



# Rule Module Configuration

---

## Overview

Rule modules consist of one or more rules. One or more rule modules are meant to be attached to a policy. This module of rules is generally configured for a particular “modular” purpose. It is in this manner that several rules can be moved together from one policy to another or exist as part of several policies.

Rule module are generally OS specific while policies are not. This way, you can scale a great many rule modules to a lesser number of policies to simplify your basic product configuration view. For example, you could have one policy for all Apache servers, but that policy would consist of several OS specific rule modules containing hundreds of rules. As an administrator, you may only be interested in your Apache policy which you can attach to a general servers groups and deploy in this manner.

This section contains the following topics.

- [About Rule Modules and Rules, page 5-3](#)
- [Rule Module Components, page 5-5](#)
- [Configuring Rule Modules, page 5-5](#)
- [Using Test Mode, page 5-7](#)
- [Using Learn Mode, page 5-11](#)
- [Setting State Conditions, page 5-13](#)
- [System State Sets, page 5-13](#)

- User State Sets, page 5-19
- Adding Rules to a Rule Module, page 5-21
- Filtering the Rules Display, page 5-24
- Copying Rules between Modules, page 5-24
- Comparing Configurations, page 5-27
- Merging or Copying Rule Modules, page 5-29
- View Change History, page 5-29
- Explanation of Rules, page 5-30
- Consistency Check, page 5-31
- Attaching Rule Modules to Policies, page 5-32
- Generating Rule Programs, page 5-34
- Common Rule Page Configuration Items, page 5-36
- Rules: Action Definitions, page 5-38
- Rules: Manipulating Precedence, page 5-41
- The Monitor Action, page 5-42
- Using the Set Action, page 5-43
- Querying the User, page 5-48
- Caching Responses, page 5-51

# About Rule Modules and Rules

Rules are the foundation of your security policies. CSA MC lets you create several rule types. Each rule type requires you to enter varying combinations of information using a specific syntax. Most Policies and Rule Modules use a combination of *Enforce* and *Detect* rules. Enforce rules are primarily access control rules that allow, deny, or terminate a process. Detect rules are monitoring, 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.

For example, the following basic *enforcement* rules require information as follows:

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

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

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

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

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

- the action you are allowing or denying
- the application attempting to write to the registry keys and values
- the modification of Key/Value pairs

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

- the action you are allowing or denying

- the application accessing the COM component

Other types of enforcement rules shipped with CSA MC provide event correlation and heuristic features which can be enabled on a per group basis, like portscan detection, SYN flood protection, the prevention of predictable TCP sequence numbers, and the blocking of malformed IP packets. (These features are located on the Network shield rule page.) This is especially useful for network servers. (See [Chapter 7, “Using Event Correlation”](#) for more information.)

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 Module Components

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

## Configuring Rule Modules

Rule modules are the building blocks for your policies. Modules are made of several different types of rules. See [Chapter 6, “Available Rule Types”](#) for information on rule types.

**Note**

Carefully read [Rules: Action Options and Precedence, page 5-37](#) so that you will understand how rule precedence works once policies are deployed. You should also refer to the chapter on configuration Variables ([Chapter 9, “Configuring Variables”](#)) to help you understand the information required by the rule text fields.

**Caution**

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

To configure a rule module, do the following.

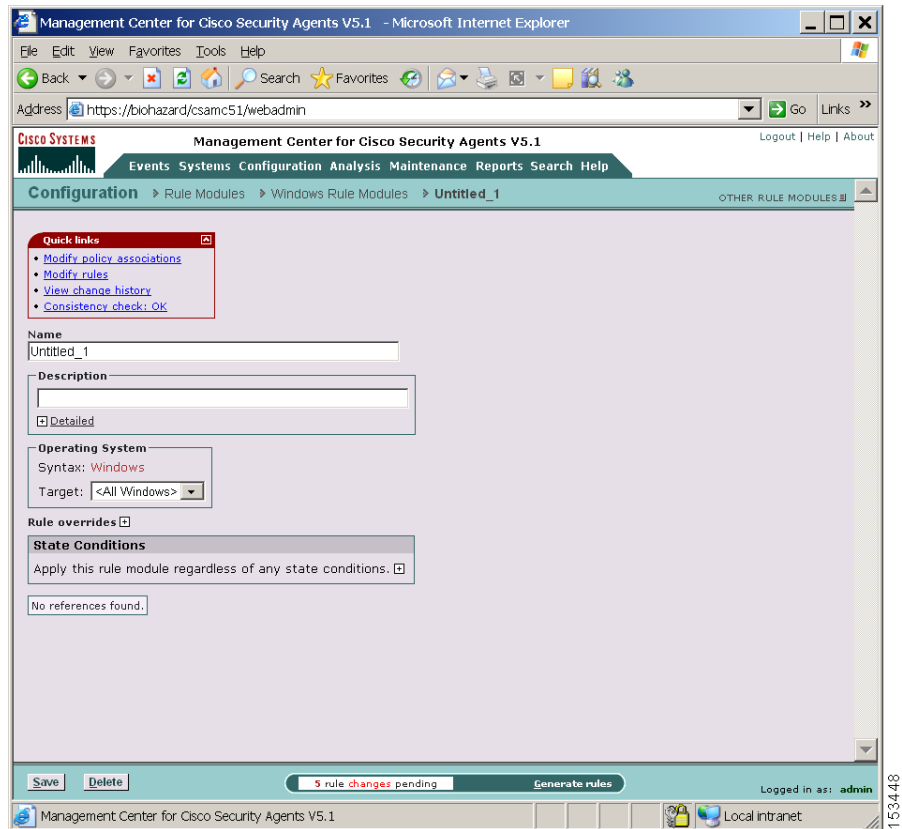
- Step 1** Move the mouse over **Configuration>Rule Modules** in the menu bar. If you have not set OS admin preferences, you must select whether this is a Windows or a UNIX rule module from the cascading menu that appears. When you make a selection, the list of existing rule modules is displayed. CSA MC ships with several pre-configured 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 listbox when you are attaching modules to policies.
- Step 4** Enter a **Description** of your module. This description is visible in the rule module list view. Optionally, expand the **+Detailed** field to enter a longer description.
- Step 5** In the **Operating System** box, optionally, you can select to target this module for a specific operating system within your Windows or UNIX classification.
- Step 6** Optionally, available under **Rule overrides**, 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 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. You can also apply test mode on the group level. See [Using Test Mode, page 5-7](#) for details.
- Step 7** Optionally, available under **Rule overrides**, you can put this rule module into **Learn mode**. This way, you can localize policy rules on the agent and prevent the flurry of query pop-ups that can appear to a user when the agent is first installed. Learn mode works in a specific manner, in combination with deployed query user rules. See [Using Learn Mode, page 5-11](#) for details.
- Step 8** Click the **Save** button.
- Step 9** Now you add rules to your module. Click the **Modify rules** link at the top of the page.  
Refer to the following sections for details on adding, copying, and configuring rules.
- Step 10** Optionally, you can impose configured **State Conditions** on rule modules. You can configuration System State conditions or User State conditions.

Figure 5-1 New Rule Module



## Using Test Mode

Test Mode is useful when you are installing a new host or are modifying a host configuration and want to understand the ramifications without actually impacting host operation. 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 Group test mode (available from the Rule Overrides category), you may 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 if test mode is on or off.

## Group Test Mode

You can turn on test mode in two places within the MC. If it is enabled on the group level (see [Figure 5-2](#)), 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 are in test mode (even if the host is part of another group that does not have test mode selected), not just the policies applied to the test group. Therefore, test mode applies to the host as a whole, not to specific policies.



