Understanding Network Analysis and Intrusion Policies

Network analysis and intrusion policies work together as part of the ASA FirePOWER module intrusion detection and prevention feature. The term *intrusion detection* generally refers to the process of passively analyzing network traffic for potential intrusions and storing attack data for security analysis. The term *intrusion prevention* includes the concept of intrusion detection, but adds the ability to block or alter malicious traffic as it travels across your network.

In an intrusion prevention deployment, when the system examines packets:

- A **network analysis policy** 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* (sometimes referred to collectively as *intrusion rules*) to examine the decoded packets for attacks based on patterns. Intrusion policies are paired with *variable sets*, which allow you to use named values to accurately reflect your network environment.

Both network analysis and intrusion policies are invoked by a parent access control policy, but at different times. As the system analyzes traffic, the network analysis (decoding and preprocessing) phase occurs before and separately from the intrusion prevention (additional preprocessing and intrusion rules) 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.

The ASA FirePOWER moduleis delivered with several similarly named network analysis and intrusion policies (for example, Balanced Security and Connectivity) that complement and work with each other. By using system-provided policies, you can take advantage of the experience of the Cisco Vulnerability Research Team (VRT). For these policies, the VRT sets intrusion and preprocessor rule states, as well as provides the initial configurations for preprocessors and other advanced settings.

You can also create custom network analysis and intrusion policies. You can tune settings in custom policies to inspect traffic in the way that matters most to you.

You create, edit, save, and manage network analysis and intrusion policies using similar policy editors. When you are editing either type of policy, a navigation panel appears on the left side of the user interface; the right side displays various configuration pages.

This chapter contains a brief overview of the types of configurations the network analysis and intrusion policies govern, explains how the policies work together to examine traffic and generate records of policy violations, and provides basic information on navigating the policy editors. This chapter also explains the benefits and limitations of using custom versus system-provided policies. For more information, see the following sections:

- Understanding How Policies Examine Traffic For Intrusions, page 18-2
- Comparing System-Provided with Custom Policies, page 18-7
- Using the Navigation Panel, page 18-13
- Resolving Conflicts and Committing Policy Changes, page 18-15

To customize your intrusion deployment, see the following for your next steps:

- Working with Variable Sets, page 2-13 explains how to configure the system's intrusion variables to accurately reflect your network environment. Even if you do not use custom policies, Cisco **strongly** recommends that you modify the default variables in the default variable set. Advanced users can create and use custom variable sets for pairing with one or more custom intrusion policies.

- Getting Started with Intrusion Policies, page 26-1 explains how to create and edit a simple custom intrusion policy.

- Controlling Traffic Using Intrusion and File Policies, page 11-1 explains how to configure the system to use intrusion policies to examine only the traffic you are interested in by associating intrusion policies with a parent access control policy. It also explains how to configure advanced intrusion policy performance options.

- Configuring Advanced Transport/Network Settings, page 24-1 explains how to configure advanced transport and network preprocessor settings that apply globally to all traffic. You configure these advanced settings in an access control policy rather than in a network analysis or intrusion policy.

- Getting Started with Network Analysis Policies, page 21-1 explains how to create and edit a simple custom network analysis policy.

- Customizing Preprocessing with Network Analysis Policies, page 20-2 explains how to change the default network analysis policy. For advanced users, this section also explains how to tailor preprocessing to specific security zones and networks by assigning custom network analysis policies to preprocess matching traffic.

- Using Layers in a Network Analysis or Intrusion Policy, page 19-1 explain how, in larger organizations or complex deployments, you can use building blocks called policy *layers* to more efficiently manage multiple network analysis or intrusion policies.

Understanding How Policies Examine Traffic For Intrusions

License: Protection

When the system analyzes traffic as part of your access control deployment, the network analysis (decoding and preprocessing) phase occurs before and separately from the intrusion prevention (intrusion rules and advanced settings) phase.

