Intrusion Policies

The following topics explain intrusion policies and the closely associated network analysis policies (NAP). Intrusion policies include rules that check traffic for threats and block traffic that appears to be an attack. Network analysis policies control traffic preprocessing, which prepares traffic to be further inspected by normalizing traffic and identifying protocol anomalies.

Because preprocessing and intrusion inspection are so closely related, the network analysis and intrusion policies examining a single packet must complement each other.

About Intrusion and Network Analysis Policies

Network analysis and intrusion policies work together to detect and prevent intrusion threats.

- A network analysis policy (NAP) governs how traffic is decoded and preprocessed so that it can be further evaluated, especially for anomalous traffic that might signal an intrusion attempt.
- An intrusion policy uses intrusion and preprocessor rules, which are collectively known as intrusion rules, to examine the decoded packets for attacks based on patterns. The rules can either prevent (drop) the threatening traffic and generate an event, or simply detect (alert) it and generate an event only.

As the system analyzes traffic, the network analysis decoding and preprocessing phase occurs before and separately from the intrusion prevention phase. Together, network analysis and intrusion policies provide broad and deep packet inspection. They can help you detect, alert on, and protect against network traffic that could threaten the availability, integrity, and confidentiality of hosts and their data.

System-Defined Network Analysis and Intrusion Policies

The system includes several pairs of same-named network analysis and intrusion policies that complement and work with each other. For example there are both NAP and intrusion policies named “Balanced Security and Connectivity,” which are meant to be used together. The system-provided policies are configured by the

Cisco Talos Intelligence Group (Talos). For these policies, Talos sets the intrusion and preprocessor rule states and provides the initial configurations for preprocessors and other advanced settings.

As new vulnerabilities become known, Talos releases intrusion rule updates. These rule updates can modify any system-provided network analysis or intrusion policy, and can provide new and updated intrusion rules and preprocessor rules, modified states for existing rules, and modified default policy settings. Rule updates might also delete rules from system-provided policies and provide new rule categories, as well as modify the default variable set.

You can manually update the rules database, or configure a regular update schedule. You must deploy an update for it to take effect. For more information on updating system databases, see Updating System Databases.

The following are the system-provided policies:

Balanced Security and Connectivity network analysis and intrusion policies

These policies are built for both speed and detection. Used together, they serve as a good starting point for most networks and deployment types. The system uses the Balanced Security and Connectivity network analysis policy as the default.

Connectivity Over Security network analysis and intrusion policies

These policies are built for networks where connectivity, the ability to get to all resources, takes precedence over network infrastructure security. The intrusion policy enables far fewer rules than those enabled in the Security over Connectivity policy. Only the most critical rules that block traffic are enabled.

Security Over Connectivity network analysis and intrusion policies

These policies are built for networks where network infrastructure security takes precedence over user convenience. The intrusion policy enables numerous network anomaly intrusion rules that could alert on or drop legitimate traffic.

Maximum Detection network analysis and intrusion policies

These policies are built for networks where network infrastructure security is given even more emphasis than is given by the Security Over Connectivity policies, with the potential for even greater operational impact. For example, the intrusion policy enables rules in a large number of threat categories including malware, exploit kit, old and common vulnerabilities, and known in-the-wild exploits.

Intrusion and Preprocessor Rules

An intrusion rule is a specified set of keywords and arguments that the system uses to detect attempts to exploit vulnerabilities in your network. As the system analyzes network traffic, it compares packets against the conditions specified in each rule, and triggers the rule if the data packet meets all the conditions specified in the rule.

The system includes the following types of rules created by Cisco Talos Intelligence Group (Talos):

- Intrusion rules, which are subdivided into shared object rules and standard text rules
- Preprocessor rules, which are rules associated with preprocessors and packet decoder detection options in the network analysis policy. Most preprocessor rules are disabled by default.

The following topics explain intrusion rules in more depth.

Intrusion Rule Attributes

When you view an intrusion policy, you see a list of all the intrusion rules available for identifying threats.

The list of rules for each policy show only those rules set to alert or drop, and those rules you explicitly disabled. Rules that are disabled by default are not shown. Although there are over 30,000 rules, you will see only a subset of all possible rules. But even for the smallest enabled rule set, scrolling through the list will take time. Rules are revealed as you scroll.