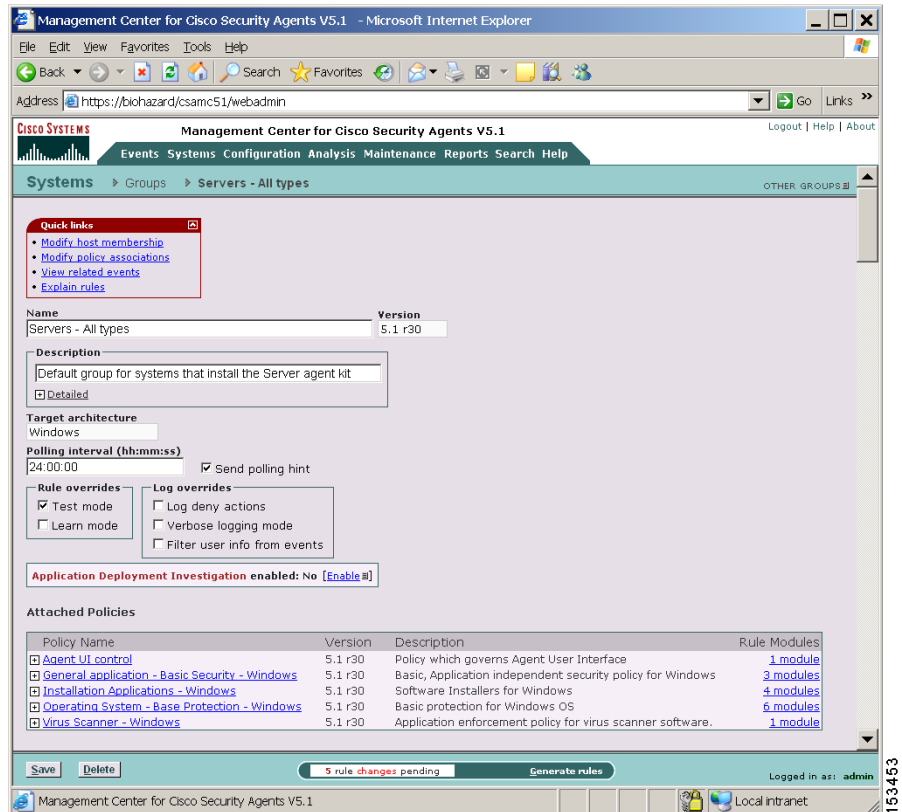
### Note

---

Using the Hosts Managing Tasks page, you can configure “timed” Test Mode. Basically, you can configure a task that causes hosts to move in and out of selected groups at timed intervals. This way, you can have all new hosts move out of a Test Mode group and into a live mode group after a 30 day pilot, for example. Refer to [Host Managing Tasks, page 3-47](#) for configuration information.

---

Figure 5-2 Group Test Mode



153453

## Rule Module Test Mode

You can also use test mode on the rule module level (see [Figure 5-3](#)). 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.

**Figure 5-3 Rule Module Test Mode**

The screenshot shows the Management Center for Cisco Security Agents V5.1 web interface. The breadcrumb navigation is Configuration > Rule Modules > Windows Rule Modules > System Hardening Module. The page displays the configuration for the System Hardening Module (Version 5.1 r30). The "Rule overrides" section has the "Test mode" checkbox checked. The "State Conditions" section is set to "Apply this rule module only if the following state conditions are met". Under "System State Conditions", the "AND" condition is selected, and a list of system state sets is shown, including "Cisco Trust Agent Infected Posture [V5.1 r30]", "Cisco Trust Agent Quarantine Posture [V5.1 r30]", "Installation in progress [V5.1 r30]", and "Management Center not reachable [V5.1 r30]". A "New" button is visible next to the list. At the bottom, a status bar indicates "5 rule changes pending" and a "Generate rules" button is present. The user is logged in as "admin".

153462

# Using Learn Mode

Learn mode is intended to localize policies on individual systems, eliminating the initial flurry of pop-up queries that users may experience when the agent is first installed on a system. This flurry of queries is a result of query rules that are deployed to let end users decide when an unknown system action is normal or abnormal. When the agent is first deployed, these pop-up queries can be numerous since the agent is seeing system actions for the first time. Unfortunately, this initial bombardment of queries for actions that are likely benign (in most cases) can train the user to respond Yes to the majority of queries they see.

To solve this problem without removing useful query rules from your deployment, you can enable Learn mode for a temporary period of time. Learn mode directs the agent not to display query pop-ups, and to instead take an immediate *Allow* query response when a query rule is triggered, and to persistently save the allow response.

If you intend to use Learn mode, the queries that qualify for learning must have certain options enabled. While Learn mode is enabled from either the Group page or the Rule module page, queries are defined from the Configuration>Variables>Query Settings menu. There are several pre-configured queries listed on the Query Settings page. In order for Learn mode to work, you must do the following:

- Enable **Learn mode**, located in the Rule overrides section, on either the Group page or on a specific Rule module page.

Like Test mode, Learn mode can be enabled for all persistent queries in all policies running on hosts (Group Learn mode) or enabled for specific persistent queries located in a particular Rule module (Rule module Learn mode).

- Access the Query Setting configuration page (see [Query Settings, page 9-25](#)) for the query you want to deploy for learning. A query qualifies for learning if:
  - The **Enable “Don’t ask again” option** is selected
  - **Allow** is selected as one of the Allowed query actions

Once query responses are taken, and Learn mode is turned off, the majority of queries no longer appear and system security provided by the agent is normalized to the individual system. At this point, users should only see query pop-ups for unusual or suspicious system behavior.

## Additional Notes on Learn Mode

- All qualifying queries will be answered with an automatic and persistent Yes for the time frame that Learn mode is enabled. (You must remember to turn Learn mode off after a reasonable amount of time.)
- An event is logged to the MC Event Log when a query is triggered and learned. This event is formatted similar to most events, but it explains that a query response was taken and that learn mode was turned on for the query in question.
- The learned allow responses are displayed in the agent User Query Responses window as they occur.
- If both Test mode and Learn mode are enabled for a query rule, learning does not occur.
- Using the Hosts Managing Tasks page, you can configure “timed” Learn Mode. Basically, you can configure a task that causes hosts to move in and out of selected groups at timed intervals. This way, you can have all new hosts move out of a Learn Mode group after a set time. Refer to [Host Managing Tasks, page 3-47](#) for configuration information.
- When in Learn Mode, the agent learns the various applications that are running on the local machine. It also learns any unusual systems calls that are typically exhibited by the machine. Consequently, any applications that have been seen running on the system during the Learn Mode period will not trigger or be classified as a “first time execute” application when it is first run after the learning period. And any unusual systems calls that were observed during learning will not trigger the unusual system calls detection in the System API rule.



---

**Note**

After the agent is initially installed, there is a non-configurable, automatic, normalization learning period of 72 hours of running time. This learning is for both running applications and unusual system calls. You can use the Reset Cisco Security Agent, Learned Information checkbox to clear all the initial learning and to start the automatic 72 hour learning period again.

---

# Setting 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 Sets

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. For example, if you apply a system state to a rule module, you can dynamically “activate” and “deactivate” rules modules based on the changing state of a system. For example, you can apply special boot time rules that apply only at boot time. After booting is complete, normal operation rule modules are applied. Also, for example, you can apply a looser set of rules if an installation is occurring on a system. Once the installation is complete, a more stringent, normal operating set of rule modules are applied.

**Note**

You do not have to specify a setting for every field in the System State page. If you do specify multiple settings, note that within single fields, multiple selections are “or-ed”. If settings are specified for multiple fields, they are “and-ed”.

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. See [Figure 5-4](#).
- Step 3** Enter a unique **Name** for your system state. You will select this name in the Rule Module page. Names are case insensitive, must start with an alphabetic character, can be up to 64 characters long and can include hyphens, underscores, and spaces.
- Step 4** Enter a **Description**.

- Step 5** In the Network Admission Control section, you can select one or more (Hold the Shift key down while you click the mouse to choose multiple, concurrent options. Hold the Ctrl key down to select non-concurrent options.) **Cisco Trust Agent posture** state conditions for a system to ensure that corporate security requirements are met on that system. This feature works in conjunction with the capabilities of Cisco's Network Admission Control (NAC) functionality. See [Creating Agent Kits, page 3-10](#) and [Generate an Alert Log File for Third Party Applications, page 10-47](#) for more information.



---

**Note** Currently, the Cisco Trust Agent (an optionally installed product available from Cisco Systems) is only supported on Windows Linux 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 anti-virus software that is not up-to-date or is disabled, the Cisco Trust Agent can report this status to ACS which can then return an “Unknown” or a “Quarantine” state. 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 ACS documentation for information on posture states and what they mean. Possible posture states are:

- <Don't care>—This state is not provided by 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.
- Transition - The Host is in the process of having its posture checked and is given interim access pending a result from a full posture validation. A transition result may be applicable when a host is booting and requires restricted access to the network to complete the booting process , but not all

posture information is available, for example, Windows machines need access to domain controller before user has logged in and before all posture applications may be running.

- **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 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 you use other criteria.

**Step 6** In the System Security section, you can select one or more (Hold the Shift key down while you click the mouse to choose multiple, concurrent options. Hold the Ctrl key down to select non-concurrent options.) **Security level** conditions. If the end user has an agent UI, you can have a Security level condition apply which allows the user to set the security sidebar on their UI to a specific level. This provides some degree of user control to manage false positives or to control security when operating remotely or on the local network. This allows the user to decide, again to some degree, how much security they require.