The following diagram shows, in a simplified fashion, the order of traffic analysis in an inline, intrusion prevention and advanced malware protection (AMP) deployment. It illustrates how the access control policy invokes other policies to examine traffic, and in which order those policies are invoked. The network analysis and intrusion policy selection phases are highlighted.
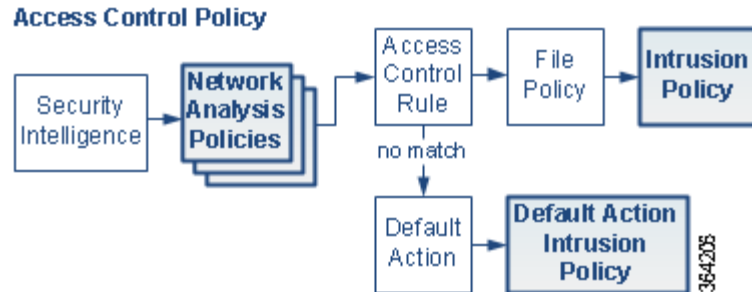
In an inline deployment, the system can block traffic without further inspection at almost any step in the illustrated process. Security Intelligence, the SSL policy, network analysis policies, file policies, and intrusion policies can all either drop or modify traffic.

Similarly, at each step of the process, a packet could cause the system to generate an event. Intrusion and preprocessor events (sometimes referred to collectively as *intrusion events*) are indications that a packet or its contents may represent a security risk.

Note that for a single connection, although the system selects a network analysis policy before an access control rule as shown in the diagram, some preprocessing (notably application layer preprocessing) occurs after access control rule selection. This does **not** affect how you configure preprocessing in custom network analysis policies.

For more information, see:

- Decoding, Normalizing, and Preprocessing: Network Analysis Policies, page 18-3
- Access Control Rules: Intrusion Policy Selection, page 18-4
- Intrusion Inspection: Intrusion Policies, Rules, and Variable Sets, page 18-5
- Intrusion Event Generation, page 18-6

Decoding, Normalizing, and Preprocessing: Network Analysis Policies

License: Protection

Without decoding and preprocessing, the system could not appropriately evaluate traffic for intrusions because protocol differences would make pattern matching impossible. As shown in the diagram in Understanding How Policies Examine Traffic For Intrusions, page 18-2, network analysis policies govern these traffic-handling tasks:

- **after** traffic is filtered by Security Intelligence
- **before** traffic can be inspected by file or intrusion policies

A network analysis policy governs packet processing in phases. First the system decodes packets through the first three TCP/IP layers, then continues with normalizing, preprocessing, and detecting protocol anomalies.

- The packet decoder converts packet headers and payloads into a format that can be easily used by the preprocessors and later, intrusion rules. Each layer of the TCP/IP stack is decoded in turn, beginning with the data link layer and continuing through the network and transport layers. The packet decoder also detects various anomalous behaviors in packet headers. For more information, see Understanding Packet Decoding, page 24-16.

- In inline deployments, the inline normalization preprocessor reformats (normalizes) traffic to minimize the chances of attackers evading detection. It prepares packets for examination by other preprocessors and intrusion rules, and helps ensure that the packets the system processes are the same as the packets received by the hosts on your network. For more information, see Normalizing Inline Traffic, page 24-6.

- Various network and transport layers preprocessors detect attacks that exploit IP fragmentation, perform checksum validation, and perform TCP and UDP session preprocessing; see Configuring Transport & Network Layer Preprocessing, page 24-1.

  Note that some advanced transport and network preprocessor settings apply globally to all traffic handled by an access control policy. You configure these in the access control policy rather than in a network analysis policy; see Configuring Advanced Transport/Network Settings, page 24-1.

- Various application-layer protocol decoders normalize specific types of packet data into formats that the intrusion rules engine can analyze. Normalizing application-layer protocol encodings allows the system to effectively apply the same content-related intrusion rules to packets whose data is represented differently, and to obtain meaningful results. For more information, see Using Application Layer Preprocessors, page 22-1.

- The Modbus and DNP3 SCADA preprocessors detect traffic anomalies and provide data to intrusion rules. Supervisory Control and Data Acquisition (SCADA) protocols monitor, control, and acquire data from industrial, infrastructure, and facility processes such as manufacturing, production, water treatment, electric power distribution, airport and shipping systems, and so on. For more information, see Configuring SCADA Preprocessing, page 23-1.

- Several preprocessors allow you to detect specific threats, such as Back Orifice, portscans, SYN floods and other rate-based attacks; see Detecting Specific Threats, page 28-1.

  Note that you configure the sensitive data preprocessor, which detects sensitive data such as credit card numbers and Social Security numbers in ASCII text, in intrusion policies; see Detecting Sensitive Data, page 28-19.

In a newly created access control policy, one default network analysis policy governs preprocessing for *all* traffic for *all* intrusion policies invoked by the same parent access control policy. Initially, the system uses the Balanced Security and Connectivity network analysis policy as the default, but you can change it to another system-provided or custom network analysis policy. In a more complex deployment, advanced users can tailor traffic preprocessing options to specific security zones and networks by assigning different custom network analysis policies to preprocess matching traffic. For more information, see Comparing System-Provided with Custom Policies, page 18-7.