Following are the attributes that define each rule:

> (Signature Description)

Click the > button in the left column to open the signature description. The description is the actual code used by the Snort inspection engine to match traffic against the rule. Explaining the code is out of scope for this document, but it is explained in detail in *Management Center Configuration Guide*; select the book for your software version from http://www.cisco.com/c/en/us/support/security/defense-center/products-installation-and-configuration-guides-list.html. Look for information on intrusion rule editing.

The signatures contain variables for certain items. For more information, see Default Intrusion Variable Set, on page 3.

GID

Generator Identifier (ID). This number indicates which system component evaluates the rule and generates events. A 1 indicates a standard text intrusion rule, a 3 indicates a shared object intrusion rule. (The difference in these rule types is not meaningful for a FDM user.) These are the main rules of interest when configuring an intrusion policy. For information on the other GIDs, see Generator Identifiers, on page 4.

SID

Snort Identifier (ID), also called signature ID. Snort IDs lower than 1000000 were created by the Cisco Talos Intelligence Group (Talos).

Action

The state of this rule in the selected intrusion policy. For each rule, "(Default)" is added to the action that is the default action for the rule within this policy. To return a rule to its default setting, you select this action. Possible actions are:

- **Alert**—Create an event when this rule matches traffic, but do not drop the connection.
- **Drop**—Create an event when this rule matches traffic, and also drop the connection.
- **Disabled**—Do not match traffic against this rule. No events are generated.

Status

If you change the default action for a rule, this column shows "Overridden." Otherwise, the column is empty.

Message

This is the name of the rule, which also appears in events triggered by the rule. The message typically identifies the threat that the signature matches. You can search the Internet for more information on each threat.

Default Intrusion Variable Set

The intrusion rule signatures contain variables for certain items. Following are the default values for the variables, with $HOME_NET and $EXTERNAL_NET being the most commonly used variables. Note that the protocol is specified separately from port numbers, so port variables are numbers only.

- $DNS_SERVERS = $HOME_NET (meaning any IP address).
- $EXTERNAL_NET = any IP address.
- $FILE_DATA_PORTS = $HTTP_PORTS, 143, 110.
- $FTP_PORTS = 21, 2100, 3535.
- $GTP_PORTS = 3386, 2123, 2152.
- $HOME_NET = any IP address.
- $HTTP_PORTS = 144 ports numbered: 36, 80-90, 311, 383, 443, 555, 591, 593, 631, 666, 801, 808, 818, 901, 972, 1158, 1212, 1220, 1414, 1422, 1533, 1741, 1830, 1942, 2231, 2301, 2381, 2578, 2809, 2980, 3029, 3037, 3057, 3128, 3443, 3507, 3702, 4000, 4343, 4848, 5000, 5117, 5222, 5250, 5450, 5600, 5814, 6080, 6173, 6767, 6988, 7000, 7001, 7005, 7071, 7080, 7144, 7145, 7510, 7770, 7777-7779, 8000, 8001, 8008, 8014, 8015, 8020, 8028, 8040, 8060, 8080-8082, 8085, 8088, 8118, 8123, 8161, 8180-8182, 8222, 8243, 8280, 8300, 8333, 8344, 8400, 8443, 8500, 8509, 8787, 8800, 8888, 8899, 8983, 9000, 9002, 9060, 9080, 9090, 9091, 9111, 9290, 9443, 9447, 9710, 9788, 9999, 10000, 11371, 12601, 13014, 15489, 19980, 23472, 29991, 33300, 34412, 34443, 34444, 40007, 41080, 44449, 50000, 50002, 51423, 53331, 55252, 55555, 56712.
- $HTTP_SERVERS = $HOME_NET (meaning any IP address).
- $ORACLE_PORTS = any
- $SHELLCODE_PORTS = 180.
- $SIP_PORTS = 5060, 5061, 5600
- $SIP_SERVERS = $HOME_NET (meaning any IP address).
- $SMTP_SERVERS = $HOME_NET (meaning any IP address).
- $SNMP_SERVERS = $HOME_NET (meaning any IP address).
- $SQL_SERVERS = $HOME_NET (meaning any IP address).
- $SSH_PORTS = 22.
- $SSH_SERVERS = $HOME_NET (meaning any IP address).
- $TELNET_SERVERS = $HOME_NET (meaning any IP address).