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

**Step 8** In the System Location section, you can use the **DNS suffix matching** field to set a condition based on DNS server domain. It is the suffix of the DNS server used to resolve names that this field refers to. If any DNS server suffix (e.g. 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 you configure.

**Step 9** In the Additional State Conditions section, you click the **Add State** link to add one or more of the following additional states to this page. (Use the pulldown menus that appear to select an option. Then use the pulldown to the right of your selected option to choose one of the following settings: <Don’t care>, Yes, No.)

- Select the **Insecure boot detected** option to set a state condition to apply if a previous system boot occurred in a non-standard manner. For example, the system was booted from a peripheral device (CD ROM) rather than from the

hard drive. This type of boot can be considered non-standard and therefore possibly suspicious. (This is one way of introducing a Trojan to a system.) This type of peripheral device insecure boot detection works in conjunction with a particular type of compatible BIOS on compliant systems. The compatible BIOS detects a non-standard boot and on the next normal boot, if you have an appropriately configured Kernel Protection rule (see [Kernel Protection, page 6-54](#)), a message is sent to the MC which logs this insecure boot detection. This, in turn, causes the system state (if configured) to trigger. A Safe Mode boot also falls into this insecure boot category. Compatible BIOS is not required for a Safe Mode boot detection.

This state condition could be met if you are using a “Set-detected boot-as insecure” Kernel Protection rule. When this “Set” rule type triggers, the system state consequently takes effect. See [Using the Set Action, page 5-43](#) for more information on this setting.

Note that this is a persistent state. This state condition, once set, can only be removed by using the Reset feature. See [Resetting Cisco Security Agents, page 3-8](#). (Most states are not persistent in this way. Most states can be switched in and out using rule triggers. A persistent states can only be switched “on” using rules and must be switched “off” manually.)

- 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. See [Installation Applications Policy, page 7-4](#) for more information.
- 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. When the agent service first starts, it assumes that the management center is unreachable. When it attempts to communicate with the management center to receive rule changes or to upload events, if it can communicate with the management center at that time, it is then considered reachable.
- 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 **Unprotected access detected** option to set a state condition when an application or service, or other system component that is marked as Unprotected does not have a corresponding Protected rule and is therefore not being protected by the agent. See [Attribute: detected access, page 5-44](#) for more information on this setting.
- Select the **Untrusted 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. This state condition could be met if you are using a “Set-detected rootkit-Untrusted” rule in a rule module. When this “Set” rule type triggers, the system state consequently takes effect. See [Attribute: detected rootkit, page 5-45](#) for more information on this setting.

Note that this is a persistent state. This state condition, once set, can only be removed by using the Reset feature. See [Resetting Cisco Security Agents, page 3-8](#). (Most states are not persistent in this way. Most states can be switched in and out using rule triggers. A persistent states can only be switched “on” using rules and must be switched “off” manually.)

- 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. This state condition could be met if you are using the “Suspected Virus Applications” built-in dynamic application class in a rule module. When a process is added to the dynamic Suspected virus applications class, this system state condition is triggered. See [Built-in Configurable Application Classes, page 8-7](#) for details on this application class.

**Step 10** Click the **Save** button.



**Note**

---

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

---



**Caution**

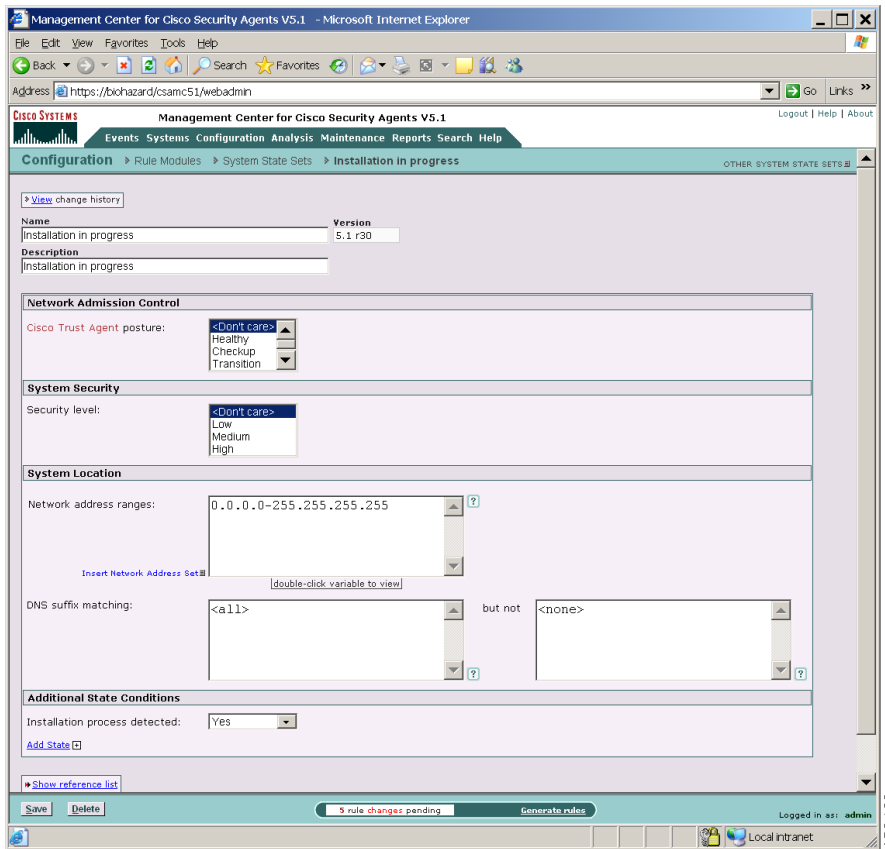
---

Remote VPN Clients - System Location and Management Center Reachable system states are checked by the agent whenever the network configuration changes on the system. Some VPN clients may make network configuration changes on a system that cause a system state to trigger. Such VPN clients can make use of System Location and Management Center Reachable settings to change the policy depending upon whether a tunnel is up or not. Other VPN clients, such as the Cisco VPN client V3.6, do not change network configurations

---

on a system. Therefore, you cannot use System Location and Management Center Reachable states to detect tunnels with these types of VPN clients. You should understand how your VPN client operates if you want to use these system states.

**Figure 5-4** System State Conditions



## User State Sets

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 in 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 (e.g. FrontPage) are allowed to alter web server content.

Another example of appropriate user state setting usage is a situation where 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. See [Figure 5-5](#).
  - Step 3** Enter a unique **Name** for your user state. You will select this name in the Rule Module 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\user account. For example, entries in this field might appear as follows:
    - `Domain_Accounting\Administrator` This represents the administrator in the Windows domain "Domain\_Accounting."
    - `W2K-jefe\Administrator` This represents user "Administrator" defined locally on the computer "W2K-jefe."
    - `*\Administrator` This represents any user administrator.

- `Domain_Accounting\*` This represents all users in the domain "Domain\_Accounting".

You can use wildcards in the Users matching/but not fields.

**Step 6** You can use the **but not** field to make specific exclusions to user matching parameters 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:

- `NT AUTHORITY\SYSTEM`
- `Domain_Accounting\Administrators`

For Windows, you can also enter SID (Security Identifier) numerical classifications into the Group 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/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**


---

User and Group names are case sensitive for UNIX. They are not case sensitive for Windows.

---

**Note**


---

It is recommended that you use Group permissions rather than User permissions because Group designations are more widely applicable.