Access Control Rules: Intrusion Policy Selection

License: Protection

After initial preprocessing, access control rules (when present) evaluate traffic. In most cases, the first access control rule that a packet matches is the rule that handles that traffic; you can monitor, trust, block, or allow matching traffic.

When you allow traffic with an access control rule, the system can inspect the traffic for malware, prohibited files, and intrusions, in that order. Traffic not matching any access control rule is handled by the access control policy's default action, which can also inspect for intrusions.

> **Note**   All packets, **regardless** of which network analysis policy preprocesses them, are matched to configured access control rules—and thus are potentially subject to inspection by intrusion policies—in top-down order. For more information, see Limitations of Custom Policies, page 18-11.

The diagram in Understanding How Policies Examine Traffic For Intrusions, page 18-2 shows the flow of traffic through a device in an inline, intrusion prevention and AMP deployment, as follows:

- The access control rule allows matching traffic to proceed. The traffic is then inspected for prohibited files and malware by a file policy, and then for intrusions by an intrusion policy.

- In this scenario, the access control policy's default action allows matching traffic. The traffic is then inspected by an intrusion policy. You can (but do not have to) use a different intrusion policy when you associate intrusion policies with access control rules or the default action.

The example in the diagram does not include any blocking or trusting rules because the system does not inspect blocked or trusted traffic. For more information, see Using Rule Actions to Determine Traffic Handling and Inspection, page 6-6 and Setting Default Handling and Inspection for Network Traffic, page 4-4.

Intrusion Inspection: Intrusion Policies, Rules, and Variable Sets

License: Protection

You can use intrusion prevention as the system's last line of defense before traffic is allowed to proceed to its destination. Intrusion policies govern how the system inspects traffic for security violations and, in inline deployments, can block or alter malicious traffic. The main function of intrusion policies is to manage which intrusion and preprocessor rules are enabled and how they are configured.

Intrusion and Preprocessor Rules

An intrusion rule is a specified set of keywords and arguments that detects attempts to exploit vulnerabilities on your network; the system uses an intrusion rule to analyze network traffic to check if it matches the criteria in the rule. The system compares packets against the conditions specified in each rule and, if the packet data matches all the conditions specified in a rule, the rule triggers.

The system includes the following types of rules created by the VRT:

- *shared object intrusion rules*, which are compiled and cannot be modified (except for rule header information such as source and destination ports and IP addresses)

- *standard text intrusion rules*, which can be saved and modified as new custom instances of the rule.

- *preprocessor rules*, which are rules associated with preprocessors and packet decoder detection options in the network analysis policy. You cannot copy or edit preprocessor rules. Most preprocessor rules are disabled by default; you must enable them to use preprocessors to generate events and, in an inline deployment, drop offending packets.

When the system processes packets according to an intrusion policy, first a rule optimizer classifies all activated rules in subsets based on criteria such as: transport layer, application protocol, direction to or from the protected network, and so on. Then, the intrusion rules engine selects the appropriate rule subsets to apply to each packet. Finally, a multi-rule search engine performs three different types of searches to determine if the traffic matches the rule:

- The protocol field search looks for matches in particular fields in an application protocol.

- The generic content search looks for ASCII or binary byte matches in the packet payload.

- The packet anomaly search looks for packet headers and payloads that, rather than containing specific content, violate well-established protocols.

In a custom intrusion policy, you can tune detection by enabling and disabling rules, as well as by writing and adding your own standard text rules.

Variable Sets

Whenever the system uses an intrusion policy to evaluate traffic, it uses an associated *variable set*. Most variables in a set represent values commonly used in intrusion rules to identify source and destination IP addresses and ports. You can also use variables in intrusion policies to represent IP addresses in rule suppressions and dynamic rule states.

The system provides a single default variable set, which is comprised of predefined default variables. Most system-provided shared object rules and standard text rules use these predefined default variables to define networks and port numbers. For example, the majority of the rules use the variable $HOME_NET to specify the protected network and the variable $EXTERNAL_NET to specify the unprotected (or outside) network. In addition, specialized rules often use other predefined variables. For example, rules that detect exploits against web servers use the $HTTP_SERVERS and $HTTP_PORTS variables.

