Applications** from the Applications in any of the selected classes pane.
- Step 11** Enter ***]*** in the Attempt to access these data sets pane.
- Step 12** Click the **Save** button. A “saved changes” message is displayed.
- Step 13** Click the **Generate rules** link at the bottom of the window. The Generate Rules Program window is displayed.
- Step 14** In the Generate Rules Program window, click the **Generate** button at the bottom of the window. A message announces that rule program generation was successful.
- Step 15** Access the protected host and log on as **administrator** with the password **attack**.
- Step 16** On the protected host desktop, right-click the red CSA flag in the system tray and select **Open Agent Panel**. The CSA applet is displayed.
- Step 17** Click the **Poll** button.
- Step 18** Click the **Status** tab and verify that the last poll time is within the last minute or so.
- Step 19** Open Internet Explorer on the Student PC.
- Step 20** Enter **10.0.1.50*]*** in the Address field and press **Enter**. Note the message on the web page displayed.
- Step 21** On the CSA MC, choose **Events > Event Monitor**. The Event Monitor is displayed. Note the new event logged regarding the ***]*** attack.

Task 6: Add a File Access Control Rule to a Rule Module

This task involves adding a file access control rule to a rule module that will log all attempts to read or write to files in the inetpub\scripts directory. Complete the following steps to add a file monitor rule to a rule module:

- Step 1** Choose **Configuration > Rule Modules > Rule Module Windows**.
- Step 2** Select the **Test_Rule_Module** link. The Test Rule Module configuration page is displayed.
- Step 3** Click the **Modify rules** link. The Rule Modules window is displayed.
- Step 4** Click the **Add rule** link at the bottom of the rule list. The Add rule drop-down menu is displayed.
- Step 5** Choose **File access control** from the Add rule drop-down menu. The file access control rule configuration window is displayed.
- Step 6** Enter **Monitor script changes** in the Description field.
- Step 7** Check the **Enabled** check box to enable this rule within the policy.

- Step 8** Select **High Priority Deny** from the Take the following action drop-down menu.
- Step 9** Check the **Log** check box, but do not check the Take precedence over other High Priority Deny rules check box.
- Step 10** Select **All Applications** in the Applications in any of the selected classes pane.
- Step 11** Select the **Write File** box next to Attempt the following operations.
- Step 12** Enter **c:\inetpub\scripts*** in the On any of these files pane.
- Step 13** Click the **Save** button. A “saved changes” message is displayed.
- Step 14** Click the **Generate rules** link at the bottom of the window. The Generate Rules Program window is displayed.
- Step 15** In the Generate Rules Program window, click the **Generate** button at the bottom of the window. A message announces that rule program generation was successful.
- Step 16** Access the protected host and log on as **administrator** with the password **attack**.
- Step 17** On the protected host desktop, right-click the red CSA flag in the system tray and select **Open Agent Panel**. The CSA applet is displayed.
- Step 18** Click the **Poll** button.
- Step 19** Click the **Status** tab and verify that the last poll time is within the last minute or so.
- Step 20** Open the command prompt by choosing **Start > Run** and enter **cmd** in the Open field.
- Step 21** Click the **OK** button. The command prompt window is displayed.
- Step 22** Enter **copy c:\testfile.txt c:\inetpub\scripts** and press **Enter**. The command prompt displays “0 files copied.”
- Step 23** Close the command prompt window.
- Step 24** On the CSA MC, choose **Events > Event Monitor**. Note the new Event Log message regarding inetpub\scripts.

Task 8: Add a Network Access Control Rule to a Rule Module

This task involves adding a network access control rule to a rule module that blocks access to the SMTP service on the protected host. Complete the following steps to add a network access control rule to a rule module:

- Step 1** Open a command prompt on the Student PC by choosing **Start > Run**.
- Step 2** Enter **cmd** in the Open field and press **Enter**. A command prompt window is displayed.
- Step 3** Enter **telnet 10.0.1.50 25**. SMTP mail server connection information is displayed.
- Step 4** Enter **quit** at the command prompt (it will not show as you type it) and press **Enter**.
- Step 5** Close the command prompt window.
- Step 6** At the CSA MC, choose **Configuration > Rule Modules > Rule Module Windows**. All existing rule modules are displayed.

- Step 7** Select the **Test_Rule_Module** link. The Test Rule Module page is displayed.
- Step 8** Click the **Modify rules** link. All existing rules are shown.
- Step 9** Click the **Add rule** link at the bottom of the rule list. The Add rule drop-down menu is displayed.
- Step 10** Choose **Network access control** from the Add rule drop-down menu. The network access control rule configuration window is displayed.
- Step 11** Enter **Block SMTP access** in the Description field.
- Step 12** Verify that the **Enabled** check box is checked, to enable this rule within the policy.
- Step 13** Choose **High Priority Deny** from the Take the following action drop-down menu.
- Step 14** Verify that the **Log** check box is checked and that the Take preference over other High Priority Deny rules check box is not checked.
- Step 15** Select **All Applications** in the Applications in any of the selected classes pane.
- Step 16** Select **server** from the Attempt to act as a <value> for network services drop-down menu.
- Step 17** Enter **tcp/25** in the network services pane.
- Step 18** Verify that **0.0.0.0-255.255.255.255** is entered in the Communicating with host addresses pane.