---

**Step 8** You can use the **but not** field to make specific exclusions to group matching parameters you configure.

**Step 9** Click the **Save** button.

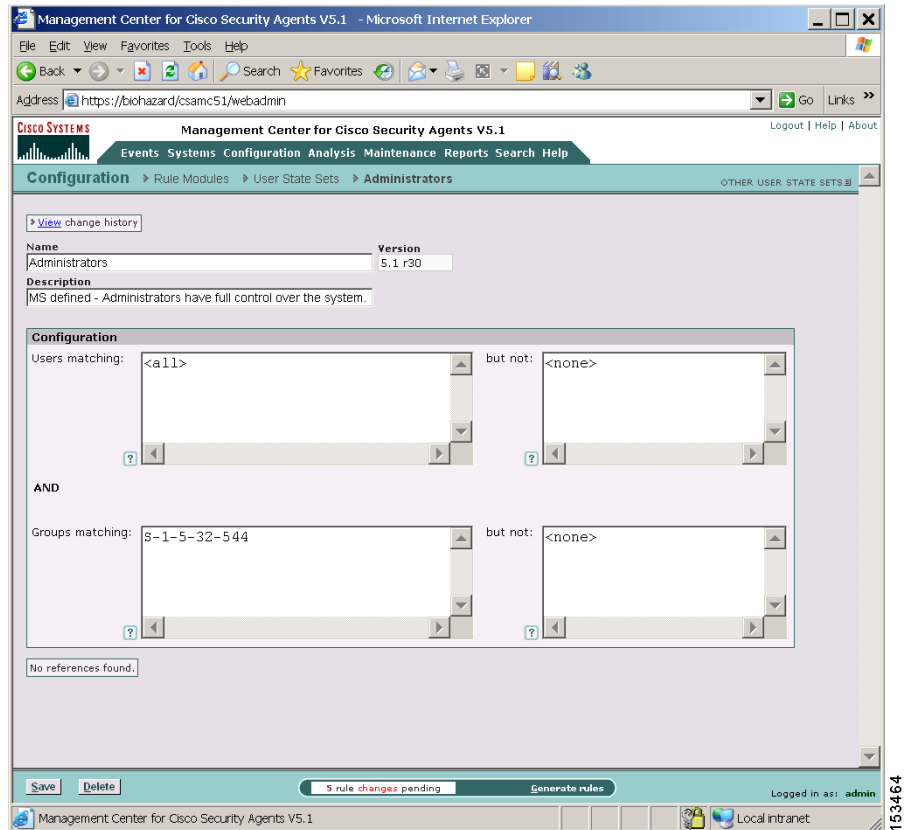
**Note**


---

In addition to user information found in logged events, you may use the host diagnostic feature to retrieve a record of observed user/group credentials on a given host. This may be useful in troubleshooting policies. It's important to note that the record of credentials that is displayed does not imply that a rule fired in that user's context.

---

Figure 5-5 User State Conditions

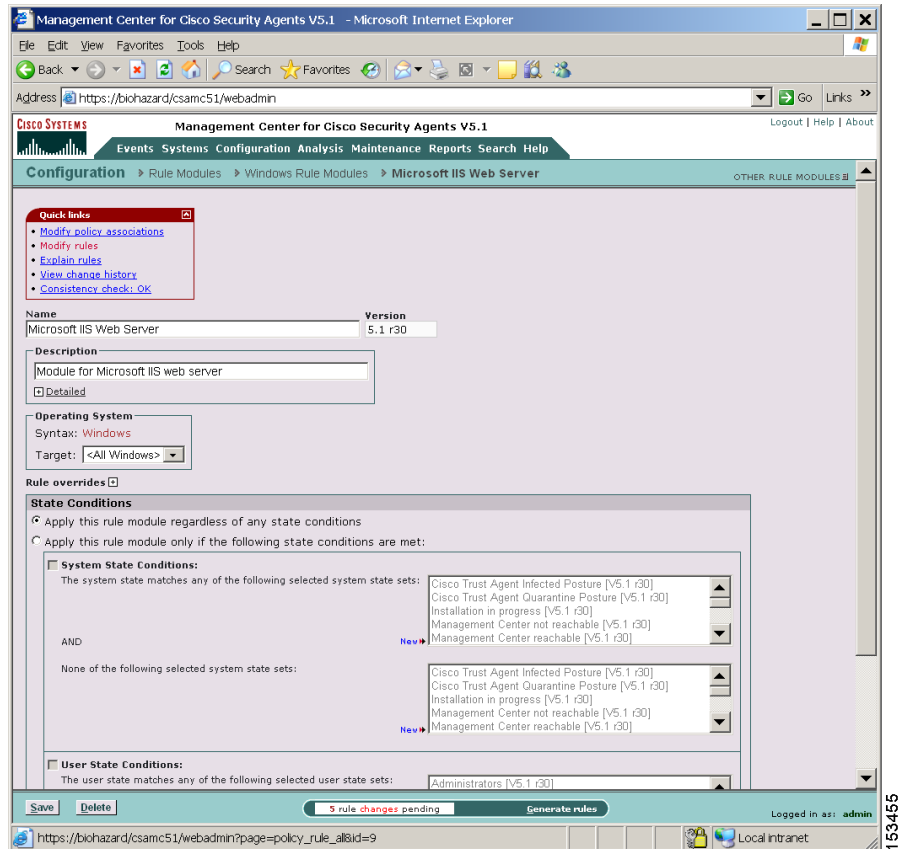


## Adding Rules to a Rule Module

First, click the **Modify rules** link at the top of the Rule Module page to go to the Rules page. See [Figure 5-6](#).

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.

Figure 5-6 Rule Module, Modify Rules

**Note**

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

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 checkbox for the rule you want to enable or disable and click the corresponding button. See [Figure 5-7](#).

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

The **Events** column in the Rules section (see [Figure 5-8](#)) 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.

**Figure 5-7 Rules List**

The screenshot shows the Management Center for Cisco Security Agents V5.1 interface. The main content area displays a list of rules for the Microsoft IIS Web Server [V5.1 r30]. The rules are listed in a table with the following columns: ID, Type, Events, Status, Action, Log, and Description. The 'Events' column contains a number of events generated by each rule, and the 'Status' column shows whether the rule is enabled or disabled. The 'Action' column shows the action taken by the rule, and the 'Log' column shows whether the rule is logged. The 'Description' column provides a brief description of the rule.

ID	Type	Events	Status	Action	Log	Description
81	File access control	Enabled	Enabled	✓	✗	IIS Server, read/write News files (*.hdr, *.ord and *.lst)
82	Network access control	Enabled	Enabled	✓	✗	IIS Server, server for HTTP and FTP services
87	Application control	Enabled	Enabled	✓	✗	IIS Web Server Dynamic Applications, invoke IIS Web Server in Isolation Mode
89	File access control	Enabled	Enabled	✓	✗	MS Management applications, read/write IIS directories
91	Registry access control	Enabled	Enabled	✓	✗	MS Management applications, write IIS keys
92	File access control	Enabled	Enabled	✓	✗	IIS Server, read/write FTP root directory
96	Network access control	Enabled	Enabled	✓	✗	Web browsers, client for HTTP services
97	Network access control	Enabled	Enabled	✓	✗	IIS Server and COM+ surrogate applications, client for HTTP services
98	File access control	Enabled	Enabled	✓	✗	IIS Server, read/write Web Server writeable files
90	File access control	Enabled	Enabled	✗	✗	All applications, write IIS executable directories
85	File access control	Enabled	Enabled	✗	✗	Vulnerable applications, write IIS data directories
93	Registry access control	Enabled	Enabled	✗	✗	All applications, write IIS keys
99	File access control	Enabled	Enabled	✗	✗	IIS Server and descendants, write all files
86	Service restart	Enabled	Enabled	-	✗	FTP Publishing Service (FTP)
88	Service restart	Enabled	Enabled	-	✗	World Wide Web Publishing Service (HTTP) - Windows 2000
94	Service restart	Enabled	Enabled	-	✗	World Wide Web Publishing Service (HTTP) - Windows XP
83	File access control	Enabled	Enabled	+	✗	Add to IIS Web Server Dynamic Application in Isolation Mode
84	File access control	Enabled	Enabled	+	✗	Add to IIS Web Server Dynamic Application
95	NT Event log	Enabled	Enabled	-	-	MS IIS Server events

At the bottom of the interface, there are buttons for 'Delete', 'Enable', and 'Disable'. A status bar indicates '5 rule changes pending' and a 'Generate rules' button. The user is logged in as 'admin'.

## Filtering the Rules Display

The Groups configuration page, Policy configuration page, and the Rules configuration page each display a table listing either the rules attached to the group or the rules included in the module (see [Figure 5-8](#)). On all of these pages, there is a **View All rules** item above the table. Clicking the **All** link here lets you filter your view of this rule list by selected rule type. When you click **All**, a pop-up appears listing the rule types present in the module or modules. Select a rule type from the pop-up, and that is now the only rule type displayed in the table. You can also select to view only enabled rules by selecting the **show enabled rules only** checkbox and then select the rule type you wish to view.



### Note

---

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

---

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

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

## Copying Rules between Modules

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

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

- 
- Step 1** From the Rule Module page (see [Figure 5-8](#)), select the checkbox for the rule or rules you want to copy to another module.

**Step 2** Beside the Copy button, **to** is the default selection in the pulldown menu. (Do not change this for copying individual rules between modules.) From the **rule module** pulldown 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 pulldown list, select the current module you are in from that same pulldown. Selected rules are cloned within the same module when you click the Copy button.