Tip   Even if you use system-provided intrusion policies, Cisco **strongly** recommends you modify key default variables in the default set. When you use variables that accurately reflect your network environment, processing is optimized and the system can monitor relevant systems for suspicious activity. Advanced users can create and use custom variable sets for pairing with one or more custom intrusion policies. For more information, see Optimizing Predefined Default Variables, page 2-13.

Intrusion Event Generation

License: Protection

When the system identifies a possible intrusion, it generates an *intrusion* or *preprocessor event* (sometimes collectively called *intrusion events*). You can view the data to gain a greater understanding of the attacks against your network assets. In an inline deployment, the system can also drop or replace packets that you know to be harmful.

Each intrusion event includes an event header and contains information about the event name and classification; the source and destination IP addresses; ports; the process that generated the event; and the date and time of the event, as well as contextual information about the source of the attack and its target. For packet-based events, the system also logs a copy of the decoded packet header and payload for the packet or packets that triggered the event.

The packet decoder, the preprocessors, and the intrusion rules engine can all cause the system to generate an event. For example:

- If the packet decoder (configured in the network analysis policy) receives an IP packet that is less than 20 bytes, which is the size of an IP datagram without any options or payload, the decoder interprets this as anomalous traffic. If, later, the accompanying decoder rule in the intrusion policy that examines the packet is enabled, the system generates a preprocessor event.

- If the IP defragmentation preprocessor encounters a series of overlapping IP fragments, the preprocessor interprets this as a possible attack and, when the accompanying preprocessor rule is enabled, the system generates a preprocessor event.

- Within the intrusion rules engine, most standard text rules and shared object rules are written so that they generate intrusion events when triggered by packets.

As the device accumulates intrusion events, you can begin your analysis of potential attacks. The system provides you with the tools you need to review intrusion events and evaluate whether they are important in the context of your network environment and your security policies.

Comparing System-Provided with Custom Policies

License: Protection

Creating a new access control policy is one of the first steps in managing traffic flow using the ASA FirePOWER module. By default, a newly created access control policy invokes system-provided network analysis and intrusion policies to examine traffic.

The following diagram shows how a newly created access control policy in an inline, intrusion-prevention deployment initially handles traffic. The preprocessing and intrusion prevention phases are highlighted.

Note how:

- A default network analysis policy governs the preprocessing of *all* traffic handled by the access control policy. Initially, the system-provided *Balanced Security and Connectivity network analysis policy* is the default.

- The default action of the access control policy allows all non-malicious traffic, as determined by the system-provided *Balanced Security and Connectivity intrusion policy.*

- The policy uses default Security Intelligence options (global whitelist and blacklist only), does not decrypt encrypted traffic with an SSL policy, and does not perform special handling and inspection of network traffic using access control rules.

A simple step you can take to tune your intrusion prevention deployment is to use a different set of system-provided network analysis and intrusion policies as your defaults. Cisco delivers several pairs of these policies with the ASA FirePOWER module.

Or, you can tailor your intrusion prevention deployment by creating and using custom policies. You may find that the preprocessor options, intrusion rule, and other advanced settings configured in those policies do not address the security needs of your network. By tuning your network analysis and intrusion policies you can configure, at a very granular level, how the system processes and inspects the traffic on your network for intrusions.

For more information, see:

Understanding the System-Provided Policies

License: Protection

Cisco delivers several pairs of network analysis and intrusion policies with the ASA FirePOWER module. By using system-provided network analysis and intrusion policies, you can take advantage of the experience of the Cisco Vulnerability Research Team (VRT). For these policies, the VRT sets intrusion and preprocessor rule states, as well as provides the initial configurations for preprocessors and other advanced settings. You can use system-provided policies as-is, or you can use them as the base for custom policies.

$\mathcal{Q}$

Tip   Even if you use system-provided network analysis and intrusion policies, you should configure the system's intrusion variables to accurately reflect your network environment. At a minimum, modify key default variables in the default set; see Optimizing Predefined Default Variables, page 2-13.

As new vulnerabilities become known, the VRT 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 may also delete rules from system-provided policies and provide new rule categories, as well as modify the default variable set.

If a rule update affects your deployment, the system marks affected intrusion and network analysis policies as out of date, as well as their parent access control policies. You must reapply an updated policy for its changes to take effect.

For your convenience, you can configure rule updates to automatically reapply affected intrusion policies, either alone or in combination with affected access control policies. This allows you to easily and automatically keep your deployment up-to-date to protect against recently discovered exploits and intrusions.