Note The host addresses are expressed as a range, so 0.0.0.0-255.255.255.255 means all IP addresses.

- Step 19** Enter **@local** in the Using these local addresses pane.
- Step 20** Click the **Save** button. A “saved changes” message is displayed.
- Step 21** Click the **Generate rules** link at the bottom of the window. The Generate Rules Program window is displayed.
- Step 22** In the Generate Rules Program window, click the **Generate** button at the bottom of the window. A message announces that rule program generation was successful.
- Step 23** Access the protected host and log on as **administrator** with the password **attack**.
- Step 24** On the protected host desktop, right-click the red CSA flag in the system tray and select **Open Agent Panel**. The CSA applet is displayed.
- Step 25** Click the **Poll** button.
- Step 26** Click the **Status** tab and verify that the last poll time is within the last minute or so.
- Step 27** On the Student PC, open the command prompt by choosing **Start > Run** and enter **cmd** in the Open field.
- Step 28** Click the **OK** button. The command prompt window is displayed.
- Step 29** Enter **telnet 10.0.1.50 25** and press **Enter**. A message is displayed that the connection failed.

- Step 30** Close the command prompt window.
- Step 31** Close the Remote Desktop Connection to the protected host.
- Step 32** On the CSA MC, choose **Events > Event Monitor**. Note the new Event Log message regarding an attempted TCP connection.

Lab 9-1: Defining Application Classes

Complete the following lab exercise to practice what you learned in this lesson.

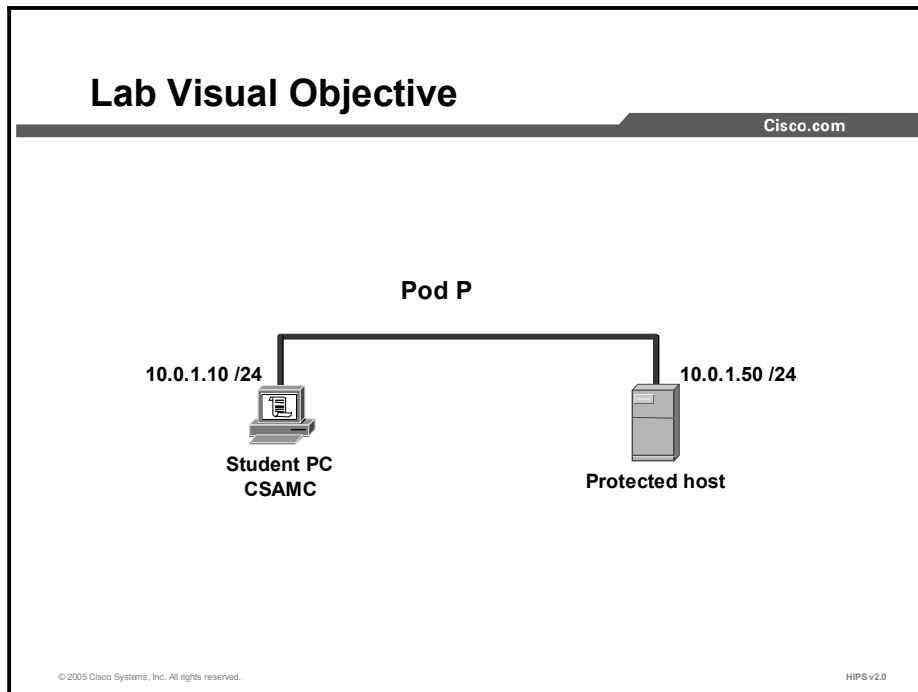
Objectives

In this lab exercise you will complete the following tasks:

- Create a static application class
- Create a dynamic application class
- Configure an application-builder rule
- Configure a rule using a dynamic application class

Visual Objectives

The following figure displays the configuration you will complete in this lab exercise.



Task 1: Create a Static Application Class

Complete the following steps to create a static application class:

- Step 1** Choose **Configuration > Applications > Application Classes [Windows]**. Existing application classes are displayed.
- Step 2** Click the **New** button to create a new application class. The new application class configuration window appears.
- Step 3** Enter **Test static application class** in the Name field and **For test purposes only** in the Description field.

- Step 4** Verify that **<All Windows>** is selected in the Operating System Target drop-down menu.
- Step 5** Select the **When created from one of the following executables** radio button.
- Step 6** Enter **winword.exe** and **notepad.exe** in the When created from one of the following executables field.
- Step 7** Do not check the After check box.
- Step 8** Leave the Seconds field blank.
- Step 9** Select the **This process and all its descendents** radio button.
- Step 10** When you are finished, click the **Save** button.

Task 2: Create a Dynamic Application Class

This task involves creating a dynamic application class. Applications will be added to the dynamic application class when their behavior matches criteria defined in the application-builder rule. Complete the following steps to create a dynamic application class:

- Step 1** Choose **Configuration > Applications > Application Classes [Windows]**. The list of existing application classes is displayed.
- Step 2** Click the **New** button to create a new application class. The New Application Class configuration window is displayed.
- Step 3** Enter **Test dyno app class** in the Name field.
- Step 4** Verify that **<All Windows>** is selected in the Operating System Target drop-down menu.
- Step 5** Select the **When dynamically defined by policy rules** radio button under the Add process to application class pane. (Do not enter any process names in the field.)
- Step 6** Check the **After** check box under Remove process from application class.
- Step 7** Enter **90** in the Seconds field.