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

Figure 5-8 Copy Rules between Modules

The screenshot shows the Management Center for Cisco Security Agents V5.1 interface. The breadcrumb navigation is: Configuration > Rule Modules > Windows Rule Modules > Microsoft IIS Web Server [V5.1 r30] > Rules. The interface displays a table of 19 rules with columns for ID, Type, Events, Status, Action, Log, and Description. A 'Copy rule' dialog box is open, showing a 'Copy' button, a 'to' dropdown menu, and a text field containing 'rule module Network Application Classification Module [V5.1 r30]'. The status bar at the bottom indicates '5 rule changes pending' and 'Generate rules'.

ID	Type	Events	Status	Action	Log	Description
81	File access control		Enabled	✓	✗	IIS Server, read/write News files (*.hdr, *.ord and *.lst)
82	Network access control		Enabled	✓	✗	IIS Server, server for HTTP and FTP services
87	Application control		Enabled	✓	✗	IIS Web Server Dynamic Applications, invoke IIS Web Server in Isolation Mode
89	File access control		Enabled	✓	✗	MS Management applications, read/write IIS directories
91	Registry access control		Enabled	✓	✗	MS Management applications, write IIS keys
92	File access control		Enabled	✓	✗	IIS Server, read/write FTP root directory
96	Network access control		Enabled	✓	✗	Web browsers, client for HTTP services
97	Network access control		Enabled	✓	✗	IIS Server and COM+ surrogate applications, client for HTTP services
98	File access control		Enabled	✓	✗	IIS Server, read/write Web Server writeable files
90	File access control		Enabled	⚠	✗	All applications, write IIS executable directories
85	File access control		Enabled	✗	✗	Vulnerable applications, write IIS data directories
93	Registry access control		Enabled	✗	✗	All applications, write IIS keys
99	File access control		Enabled	✗	✗	IIS Server and descendants, write all files
86	Service restart		Enabled	-	✗	FTP Publishing Service (FTP)
88	Service restart		Enabled	-	✗	World Wide Web Publishing Service (HTTP) - Windows 2000
94	Service restart		Enabled	-	✗	World Wide Web Publishing Service (HTTP) - Windows XP
83	File access control		Enabled	+	✗	Add to IIS Web Server Dynamic Application in Isolation Mode
84	File access control		Enabled	+	✗	Add to IIS Web Server Dynamic Application
95	NT Event log		Enabled	-	-	MS IIS Server events

Copy rule to rule module Network Application Classification Module [V5.1 r30]

5 rule changes pending Generate rules

153457

## Comparing Configurations

When you select the checkbox next to 2 items (you cannot compare more than 2 configurations at a time) and click the **Compare** button, CSA MC displays the configurations side by side and highlights the differences in red (see [Figure 5-9](#)). Once you've 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. Additionally, you can attach and detach groups and policies. (You can compare application classes and variables, but you can only copy and merge rules from the compare page.)

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

Feature notes:

- 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 modules and no corresponding similar rule in the second modules, 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/Deny actions. Logging and allow/deny actions change the priority of the rule within the policy. If the priority is not the same for each rule, they are not displayed side by side.

Figure 5-9 Compare Rule Modules

The screenshot displays the Management Center for Cisco Security Agents V5.1 interface in Microsoft Internet Explorer. The browser address bar shows the URL: https://biohazard/csamc51/webadmin. The page title is 'Management Center for Cisco Security Agents V5.1'. The navigation menu includes: Configuration, Rule Modules, Compare, Email Client Module - High Security [W, V5.1 r30] and Email Client Module - Medium Security [W, V5.1 r30].

The main content area compares two rule modules:

Property	Email Client Module - High Security	Email Client Module - Medium Security
<b>Name</b>	Email Client Module - High Security	Email Client Module - Medium Security
<b>Version</b>	5.1 r30	5.1 r30
<b>Description</b>	Email client behavior enforcement, when a high level of security is requested (or required)	Email client behavior enforcement, when a medium level of security is requested (and a high level is not required)
<b>Detailed description</b>		
<b>Syntax</b>	Windows	Windows
<b>Operating System</b>	All OS types	All OS types
<b>Test mode</b>	No	No
<b>Learn mode</b>	No	No
<b>Included system state sets</b>	<ul style="list-style-type: none"> <li>\$Cisco Trust Agent Infected Posture [V5.1 r30]</li> <li>\$Cisco Trust Agent Quarantine Posture [V5.1 r30]</li> <li>\$Security Level High [V5.1 r30]</li> <li>\$Virus detected [V5.1 r30]</li> </ul>	\$Security Level Medium [V5.1 r30]
<b>Excluded system state sets</b>		<ul style="list-style-type: none"> <li>\$Cisco Trust Agent Infected Posture [V5.1 r30]</li> <li>\$Cisco Trust Agent Quarantine Posture [V5.1 r30]</li> <li>\$Virus detected [V5.1 r30]</li> </ul>
<b>User state sets</b>		
<b>Rules</b>	<a href="#">7 items</a>	<a href="#">7 items</a>

Below the comparison, there are instructions:

- Use the checkboxes to merge or copy rules to a new rule module or to existing rule modules.
- All rules displayed ([show only similar rules with detailed differences](#))

The 'Application control' section is expanded, showing a table with columns for 'Application control' and 'Description'. The 'Description' column contains the following text:

<b>Enabled</b>	<input type="checkbox"/> Yes
<b>Description</b>	Email Applications, invoking All Applications except MS Office
<b>Detailed Description</b>	
<b>Action</b>	Query User (Default Deny)
<b>Query Setting</b>	Application Control - Tied to parent and child name [V5.1 r30]
<b>Log</b>	Yes
<b>Application Classes</b>	<All Applications>
<b>Excluded Application Classes</b>	MS Office applications [V5.1 r30]
<b>Parent Application Classes</b>	Email applications [V5.1 r30]

At the bottom of the interface, there are buttons for 'Copy', 'Delete', and 'Generate rules'. A status bar indicates '5 rule changes pending'. The user is logged in as 'admin'.

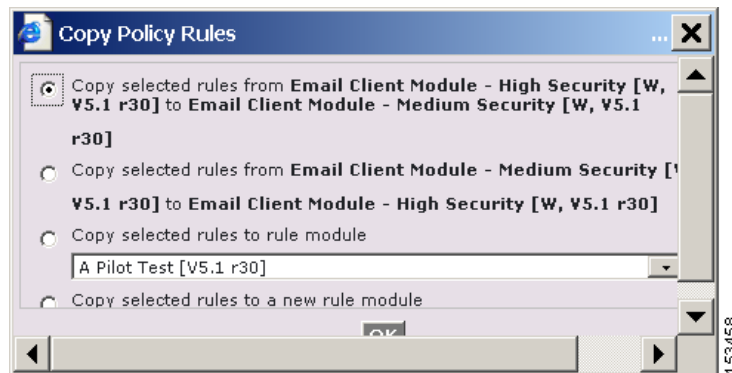
150456

## Merging or Copying Rule Modules

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

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

**Figure 5-10** Copy Rule Module Pop-up Box



## View Change History

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

## Explanation of Rules

CSA MC provides an explanation, in paragraph form, of the policy in question, describing each rule and its role in the policy. Clicking the **Explain rules** link in the Groups, Host, Rule Modules, or Policy page, takes you to this paragraph explanation. See [Figure 5-11](#).

**Figure 5-11** Rule Module Explanation Page

The screenshot shows the Management Center for Cisco Security Agents V5.1 interface. The browser address bar shows the URL: <https://biohazard/csamc51/webadmin>. The page title is "Management Center for Cisco Security Agents V5.1". The navigation menu includes: Events, Systems, Configuration, Analysis, Maintenance, Reports, Search, Help. The breadcrumb trail is: Configuration > Rule Modules > Windows Rule Modules > Email Client Module - Medium Security [V5.1 r30] > Explanation. The main content area is titled "Explanation of rule module Email Client Module - Medium Security [V5.1 r30]". It contains the following text:

The detect rules [Monitor](#), [Add Process to Application Class](#), [Remove Process from Application Class](#), and [Set](#) are always evaluated **after** the enforce rules.

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

- the system state matches system state set [Security Level Medium \[V5.1 r30\]](#) but not system state sets [Cisco Trust Agent Quarantine Posture \[V5.1 r30\]](#), [Virus detected \[V5.1 r30\]](#), [Cisco Trust Agent Infected Posture \[V5.1 r30\]](#).