Generator Identifiers

The generator identifier (GID) identifies the subsystem that evaluates an intrusion rule and generates events. Standard text intrusion rules have a generator ID of 1, and shared object intrusion rules have a generator ID of 3. There are also several sets of rules for various preprocessors. The following table explains the GIDs.

Table 1: Generator IDs

ID	Component
1	Standard Text Rule.

ID	Component
2	Tagged Packets. (Rules for the Tag generator, which generates packets from a tagged session.)
3	Shared Object Rule.
102	HTTP Decoder.
105	Back Orifice Detector.
106	RPC Decoder.
116	Packet Decoder.
119, 120	HTTP Inspect Preprocessor. (GID 120 rules relate to server-specific HTTP traffic.)
122	Portscan Detector.
123	IP Defragmentor.
124	SMTP Decoder. (Exploits against SMTP verbs.)
125	FTP Decoder.
126	Telnet Decoder.
128	SSH Preprocessor.
129	Stream Preprocessor.
131	DNS Preprocessor.
133	DCE/RPC Preprocessor.
134	Rule Latency, Packet Latency. (Events for these rules are generated when rule latency suspends (SID 1) or re-enables (SID 2) a group of intrusion rules, or when the system stops inspecting a packet because the packet latency threshold is exceeded (SID 3).)
135	Rate-Based Attack Detector. (Excessive connections to hosts on the network.)
137	SSL Preprocessor.
138, 139	Sensitive Data Preprocessor.
140	SIP Preprocessor.
141	IMAP Preprocessor.

ID	Component
142	POP Preprocessor.
143	GTP Preprocessor.
144	Modbus Preprocessor.
145	DNP3 Preprocessor.

Network Analysis Policies

Network analysis policies control traffic preprocessing. Preprocessors prepare traffic to be further inspected by normalizing traffic and identifying protocol anomalies. Network analysis-related preprocessing occurs after Security Intelligence drops and SSL decryption, but before access control and intrusion or file inspection.

By default, the system uses the Balanced Security and Connectivity network analysis policy to preprocess all traffic handled by the access control policy. However, if you configure an intrusion policy on any access control rule, the system uses the network analysis policy that matches the most aggressive intrusion policy applied. For example, if you use both Security over Connectivity and Balanced policies in your access control rules, the system uses the Security over Connectivity NAP for all traffic.

License Requirements for Intrusion Policies

You must enable the **Threat** license to apply intrusion policies in access control rules. For information on configuring licenses, see Enabling or Disabling Optional Licenses.

No extra license is needed for network analysis policies.

Applying Intrusion Policies in Access Control Rules

To apply intrusion policies to network traffic, you select the policy within an access control rule that allows traffic. You do not directly assign intrusion policies.

You can assign different intrusion policies to provide variable intrusion protection based on the relative risks of the networks you are protecting. For example, you might use the more stringent Security over Connectivity policy for traffic between your inside network and external networks. On the other hand, you might apply the more lenient Connectivity over Security policy for traffic between inside networks.

You can also simplify your configuration by using the same policy for all networks. For example, the Balanced Security and Connectivity policy is design to provide good protection without excessively impacting connectivity.

Procedure

Step 1 Select **Policies** > **Access Control**.

Step 2 Either create a new rule, or edit an existing rule, that **allows** traffic.

If the default action is allow, you can also specify an intrusion policy in the default action.

You cannot apply intrusion policies to rules that trust or block traffic.

Step 3 Click the **Intrusion Policy** tab.

Step 4 Select **Intrusion Policy** > **On** and select the intrusion inspection policy to use on matching traffic.

Configuring Syslog for Intrusion Events

You can configure an external syslog server for intrusion policies to send intrusion events to your syslog server. You must configure the syslog server on the intrusion policy to get intrusion events sent to the server. Configuring a syslog server on an access rule sends connection events only to the syslog server, not intrusion events.

Intrusion events have the message ID 430001.

Procedure

Step 1 Select **Policies** > **Intrusion**.

Step 2 Click the **Edit Logging Settings** button ( ) to configure syslog.

Step 3 Click in the **Send Connection Events To** field and select the server object that defines the syslog server. If the required object does not already exist, click **Create New Syslog Server** and create it.