Note Applications that are added to the Test dyno app class by virtue of their behavior will be removed after 90 seconds.

- Step 8** Select the **This process and all its descendents** radio button.
- Step 9** Click the **Save** button.

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.

Task 3: Configure an Application-Builder Rule

This task involves creating an application-builder rule in the rule module. This application-builder rule will add applications to the Test dyno app class when the applications perform a file read on a specific file: c:\test.txt. Complete the following steps to configure an application-builder rule:

- Step 1** Choose **Configuration > Rule Modules > Rule Module Windows**. The existing rule modules are displayed.
- Step 2** Select the **Test Rule Module** link. The test rule module configuration window is displayed.
- Step 3** Click the **Modify rules** link. The rules configuration window is displayed.
- Step 4** Click the **Add rule** link. The rules drop-down menu is displayed.
- Step 5** Choose **File access control** from the menu. The file access control rule configuration window is displayed.
- Step 6** Enter **Test app-builder** in the Description field.
- Step 7** Verify that the **Enabled** check box is checked.
- Step 8** Choose **Add process to application class** from the Take the following action drop-down menu.
- Step 9** Choose **Test dyno app class** from the Dynamic Application Class drop-down menu.
- Step 10** Check the **Log** check box.
- Step 11** Verify that the default option, **All Applications**, is selected in the Applications in any of the selected classes pane.
- Step 12** Click the **Read File** check box by Attempt the following operations.
- Step 13** Enter **c:\test.txt** in the On any of these files pane.
- Step 14** Click the **Save** button.

Task 4: Configure a Rule Using a Dynamic Application Class

This task involves using the Test dyno app class in a file access rule. This file access rule will block read access to a second file, c:\test2.txt, to applications that have been placed in the Test dyno app class by their behavior. Complete the following steps to configure a rule using a dynamic application class:

- Step 1** Choose **Configuration > Rule Modules > Rule Module Windows**. The existing rule modules are displayed.
- Step 2** Select the **Test Rule Module** link. The test rule module configuration window is displayed.
- Step 3** Click the **Modify rules** link. The rules configuration window is displayed.
- Step 4** Click the **Add rule** link. The rules drop-down menu is displayed.
- Step 5** Choose **File access control** from the menu. The file access control rule configuration window is displayed.
- Step 6** Enter **Block access to test2.txt** in the Description field.
- Step 7** Verify that the **Enabled** check box is checked.
- Step 8** Choose **High Priority Deny** from the Take the following action drop-down menu.

- Step 9** Verify that the **Log** check box is checked and that the Take precedence over other High Priority Deny rules check box is not checked.
- Step 10** Select **Test dyno app class** from the Applications in any of the selected classes pane.
- Step 11** Check the **Read File** check box by Attempt the following operations.
- Step 12** Enter **c:\test2.txt** in the On any of these files pane.
- Step 13** Click the **Save** button.
- Step 14** Select the **Generate rules** link. The Generate Rules Program window is displayed.
- Step 15** Click the **Generate** button. A message will announce that rule program generation was successful.
- Step 16** Access the protected host and log in as **administrator** with password **attack**.
- Step 17** Right-click the red CSA flag in the system tray and select **Open Agent Panel**. The Cisco Security Agent applet is displayed.
- Step 18** Click the **Status** link.
- Step 19** Click the **Poll** button.
- Step 20** Click the **Status** link and verify that the last poll time is within the last minute or so.
- Step 21** Open the command prompt by choosing **Start > Run** and enter **cmd** in the Open field.
- Step 22** Click the **OK** button. The command prompt window is displayed.
- Step 23** At the command prompt, enter **edit c:\test.txt** and press **Enter**. The Edit window is displayed.
- Step 24** At the edit window, enter **When this file is accessed**.
- Step 25** Save the file by pressing **Alt-F** and then pressing **Alt-S**.
- Step 26** Open a new file by pressing **Alt-F** and then pressing **Alt-N**.
- Step 27** At the edit window enter **--then this file will become inaccessible for 90 seconds**.
- Step 28** Save the new file by pressing **Alt-F** and then pressing **Alt-A**.
- Step 29** Enter **test2.txt** to name the new file, and press **Enter**.
- Step 30** Exit the edit program by pressing **Alt-F** and then pressing **Alt-X**.
- Step 31** At the command prompt, enter **c:\test2.txt** and press **Enter**. The text file is displayed.

Note If a user query window appears, select the **Yes** radio button and click **Apply**.

- Step 32** Activate the Test app-builder rule by entering **c:\test.txt** at the command prompt.

Note The cmd.exe process will now be moved to the Test dyno app class, and the Block access to test2.txt file access rule will then apply to it. Remember that cmd.exe will be moved to the Test dyno app class for only 90 seconds.

Step 33 At the command prompt, enter **c:\test2.txt** again. A message that access is denied is displayed.

Step 34 Wait for 90 seconds and then enter **c:\test2.txt** at the command prompt again. The text file is displayed.

Step 35 Close the command prompt window.

Step 36 On the CSA MC, choose **Events > Event Monitor**. Note the new Event Log message regarding the failed attempt to access c:\test2.txt with cmd.exe.