Application control

[COM\\_component\\_access\\_control](#)

[File\\_access\\_control](#)

[System\\_API\\_control](#)

**Control execution of applications**

**Monitor** In the absence of any applicable 'priority deny', 'priority terminate process' or 'allow' rules, Attempts to invoke processes in application class [All Applications](#) but not in application class [MS\\_Office\\_applications \[V5.1 r30\]](#), by processes in application class [Email\\_applications \[V5.1 r30\]](#) will be denied, unless overridden by the user. An event will be logged when the rule is triggered. [339](#)

**COM\_access\_control**

**Monitor** In the absence of any applicable 'priority deny', 'priority terminate process' or 'allow' rules, Attempts to access COM sets [MS\\_wscript\\_Objects \[V5.1 r30\]](#), [MS\\_vbscript\\_Objects \[V5.1 r30\]](#), [MS\\_File\\_System\\_Objects \[V5.1 r30\]](#) by processes in application class [Email\\_applications \[V5.1 r30\]](#) will be denied, unless overridden by the user. An event will be logged when the rule is triggered. [343](#)

**File\_access\_control**

**Monitor** In the absence of any applicable 'priority deny', 'priority terminate process' or 'allow' rules, Attempts to write files matching file sets [Email\\_files - Dangerous \[V5.1 r30\]](#), [Email\\_files - Suspicious \[V5.1 r30\]](#) by processes in application class [Email\\_applications \[V5.1 r30\]](#) will be denied, unless overridden by the user. An event will be logged when the rule is triggered. [342](#)

Attempts to read files matching file sets [MS\\_Script\\_Runtime \[V5.1 r30\]](#) by processes in application class [Email\\_applications \[V5.1 r30\]](#) will be denied, unless overridden by the user. An event will be logged when the rule is triggered. [344](#)

**System\_API\_control**

**Monitor** In the absence of any applicable 'priority deny', 'priority terminate process' or 'allow' rules, Attempts to

- invoke system functions from code executing in data or stack space

by processes in application class [Email\\_applications \[V5.1 r30\]](#) will cause the process to be terminated, unless overridden by the user. An event will be logged when the rule is triggered.

5 rule changes pending [Generate rules](#) Logged in as: admin

153460

## Consistency Check

The main rule module page provides an OS consistency check for variables that are part of the rule. For example, it makes sure that Linux applications classes are attached to a UNIX module that has Linux or All UNIX as its target OS. If the rule module has a target OS of Solaris, the consistency check will fail if the application class is marked as Linux.

This consistency check also ensures that modules of a specified OS type are attached to similar OS policies. You are allowed to save modules that do not pass the consistency check so that you can clone items or make multiple policy edits, but you are not allowed to attach and deploy inconsistent items.

# Attaching Rule Modules to Policies

**Note**

This is the same procedure provided in [Chapter 4, “Building Policies”](#). It is included here as well for your convenience.

When you configure a rule module, you are combining access control rules and/or tagging and monitoring rules under a common name. That rule module name is then attached to a policy. That policy uses the rules that comprise the module to control the actions that are allowed and denied on hosts. See [Configuring Rule Modules, page 5-5](#).

CSA MC gives you the option of attaching a rule module to a policy using the **Modify policy associations** link in the Rule Module configuration page or attaching a policy to a rule module 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 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 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. See [Figure 5-12](#). The left box contains the policies the rule module is not attached to. The right box 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 left box and click the **Add** button. The selected rule module moves to the right box 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 only need to maintain that one Apache policy.

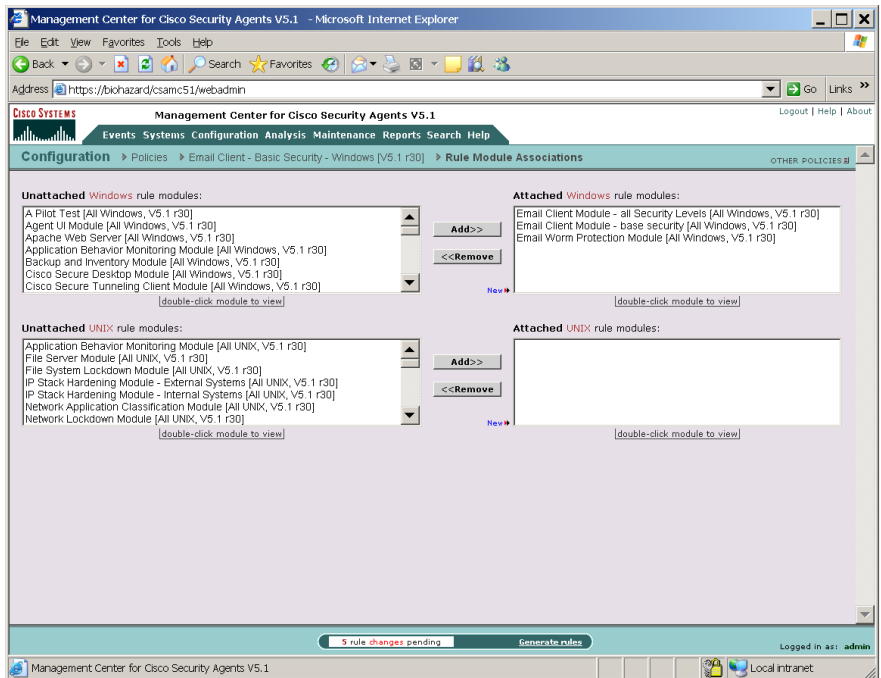
---

**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.

---

Figure 5-12 Rule Module Associations



## Generating Rule Programs



### Caution

When you make changes to existing CSA MC configurations, they are saved in the database, but they are not yet distributed to the agents across your network. You *must* click the **Generate rules** link in the bottom frame of CSA MC to first view all new and edited configurations and then distribute them to the agents. (When you have pending changes, the line beneath Generate rules link flashes.)

The Generate rule programs view displays the status of all non-distributed database items with the name of the administrator who made the configuration changes. A **Details** link appears beside each edited configuration item. Click this link to view what modifications were made to the configuration in question.

Once you've checked these modifications, you can either go back and change or delete configurations or you can click the **Generate** button (in the bottom frame) to distribute all updates.

**Note**

---

Before you generate rule programs and distribute them to agents, you can view all database changes, including the time the changes were made and the administrator who made them by accessing the **Audit Trail** view from the Reports drop-down list. See [page 2-15](#) for information.

---

Figure 5-13 Generate Configuration

Management Center for Cisco Security Agents V5.1 - Microsoft Internet Explorer

Address: <https://biohazard/csamc51/webadmin>

**Management Center for Cisco Security Agents V5.1**

Events Systems Configuration Analysis Maintenance Reports Search Help

### Generate Rule Programs

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

- [Application Behavior](#)
- [Email client - Linux](#)
- [Email Client - Multi-level Security - Windows](#)
- [General application - Multi-level Security - Linux](#)
- [General application - Multi-level Security - Windows](#)
- [Insecure boot - sample only](#)
- [Instant Messenger - Windows](#)
- [Network Quarantine](#)
- [New Policy](#)
- [Pilot Test](#)
- [Samba Server - Linux](#)
- [Security Classification](#)

5 changes since the last rule program generation:

#	Action	Time	Administrator
5	Create rule module <a href="#">Untitled_1 [W]</a>	2/6/2006 1:01:01 PM	admin
4	Add rule module 'Apache Web Server [W, V5.1 r30]' to policy <a href="#">New Policy</a>	2/6/2006 11:44:11 AM	admin
3	Modify policy <a href="#">New Policy</a> [Details]	2/6/2006 11:43:48 AM	admin
2	Modify policy <a href="#">New Policy</a> [Details]	2/6/2006 11:43:42 AM	admin
1	Create policy <a href="#">New Policy</a>	2/6/2006 11:28:44 AM	admin

Press the **Generate** button to create and distribute rule programs based on the current configuration.

5 rule changes pending

Logged in as: admin

Management Center for Cisco Security Agents V5.1

Local Intranet

153452

## Common Rule Page Configuration Items

The following sections provide information on the common fields found on most rule type pages. These include the action options that determine rule precedence as well as a description of how query rules work. The unique rule types themselves are described in the next chapter.

## Rules: Action Options and Precedence

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	Priority Terminate Process
Priority 2	Priority Deny
Priority 3	Priority 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/Not applicable	Monitor
Priority/Not applicable	Set
Priority/Not applicable	Add process to application class
Priority/Not applicable	Remove process from application class

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



### Note

The default action of the agent is to allow an operation (priority 9 in the previous table) in the absence of any applicable rule. An exception to this occurs when attempts are made to modify the Cisco Security Agent resources. Due to agent self-protection, these requests are denied by default.