Step 4 Click **OK**.

Managing Intrusion Policies

You can apply any of the pre-defined intrusion policies. Each of these policies includes the same list of intrusion rules (also known as signatures), but they differ in the actions taken for each rule. For example, a rule might be active in one policy, but disabled in another policy.

If you find that a particular rule is giving you too many false positives, where the rule is blocking traffic that you do not want blocked, you can disable the rule without needing to switch to a less-secure intrusion policy. You could alternatively change it to alert on matches without dropping traffic.

However, if a rule is disabled by default in the intrusion policy, you cannot change it to drop or alert on matching traffic. You can change the action only on enabled policies or on policies that you previously disabled.

Use the intrusion related dashboards, and the Event Viewer (both on the **Monitoring** page), to evaluate how intrusion rules are impacting traffic. Keep in mind that you will see intrusion events and intrusion data only for traffic that matches intrusion rules set to alert or drop; disabled rules are not evaluated.

The following topics explain intrusion policies and rule tuning in more detail.

Changing Intrusion Rule Actions

Each pre-defined intrusion policy has the same rules. The difference is the action taken for each rule can be different from policy to policy.

Within a given policy, you can change the default action for a rule only if it is enabled, that is, set to alert or drop. By changing the default action, you can disable rules that are giving you too many false positives, or you can change whether the rule alerts on or drops matching traffic.

Note If you change an action from the default, the next time the intrusion rules database is updated, the system resets the rule default to the action you selected. At that point, your selection becomes the new default, and the status no longer shows the action as Overridden.

Procedure

Step 1 Select **Policies** > **Intrusion**.

Step 2 Click the tab for the intrusion policy whose rule actions you want to change.

The pre-defined policies are:

- Connectivity over Security
- Balanced Security and Connectivity
- Security over Connectivity
- Maximum Detection

Step 3 Find the rule whose action you want to change.

The rules are sorted with the overridden ones listed first, and sorted by action within the group of overridden rules. Otherwise, the rules are sorted by GID, then SID.

The list of rules for each policy show only those rules set to alert or drop, and those rules you explicitly disabled. Rules that are disabled by default are not shown.

Use the search box to locate the rule you want to change. Ideally, you can get the Snort identifier (SID) and generator identifier (GID) from an event or from Cisco Technical Support, if you are working with them on an issue.

For detailed information about the elements of each rule, see Intrusion Rule Attributes, on page 2.

To search the list:

a) Click in the **Search** box to open the search attributes dialog box.
b) Enter a combination of Generator ID (**GID**), Snort ID (**SID**), or rule **Action**, and click **Search**.

For example, you could select **Action = Drop** to see all the rules in the policy that will drop matching connections. The text beside the search box indicates how many rules match your criteria, for example, “8937 of 9416 rules found.”

To clear a search criteria, click the x for the criteria in the search box.

Step 4 Click the **Action** column for the rule and select the required action:

- **Alert**—Create an event when this rule matches traffic, but do not drop the connection.
- **Drop**—Create an event when this rule matches traffic, and also drop the connection.

- **Disabled**—Do not match traffic against this rule. No events are generated.

The default action for the rule is indicated by "(Default)" added to the action. If you change the default, the status column indicates "Overridden" for that rule.

Monitoring Intrusion Policies

You can find intrusion policy statistics on the **Attackers** and **Targets** dashboards on the **Monitoring** page. You must apply an intrusion policy to at least one access control rule to see any information on these dashboards. See Monitoring Traffic and System Dashboards.

To see intrusion events, select **Monitoring** > **Events**, then click the **Intrusion** tab. You can hover over an event and click the **View Details** link to get more information. From the details page, you can click the **View IPS Rule** to go to the rule in the relevant intrusion policy, where you can change the rule action. This can help you reduce the impact of false positives, where a rule is blocking too many good connections, by changing the action from drop to alert. Conversely, you can change an alert rule into a drop rule if you are seeing a lot of attack traffic for a rule.

If you configure a syslog server for the intrusion policy, intrusion events have the message ID 430001.

Examples for Intrusion Policies

The use case chapter includes the following examples of implementing intrusion policies.

- How to Block Threats
- How to Passively Monitor the Traffic on a Network