Lab 10-1: Working with Variables

Complete the following lab exercise to practice what you learned in this lesson.

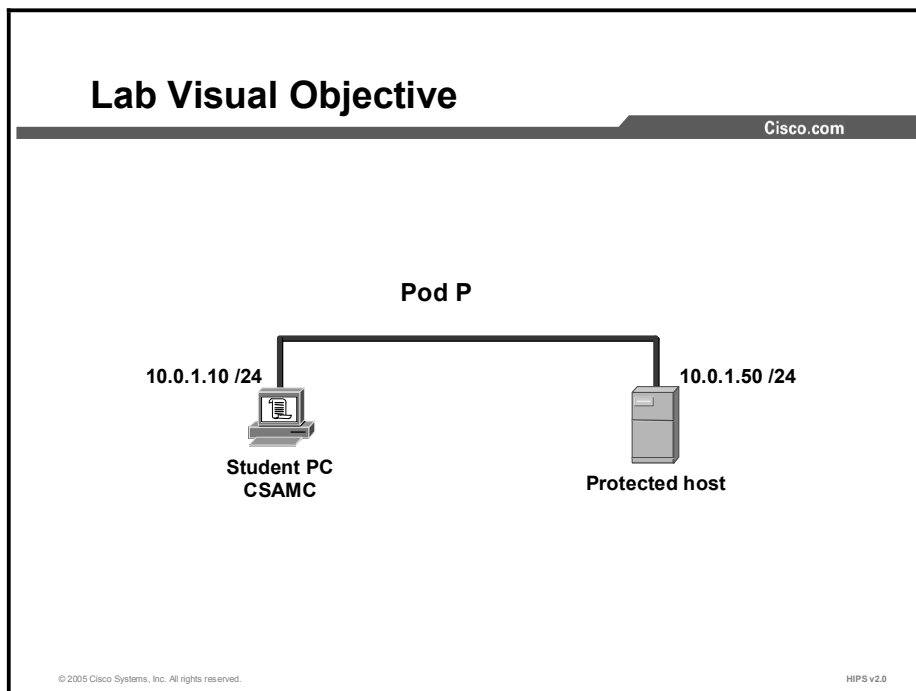
Objectives

In this lab exercise you will complete the following tasks:

- Configure a data set
- Configure a file set
- Configure a network address set
- Configure a network services set
- Configure a registry set
- Configure a COM component set

Visual Objectives

The following figure displays the configuration that you will complete in this lab exercise.



Task 1: Configure a Data Set

Complete the following steps to configure a data set that contains the metacharacter string `*[*:`

- Step 1** Choose **Configuration > Variables > Data Sets**. Existing data set configurations are shown.
- Step 2** Click the **New** button to create a new data set.
- Step 3** Enter **Test data set** in the Name field and **Test data set** in the Description field.

- Step 4** Enter `*[*` in the Patterns matching field.
- Step 5** Leave the But not field at its default value of `<none>`.
- Step 6** Click the **Save** button to save your data set in the CSA MC database.
- Step 7** Verify the creation of the new data set by choosing **Configuration > Variables > Data Sets**. Scroll down the window and verify that the Test data set is listed.

Note Clicking the Save button will save a new variable to the CSA MC database, and it will be visible under Configuration > Variables; however, the new variable will not be effective in any rules until the Generate Rules utility is run.

Task 2: Configure a File Set

Complete the following steps to configure a file set that contains all DLL files in the spool directory except sfmpsport.dll:

- Step 1** Choose **Configuration > Variables > File Sets [Windows]**. Existing file set configurations are shown.
- Step 2** Click the **New** button to create a new file set.
- Step 3** Enter **Test file set** in the Name field and **Test file set** in the Description field.
- Step 4** Verify that `<All Windows>` is selected in the Operating System Target drop-down menu.
- Step 5** Enter `C:\WINNT\system32\spool*` in the Directories matching field.
- Step 6** Enter `*.dll` in the Files Matching field.
- Step 7** Enter `sfmpsprt.dll` in the But not field.
- Step 8** When all required information is entered, click the **Save** button to save your file.
- Step 9** Verify the creation of the new file set by choosing **Configuration > Variables > File Sets [Windows]**. Scroll down the window and verify that the Test file set is listed.

Task 3: Configure a Network Address Set

Complete the following steps to configure a network address set that includes network addresses from 192.168.1.0 to 192.168.1.254:

- Step 1** Choose **Configuration > Variables > Network Address Sets**. Existing network address set configurations are shown.
- Step 2** Click the **New** button to create a new network address set.
- Step 3** Enter **Test network address set** in the Name field and **Test network address set** in the Description field.
- Step 4** Enter `192.168.1.0-254` in the Address Ranges field.
- Step 5** When all required information is entered, click the **Save** button to save your network address set.

- Step 6** Verify the creation of the new network address set by choosing **Configuration > Variables > Network Address Sets**. Scroll down the window and verify that the Test network address set is listed.