## Rules: Action Definitions

When you configure your access control rules, you must select an action for that rule. The following is a description of all possible action types. You should note that not all action types are available for all rules.

### *Enforcement Actions*

- **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. Note that all processes cannot be safely terminated (e.g. winlogon). If it is not safe to terminate the process, the action will be denied but not terminated.
- **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.
- **Priority Allow**—Select this action type to create a rule that allows the action you specify to take place. Because the default action of all rules is "allow", generally, you'll 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 you indicate occurs. The user then has 5 minutes to answer Yes, No, or Terminate. By default, it will be denied and the process will be terminated unless the user decides otherwise. See Query User for more details. (Query User options are not available for Solaris rules.)
- **Query User (Default Deny)**—Select this action type to prompt the user when the action you indicate occurs. The user then has 5 minutes to answer Yes, No, or Terminate. By default, it will be denied unless the user decides otherwise. See Query User for more details. (Query User options are not available for Solaris rules.)

- **Query User (Default Allow)**—Select this action type to prompt the user when the action you indicate occurs. The user then has 5 minutes to answer Yes, No, or Terminate. By default, it will be allowed unless the user decides otherwise. See Query User for more details. (Query User options are not available for Solaris rules.)

**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 pop-up 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 that all processes cannot be safely terminated (e.g. winlogon). If it is not safe to terminate the process, the action will be denied but the process will not be terminated.
- **Deny**—Select this action type to create a rule that stops the action you specify from occurring on systems. (When you select Deny for this rule, if the user attempts to run the application in question, he/she is notified with a pop-up box explaining that the application is forbidden to run.)

#### *Detection Actions*

Note that Detection rules are always checked, even in the presence of a priority enforcement rule which governs the same resources triggering first.



#### **Note**

---

All rules are evaluated before any dynamic application classifications are applied to processes. This ensures that application class memberships are consistently applied for all rules being evaluated for a given request.

---

- **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 if the resource access action is an allow or a deny. Additionally, you can configure

monitor rules to trigger only when an enforcement action of a certain type occurs. This way, you are only monitoring “deny actions on a file resource”, for example. See [The Monitor Action, page 5-42](#) for more information,

- **Set**—Use the "Set" action in a rule to cause a particular configuration action to occur when the criteria configured in the rule occurs on a system. See [Using the Set Action, page 5-43](#).
- **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 by a specific application executable name. A process will be added to a dynamic class if the action specifying access to the resource matches any of the parameters in this rule (i.e., allow, deny, terminate). See [Building Classes as Rule Consequences, page 8-12](#) for details.
- **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 if the action specifying access to the resource matches any of the parameters in this rule (i.e., allow, deny, terminate). See [Building Classes as Rule Consequences, page 8-12](#) for details.




---

**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 dynamic application classifications are ephemeral and are constantly being re-evaluated and classified on the system.

---

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, it is unlikely that you would write allow rules unless they are to make exceptions to deny rules you are writing within a module or for monitoring purposes (see [The Monitor Action, page 5-42](#)). If you do write a stand-alone allow rule, because the default action is allow, the allow rule itself is then essentially irrelevant.

A good model for configuring rules within modules would be to take the priority levels into account and work from the bottom up, lowest priority to highest priority. Before you even add a single parameter to a rule, by default, 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 allowing the user to decide if an action should be allowed or denied. Lastly write any high priority rules you might need.

## Rules: Manipulating Precedence

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

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

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

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

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

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

**Caution**

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

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

**Note**

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

**Note**

Test Mode and Learn Mode do not affect rule precedence. For example, if you have two rules that are exactly the same, except one is in Test Mode and one is not, the manner in which CSA MC orders those rules in the list dictates which rule will fire first. That order can be as simple as - the rule that was created by the administrator first fires first.

## The Monitor Action

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 whether the resource access action is an allow or a deny.

Additionally, you can configure monitor rules to trigger only when an enforcement action of a certain type occurs. This way, you are only monitoring “deny actions on a file resource”, for example.

## Using the Set Action

Set is a singular configuration action that causes a particular, one-time, configuration item to occur when the criteria configured in the rule triggers on a system. For example, when a rule with “Set” configured triggers, a specific action occurs, such as the security level being set to low. This is different from Add and Remove process tagging. Add and Remove process tagging cause a tag to be bound to a process for the life of that process or until the tag is removed. Set causes a one-time action to occur.



---

**Note**

Set is similar to Add and Remove process tagging in one respect. You can configure the Set to occur based on a user query response. However, the only valid query response for Set is "Allow". If the connection were to be denied, there would be nothing to mark.

---



---

**Note**

In some cases the resulting Set action may configure a global system state (Set Host Address as Untrusted). In other cases, the output of the Set action is to simply generate an event (such as Set Rule Module protection.)

---

When you select Set as the action type for a rule, an Attribute pulldown menu and a Value pulldown menu appear. For every attribute you can set, there are corresponding values that must also be set. There are several attribute and value pairs you can set for a rule.



---

**Note**

Not all attributes are available for every rule type. For example, Differentiated Service options are only available for Network access control rules. The applicable set attributes for the rule type in question will be the only types you can select for that rule.

---

Possible attribute and value pairs are:

## Attribute: detected access

**Values:** Protected, Unprotected

This attribute is intended to notify you (via the event log) and optionally take action (via a system state) when an application or service, or other system component that is marked as Unprotected does not have a corresponding Protected rule and is therefore not being protected by the agent.

To use this as a protection auditing tool, you would include a rule that is configured with “Set-detected access-Protected” for the resource that you are protecting in a policy. This rule marks the resource as being protected by the agent.

Subsequently, if you want to make sure that certain resources are being protected on systems, you would configure a “Set-detected access-Unprotected” rule for the resource in question and include it, for example, in a policy attached to the <All OS type> groups. This way, if a resource that requires protection is accessed on a system and there a “Set-detected access-Unprotected” rule without a corresponding “Set-detected access-Protected” rule for that resource, a message is sent to the event log informing you that the resource is not protected.

For example, you could use this feature to ensure that hosts running Web servers are properly protected. To accomplish this, you would write a set rule in the All Windows Group that said “Set-detected access-Unprotected” when any application acts as a server for TCP/80. Another rule would be added to the Microsoft IIS Web Server module specifying to “Set-detected access-Protected” when IIS accepts a connection for TCP/80. (A similar rule could be added for Apache in the corresponding module.) Hosts running an IIS Web Server that have both policies attached would meet the protection criteria and no event would be logged. Hosts with policies that specify the requirement for protection (Unprotected) with no corresponding protection provided (Protected) would have a policy mismatch and would generate an event as a result.

You can configure a System State to apply if a corresponding Set rule of this type triggers. See [System State Sets, page 5-13](#).

## Attribute: detected boot

**Values:** Insecure, Secure

This attribute is intended to detect when a previous system boot occurred in a non-standard manner. For example, the system was booted from a peripheral device (CD ROM) rather than from the hard drive. This type of boot can be considered non-standard and therefore possibly suspicious. (This is one way of introducing a Trojan to a system.) This type of peripheral device insecure boot detection works in conjunction with a particular type of compatible BIOS on compliant systems. The compatible BIOS detects a non-standard boot and on the next normal boot, if you have an appropriately configured Kernel Protection rule (see [Kernel Protection, page 6-54](#)), a message is sent to the MC which logs this insecure boot detection. This, in turn, causes the system state (if configured) to trigger. A Safe Mode boot also falls into this insecure boot category. Compatible BIOS is not required for a Safe Mode boot detection.

You can configure a System State to apply if a corresponding Set rule of this type triggers. See [System State Sets, page 5-13](#).

## Attribute: detected rootkit

**Values:** Trusted, Untrusted

A rootkit may be detected when module loads after boot time or a module attempts to modify kernel functionality. Note that if a rootkit gets marked as both trusted and untrusted, a trusted rootkit gets precedence over an untrusted rootkit tag.

You can configure a System State to apply if a corresponding Set rule of this type triggers. See [System State Sets, page 5-13](#).

## Attribute: Differentiated Service (Trusted QoS)

**Values:** priority Best Effort (0,0), priority Scavenger (8, CS1), application specified, IP routing (48, CS6), Voice (46, EF), Interactive Video (34, AF41), Streaming Video (32, CS4), Mission Critical Data (26, AF31), Call Signaling (24, CS3), Transactional Data (18, AF21), Network Management (16, CS2), Bulk Data (10, AF11), Best Effort (0,0). Scavenger (8, CS1)

You specify Differentiated Service for a certain traffic flow by setting a QoS marking which is a recognizable value in an IP packet. This allows routers and switches to identify and take action on QoS-marked traffic, providing finer granularity of control in forwarding traffic.