Task 4: Configure a Network Services Set

Complete the following steps to configure a network services set that includes initial connections to port 8888 and subsequent client or server connections to all ports above 1024:

- Step 1** Choose **Configuration > Variables > Network Services**. Existing network services configurations are shown.
- Step 2** Click the **New** button to create a new network services set.
- Step 3** Enter **Test network services set** in the Name field and **Test network services set** in the Description field.
- Step 4** Enter **TCP/8888** in the Protocol Ports field used for initial connection field.
- Step 5** When all required information is entered, click the **Save** button to save your network services set in the CSA MC database.
- Step 6** Verify the creation of the new network services set by choosing **Configuration > Variables > Network Services**. Scroll down the window and verify that the Test network service is listed.

Task 5: Configure a Registry Set

Complete the following steps to configure a registry set that includes all registry keys under HKEY_LOCAL_MACHINE\SOFTWARE\CISCO except for the registry keys under HKEY_LOCAL_MACHINE\SOFTWARE\CISCO\CSAgent:

- Step 1** Choose **Configuration > Variables > Registry Sets**. Existing registry set configurations are shown.
- Step 2** Click the **New** button to create a new registry set.
- Step 3** Enter **Test registry set** in the Name field and **Test registry set** in the Description field.
- Step 4** Enter **HKLM\SOFTWARE\CISCO**** in the Registry keys matching field.
- Step 5** Enter ****\CSAgent**** in the But not field, to create an exception to the registry keys entered in the previous step.
- Step 6** Leave the Registry values matching field at its default value of **<all>**.
- Step 7** Leave the But not field at its default value of **<none>**.
- Step 8** When all required information is entered, click the **Save** button to save your registry set.
- Step 9** Verify the creation of the new registry set by choosing **Configuration > Variables > Registry Sets**. Scroll down the window and verify that the Test registry set is listed.

Task 6: Configure a COM Component Set

Complete the following steps to configure a COM component set that includes the ADOB.Connection PROGID:

- Step 1** Choose **Configuration > Variables > COM Component Sets**. Existing COM component set configurations are shown.
- Step 2** Click the **New** button to create a new COM component set.
- Step 3** Enter **Test COM component set** in the Name field and **Test COM component set** in the Description field.
- Step 4** Enter **ADODB.Connection** in the PROGID's/CLSID's matching field.
- Step 5** Leave the But not field at its default value of **<none>**.
- Step 6** When all required information is entered, click the **Save** button to save your COM component set.
- Step 7** Verify the creation of the new COM component set by choosing **Configuration > Variables > COM Component Sets**. Scroll down the window and verify that the Test COM component set is listed.

Lab 11-1: Using Cisco Security Agent Analysis

Complete the following lab exercise to practice what you learned in this lesson.

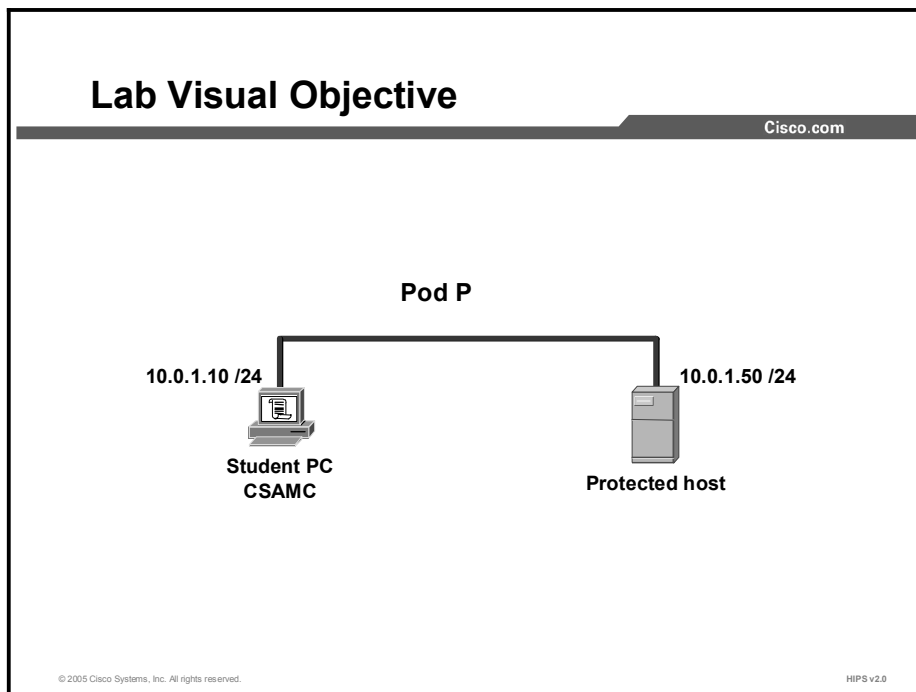
Objectives

In this lab exercise you will complete the following tasks:

- Configure an analysis job
- Start an analysis job
- Import and review an analysis job
- Review an analysis report

Visual Objectives

The following figure displays the configuration that you will complete in this lab exercise.



Task 1: Configure an Analysis Job

Complete the following steps to configure an analysis job:

- Step 1** First, build a new application class to include all applications by choosing **Configuration > Applications > Application Classes [Windows]**. All existing application classes are displayed.
- Step 2** Click the **New** button.
- Step 3** Enter **All apps** in the Name field.

- Step 4** Verify that **<All Windows>** is selected from the Operating System Target drop-down menu.
- Step 5** Select the **Insert File Set** link by the When created from one of the following executables pane. The file set drop-down menu is displayed.
- Step 6** Choose **\$ALL files** from the file set drop-down menu. “\$ALL files” will appear in the When created from one of the following executables pane.
- Step 7** Click the **This process and all its descendents** radio button. The file set drop-down menu will disappear.
- Step 8** Click the **Save** button. A message will announce “Saved changes.”
- Step 9** Choose **Analysis > Application Behavior Investigation > Behavior Analysis [Windows]**. The list of existing jobs (if any) is displayed.
- Step 10** Click the **New** button to create a new analysis job.
- Step 11** Enter **Test** in the Name field and **Analyze system behavior** in the Description field.
- Step 12** Check the **Verbose logging mode** check box.
- Step 13** Select **All apps** in the Perform an analysis of the selected application classes pane.
- Step 14** Choose the protected host from the For the selected host drop-down menu.
- Step 15** Check the **Disable policy rule enforcement** check box. Otherwise, some events may be denied by rules and the analysis may not be complete.
- Step 16** Choose **now** from the Start job at time drop-down menu.
- Step 17** Choose **after 15 minutes** from the End job at time drop-down menu.
- Step 18** Click the **Save** button. A “saved changes” message will be displayed.
- Step 19** Click the **Generate rules** link. The Generate Rules Program window will be displayed.
- Step 20** Click the **Generate** link to prepare the analysis job for distribution to the selected host when the host polls the CSA MC. A message will announce that rule program generation was successful.

Task 2: Start an Analysis Job

Complete the following steps to start an analysis job:

- Step 1** On the CSA-protected host, right-click the CSA red flag icon in the system tray.
- Step 2** Select **Open Agent Panel**. The CSA applet is displayed.
- Step 3** Select the **Status** link and click the **Poll** button. The CSA will begin an unscheduled poll of the CSA MC.
- Step 4** On the CSA MC, choose **Events > Event Monitor**. When the analysis job has been sent to the host, the Event Monitor will display the message “Logging for analysis ‘Test’ has started.”
- Step 5** Choose **Start > Run** on the Student PC when the analysis job has begun logging.

- Step 6** Enter **cmd** in the Open field.
- Step 7** Click **OK**.
- Step 8** At the command prompt, enter **telnet 10.0.1.50 53**. A blank window without a command prompt is displayed.
- Step 9** Close the command prompt window.
- Step 10** Choose **Start > Run** on the Student PC. The Run window is displayed.
- Step 11** Enter **cmd** in the Open field.
- Step 12** Click **OK**.
- Step 13** At the command prompt, enter **telnet 10.0.1.50 139**. A blank window without a command prompt is displayed.
- Step 14** Close the command prompt window.
- Step 15** Choose **Start > Run** on the Student PC. The Run window is displayed.
- Step 16** Enter **cmd** in the Open field.
- Step 17** Click **OK**.
- Step 18** At the command prompt, enter **telnet 10.0.1.50 25**. The command prompt window displays “220 ProtectedHost Microsoft ESMTP MAIL Service, Version: 5.02195.6713 ready at (day, date, time) -500.”
- Step 19** Enter **quit** at the command prompt.
- Step 20** Close the command prompt window.
- Step 21** Open Internet Explorer and enter **10.0.1.50** in the Address field. An Under Construction page is displayed.
- Step 22** Close Internet Explorer.
- Step 23** Choose **Start > Run** on the Student PC. The Run window is displayed.
- Step 24** Enter **\\10.0.1.50** in the Open field.
- Step 25** Click **OK**. A 10.0.1.50 Explorer window will be displayed.
- Step 26** Double-click the **Scheduled Tasks** folder. The Scheduled Tasks on 10.0.1.50 window is displayed.
- Step 27** Close the Scheduled Tasks on 10.0.1.50 window.
- Step 28** Access the Protected host and log in as **administrator** using the password **attack**.
- Step 29** Open Internet Explorer. A message appears that states, “The page cannot be displayed.”
- Step 30** Close Internet Explorer.
- Step 31** Open Outlook Express by choosing **Start > Programs > Outlook Express**. The Outlook Express window is displayed.
- Step 32** Click the **Inbox** icon. The Inbox is displayed.

- Step 33** Double-click the single e-mail message in the Inbox. The e-mail message is displayed in a new window.
- Step 34** Close the e-mail window, and close Outlook Express.
- Step 35** On the CSA MC, choose **Events > Event Monitor**.
- Step 36** When the Event Monitor displays the event message “Log files for analysis ‘Test’ were sent to the analysis workstation,” begin the data analysis of the logging information on the analysis workstation by choosing **Analysis > Application Behavior Investigation > Behavior Analysis [Windows]**.
- Step 37** Select the **Test** link. The Test analysis job configuration window is displayed.
- Step 38** Click the **Start analysis** button. An Event Monitor message appears, “Rule module creation for analysis ‘Test’ has started.” Analysis of the Test job may take several minutes.
- Step 39** Choose **Events > Event Monitor**. When the analysis is complete, the event monitor file displays the message “Rule module creation for analysis ‘Test’ created successfully.”