The available markings provide a DSCP (Differentiated Services Code Point) setting and a PHB (Per Hop Behavior) setting. The DSCP value matches the service field in the IP header. The PHB value matches the way routers and switches handle traffic on a hop by hop basis. These values appear in parenthesis beside each value option in the following order (DSCP,PHB).

**Note**

While all provided value markings are not described here, you should note that by default, applications transfer with a Best Effort value. The provided Scavenger value is lower than Best Effort. You would use a Scavenger marking to significantly downgrade a certain type of traffic flow. The **priority** Best Effort and **priority** Scavenger values are provided to allow you set a prioritization between multiple rules that may be using multiple types of markings.

**Important Details about CSA Differentiated Service Functionality**

In the absence of any rules on the agent that provision QoS markings, the general default of all systems is to mark traffic flows as “application specified”. In the presence of rules on the agent that provision QoS markings, the agent will select and provision DSCP markings accordingly. If there is a Differentiated Service rule conflict on the agent (i.e. more than one applicable Differentiated Service set rule for a traffic flow), the agent picks the highest marking to provision. (The precedence order of markings is the available pulldown list, top to bottom.)

Note that the agent is only marking packets that it transmits. The agent cannot mark packets that it receives. It is the responsibility of the remote host to mark those packets appropriately.

When the agent is disabled, stopped, or turned to Off using the security slide bar, existing sessions are no longer QoS provisioned by the agent. The agent stops marking existing flows. Existing flows then revert to “application specified”. If the agent is re-enabled, authorized flows resume the previous marking provisioned by the agent. If a new session is created during the time the agent is disabled, the new flow is not interfered with when the agent is back in the picture. That flow is automatically given the general system default of allow with an application specified marking and it retains that marking for the life of the connection.

Note that we only authorize at the beginning of a connection.

Note that Test Mode has no bearing on Differentiated Service functionality. The agent provisions the specified markings even if Test Mode is enabled.

Refer to RFC 2475 for general information on Differentiated Services.

You can configure a System State to apply if a corresponding Set rule of this type triggers. See [System State Sets, page 5-13](#). For example, if your network is under attack (i.e. "virus detected") you can configure a system state to trigger that uses rules to downgrade all traffic flows.

## Attribute: file

**Values:** Trusted, Untrusted

A file may be marked as Untrusted Content if it was potentially downloaded from the network or from some other untrusted source. Executable files that are persistently tracked in the agent UI Untrusted Applications window (see [The Agent User Interface, page A-8](#)) are localized to each system. Applications that appear in that localized list are automatically added to the built-in application class <\*Processes Executing Untrusted Content> (see [Built-in Application Classes, page 8-4](#)) and, in turn, that populated built-in application class can be used in rules.

Note that if a file gets marked as both trusted and untrusted, a trusted file tag gets precedence over an untrusted file tag.

## Attribute: Host address

**Values:** Untrusted host (locally and globally), Untrusted host (locally)

This Host address attribute is intended to mark the IP addresses of hosts as untrusted when they violate security policies or exhibit malicious behavior. Being classified as an untrusted host (locally) causes that host to be temporarily added (for one hour) to the @dynamic list on the local machine. Additionally, if the untrusted global attribute is used, an event is sent to CSA MC to make the host address a candidate for global event correlation. If this address is centrally correlated, it's permanently added to the global quarantine list on all hosts (if the configured event correlation threshold is reached).



### Tip

You may want to use the local untrusted value for an external web server that is continually hit by external hosts. This way, those hosts that appear malicious wouldn't become globally quarantined for your entire internal network, but they would be temporarily prevented from communicating with the web server.

You can configure a System State to apply if a corresponding Set rule of this type triggers. See [System State Sets, page 5-13](#).

## Attribute: Security level

**Values:** High, Medium, Low

Set the Security Level attribute to programmatically change the agent security level based on the current running state of the system. For example, if the agent security level is low and a virus is detected on the system, this can trigger a system state policy that will automatically be applied when the state has been moved to high. On a high setting, you may enforce a rule that denies the virus-infected system from making outgoing network connections.

You can configure a System State to apply if a corresponding Set rule of this type triggers. See [System State Sets, page 5-13](#).

## Querying the User

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

Query configurations are a Variable setting which allows you to decide which radio button options are displayed in the pop-up query box, which action is the default, whether the answer given by the user is to be remembered, and what the query text to be displayed will be.

For a Query setting, the response to the query is relevant to the question, not 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.

See [Query Settings, page 9-25](#) for configuration details.

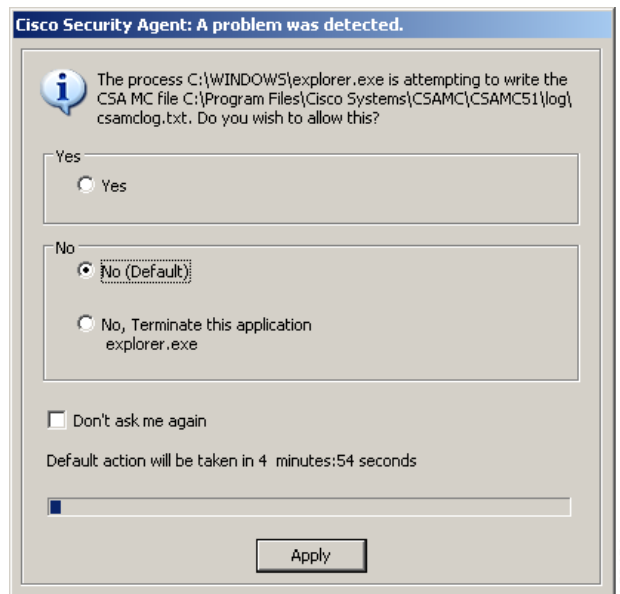
**Caution**

For Solaris rules, Query user actions are selectable on the MC but query pop-ups are not displayed on Solaris agent host systems. Instead, if you configure a Query user rule for a Solaris system, the default action is immediately taken on the system if the rule is triggered.

For Windows and Linux agents, agent settings (including user queries) are configurable by the administrator. If the agent UI is hidden for the group, there are no query user pop-up boxes displayed. The default is immediately taken on all query user rules and heuristics that are present in the assigned policies.

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

**Figure 5-14** Query User Pop-up Box



From the Query Settings page, accessible from the **Configuration>Variables** menu, you configure the query text and the query radio buttons that appear in the pop-up box the end user will see. In the Query pop-up box, the user reads the information given on the attempted action and selects one of the following possible choices and clicks Apply:

- **Yes**—Allows the application access to the resource in question.
- **No**—Denies the application access to the resource in question.
- **No, Terminate this application**—Denies the application access to the resource in question and also attempts to terminate the application process. The name of the application in question is displayed with the terminate option. (Some processes cannot be safely terminated, such as winlogon.)

**Default Action**—You chose one of the radio buttons displayed in the query pop-up 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 me again**—In addition to the buttons that will appear on the query box, you can decide to also display a **Don't ask me again** checkbox so that the user's query response is remembered. If the user selects that checkbox when he/she responds to the query, and the same query is triggered, the remembered response is automatically taken and the user is not queried again.

**Query challenge**—For added security, you can issue a query challenge on the query pop-up 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 pop-up box itself. (See [Query Settings, page 9-25.](#))



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

**Caution**

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

## Caching Responses

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 pop-up query box appearing again either permanently or for some period of time.

For example, if a user is queried as to whether an application can talk on the network and the user responds by selecting the **Yes** radio button and clicking a **Don't ask again** checkbox, the Yes response is remembered permanently and that response appears in the edit field in the agent UI query response window. But if the user is queried as whether setup.exe can install software on the system and the user responds by selecting the **Yes** radio button, but there is no **Don't ask again checkbox** or it is there but the user does not select it, this response is remembered temporarily and it does not appear in the agent UI query response window.

If the user response is only cached temporarily (for approximately an hour) the user can click the **Clear** button in this 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 press the Delete key.

### Notes about Query Caching

- Permanent responses are remembered across reboots.
- Temporarily cached responses are not remembered across reboots.
- When a query response is cached temporarily for one hour, if during that one hour time frame, the cached response is used by a subsequent rule request, the window is extended by an hour.
- A query response is tied to the user who responded. On multi-user machines, multiple users may be asked the same question.

## Query Rule Priority Information

You should note how CSA MC manages rule priority if there are multiple similar query rules which need to be evaluated.

Base Priority: Action=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