Task 3: Review a Profiler Report

Complete the following steps to review a Profiler report:

- Step 1** Choose **Analysis > Application Behavior Reports > Behavior Reports [Windows]**.
- Step 2** Choose the **Test** link.
- Step 3** Explore the reports available under **File Events**, **Registry Events**, **COM Events**, and **Network Events** in the **Analysis Reports** toolbar.
- Step 4** Choose **Summary Reports > Behavior Summary** to view all actions of any type that were taken during the analysis logging.
- Step 5** Choose **Summary Reports > Behavior Summary by Process** to view all actions of any type, segregated by process, that were taken during the analysis logging.

Lab 12-1: Using Event Logs and Generating Reports

Complete the following lab exercise to practice what you learned in this lesson.

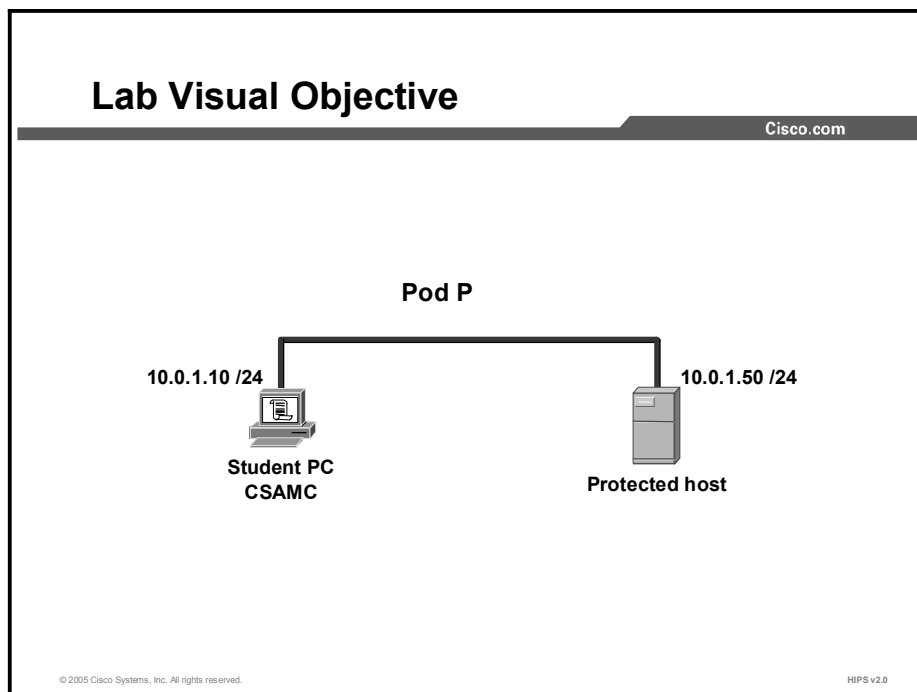
Objectives

In this lab exercise you will complete the following tasks:

- Launch attacks against the CSA-protected host
- Use the Event Log to analyze attacks
- Configure Event Log management
- Configure event sets
- Configure alerts
- Generate a report

Visual Objectives

The following figure displays the configuration that you will complete in this lab exercise.



Task 1: Launch Attacks Against the CSA-Protected Host

Complete the following steps to use hacker tools provided on the Student PC to launch attacks against your CSA-protected host:

- Step 1** Verify that the web server is running on the protected host by opening Internet Explorer on the Student PC and entering **10.0.1.50** in the Address field. Click the **Go** button. An “Under Construction” web page will be displayed.

- Step 2** Enter http://10.0.1.50/*cmd.exe in the Address field. Click the **Go** button. A web page advising that you are not authorized to view this page will be displayed.
- Step 3** Open the Event Monitor in the CSA MC by choosing **Events > Event Monitor**. Note whether there has been a new event recording your attempt to access data on the protected host.
- Step 4** Launch the CSA MC, if it is not running.
- Step 5** At the CSA MC, choose **Events > Event Monitor**. Note whether there has been a new event recording the restart of the W3SVC service by the CSA.
- Step 6** Double-click **My Computer**.
- Step 7** Choose **C: > WINNT > System32** to open the System32 folder.
- Step 8** Right-click the **tftp.exe** file.
- Step 9** Select **Rename**.
- Step 10** Enter **DLLHOST.EXE** to rename the tftp.exe file. Notice a popup window stating that an attempt to access a file has been prevented.
- Step 11** At the CSA MC, choose **Events > Event Monitor**. Note whether there has been a new event recording your attempt to rename the tftp.exe file on the protected host.

Task 2: Use Event Log to Analyze Attacks

Complete the following steps to gather information on attacks against the host using tools in the Event Log:

- Step 1** Choose **Events > Event Log** to view host events on the Event Log. The following is the information available at the Event Log:
- Date and time of attacks
 - Host
 - Severity level
 - A brief explanation of the attacks
- Step 2** Click the **(protected host name)** link in a logged event to open the Host detail window. Note the status information on the host:
- The number of events that have been issued by the host in the last 24 hours
 - Which group the host has membership
 - The number of policy rules enforced by the CSA on the host that use high priority deny as an event response action
- Step 3** Use the **Back** button in the browser to return to the Event Log.
- Step 4** Note the severity level of the events.
- Step 5** Click the **Details** link in the Event column to open the Event details window. Note the detailed information available on the event:
- The listed code of the event
 - The type of event

- Whether an executable is listed under PString

Step 6 Close the Event details window to return to the Event Log.

Step 7 Click the **Rule #** link in the Event column.

Note The **Rule #** link may not be available in all events.

Step 8 Examine the structure of the rule that triggered the event.

Step 9 Use the **Back** button in the browser to return to the Event Log.

Step 10 Click the **Wizard** link in the Event column to open the Event Management Wizard.

Note The Wizard link may not be available in all events.

Step 11 Follow the steps of the Event Management Wizard by clicking **Next** in the lower-right corner of the window to observe the available changes in event handling. Do not click **Finish** in the final Event Management Wizard window.

Step 12 Click **Cancel** to return to the Event Log.

Step 13 Click the **Find Similar** link in the Event column.

Step 14 To explore the results of running a search for similar events using different criteria, return to the Event Log by choosing **Events > Event Log**.

Step 15 Click **Change filter** in the upper-left corner of the Event Log window.

Step 16 Observe the result of different event filters by selecting the **Filter by event set** radio button.

Step 17 Choose **Significant events of all types** from the drop-down menu.

Step 18 Click **View** to display the output of the filter.

Step 19 Click **Change filter** to configure another event filter.

Step 20 Select the **Define filter** radio button.

Step 21 Enter **data** in the Filter text field.

Step 22 Verify that the **Include** radio button is selected.

Step 23 Click **View** to display the output of the filter.

Task 3: Configure Event Log Management

Complete the following steps to manage the size of the Event Log database using the Event Log Management utility:

Step 1 Choose **Events > Event Log Management** to open the Event Log Management utility.

Step 2 Click the **New** button to open the Event Log task configuration window.

- Step 3** Enter a name in the Name field and a description in the Description field for the Event Log filter.
- Step 4** Type **10** in the **Delete events after (value) days** field.
- Step 5** Click the **Save** button. Because you did not configure a deletion time, the events selected for deletion will be deleted at the default time of midnight.

Task 4: Configure Event Sets

Complete these steps to configure event sets to provide greater granularity for alerts, reports, and Event Logs:

- Step 1** Choose **Events > Event Sets** from the main menu bar. All existing event set configurations are displayed.
- Step 2** Click the **New** button to create a new event set. This action takes you to the configuration view.
- Step 3** Enter **Test data filter event set** in the Name field and enter **for test purposes only** in the Description field.
- Step 4** Select the **Include only the following selected event types** radio button.
- Step 5** Choose **Data access control: Deny action** from the pane.
- Step 6** Click the **Save** button.
- Step 7** To see the events allowed in this event set, click the **View** button. This event set is now available at the Change Filter > Policy drop-down menu.
- Step 8** Next, create another new event set. Choose **Events > Event Sets** from the main menu bar.
- Step 9** Click the **New** button to create a new event set. This action takes you to the configuration view.
- Step 10** Enter **Test alert event set** in the Name field and **for test purposes only** in the Description field.
- Step 11** Select the **Include only the following selected severity levels** radio button.
- Step 12** Select **Alert** from the severity levels pane.
- Step 13** Click the **Save** button.
- Step 14** To see the events allowed in this event set, click the **View** button.
- Step 15** Click the **Change filter** link.
- Step 16** Choose **Test data filter event set** from the drop-down menu.
- Step 17** Click **View**.
- Step 18** Open the Test data filter event set configuration window by choosing **Events > Event Sets > Test data filter event set**.
- Step 19** To see the events allowed in this event set, click the **View** button. Note the number of events in the Event Log after being filtered by the Test data filter event set in the upper-left corner of the Event Log window.

- Step 20** Click the browser **Back** button to return to the Test data filter event set configuration window.
- Step 21** Delete all events specified by this event set by clicking the **Purge events** button. A popup window asks you to verify that you want to purge these events.
- Step 22** Click the **OK** button. The window displays the message “Event purging completed successfully ([number] events deleted).”

Task 5: Configure Alerts

Complete the following steps to configure an e-mail alert for critical system events:

- Step 1** Choose **Events > Alerts** from the main menu bar.
- Step 2** Click the **New** button to create a new alert. This action takes you to the alerts configuration window.
- Step 3** In the alerts configuration window, enter “**Critical systems**” in the Name field and “**Critical systems events**” in the Description field.
- Step 4** Select **Events from mission critical systems** from the Send alerts for the following event set list pane.
- Step 5** Check the **Email** check box. Enter admin@yournetwork.com in the Recipient(s) email address(es) field. Enter CSAMC@yournetwork.com in the Sender address to use field. Enter mailserver.yournetwork.com in the Address of mail server field.
- Step 6** Click the **Save** button to save the alert.

Task 6: Generate a Report

Complete the following steps to generate a report based on event severity:

- Step 1** Choose **Reports > Events by Severity** to display all existing reports.
- Step 2** Click the **New** button to create a new report. This action takes you to the events by severity configuration window.
- Step 3** Enter **Test events report** in the Name field and **Report on testing events** in the Description field.
- Step 4** Choose **Event Filter > All events** from the drop-down menu.
- Step 5** In the Sort by drop-down menu, leave the default value of **Time** selected.
- Step 6** Check the **Ascending** check box to select the viewing order of your report.
- Step 7** At the Viewer type drop-down menu, leave the default selection of **ActiveX**. ActiveX is the recommended viewer.
- Step 8** Click the **Save** button to save the parameters that you have just configured.
- Step 9** Click the **View Report** button to display the report in a new browser window.