To ensure up-to-date preprocessing settings, you **must** reapply access control policies, which also reapplies any associated SSL, network analysis, and file policies that are different from those currently running, and can also can update default values for advanced preprocessing and performance options. For more information, see Importing Rule Updates and Local Rule Files, page 46-9.

Cisco delivers the following network analysis and intrusion policies with the ASA FirePOWER module:

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 organizations. The system uses the Balanced Security and Connectivity policies and settings as defaults in most cases.

Connectivity Over Security network analysis and intrusion policies

These policies are built for organizations where connectivity (being able 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 organizations 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.

No Rules Active intrusion policy

In the No Rules Active intrusion policy, all intrusion rules and advanced settings are disabled. This policy provides a starting point if you want to create your own intrusion policy instead of basing it on the enabled rules in one of the other system-provided policies.

⚠

Caution   Cisco uses another policy, Experimental Policy 1, for testing purposes. Do not use it unless instructed to do so by a Cisco representative.

Benefits of Custom Policies

License: Protection

You may find that the preprocessor options, intrusion rules, and other advanced settings configured in the system-provided network analysis and intrusion policies do not fully address the security needs of your organization.

Building custom policies can improve the performance of the system in your environment and can provide a focused view of the malicious traffic and policy violations occurring on your network. By creating and tuning custom policies you can configure, at a very granular level, how the system processes and inspects the traffic on your network for intrusions.

All custom policies have a base policy, also called a base layer, which defines the default settings for all configurations in the policy. A layer is a building block that you can use to efficiently manage multiple network analysis or intrusion policies; see Using Layers in a Network Analysis or Intrusion Policy, page 19-1.

In most cases, you base custom policies on system-provided policies, but you can use another custom policy. However, all custom policies have a system-provided policy as the eventual base in a policy chain. Because rule updates can modify system-provided policies, importing a rule update may affect you even if you are using a custom policy as your base. If a rule update affects your policies, the module interface marks affected policies as out of date. For more information, see Allowing Rule Updates to Modify a System-Provided Base Policy, page 19-4.

For more information, see:

- Benefits of a Custom Network Analysis Policy, page 18-9
- Benefits of Custom Intrusion Policies, page 18-10

Benefits of a Custom Network Analysis Policy

License: Protection

By default, one network analysis policy preprocesses all unencrypted traffic handled by the access control policy. That means that all packets are decoded and preprocessed according to the same settings, regardless of the intrusion policy (and therefore intrusion rule set) that later examines them.

Initially, the system-provided Balanced Security and Connectivity network analysis policy is the default. A simple way to tune preprocessing is to create and use a custom network analysis policy as the default; see Setting the Default Network Analysis Policy for Access Control, page 20-3.

Tuning options available vary by preprocessor, but some of the ways you can tune preprocessors and decoders include:

- You can disable preprocessors that do not apply to the traffic you are monitoring. For example, the HTTP Inspect preprocessor normalizes HTTP traffic. If you are confident that your network does not include any web servers using Microsoft Internet Information Services (IIS), you can disable the preprocessor option that looks for IIS-specific traffic and thereby reduce system processing overhead.

> **Note** If you disable a preprocessor in a custom network analysis policy, but the system needs to use that preprocessor to later evaluate packets against an enabled intrusion or preprocessor rule, the system automatically enables and uses the preprocessor although the preprocessor remains disabled in the network analysis policy user interface.

- Specify ports, where appropriate, to focus the activity of certain preprocessors. For example, you can identify additional ports to monitor for DNS server responses or encrypted SSL sessions, or ports on which you decode telnet, HTTP, and RPC traffic

For advanced users with complex deployments, you can create multiple network analysis policies, each tailored to preprocess traffic differently. Then, you can configure the system to use those policies to govern the preprocessing of traffic using different security zones or networks.

> **Note** Tailoring preprocessing using custom network analysis policies—especially multiple network analysis policies—is an advanced task. Because preprocessing and intrusion inspection are so closely related, you **must** be careful to allow the network analysis and intrusion policies examining a single packet to complement each other. For more information, see Limitations of Custom Policies, page 18-11.

Benefits of Custom Intrusion Policies

License: Protection

In a newly created access control policy initially configured to perform intrusion prevention, the default action allows all traffic, but first inspects it with the system-provided Balanced Security and Connectivity intrusion policy. Unless you add access control rules or change the default action, all traffic is inspected by that intrusion policy; see the diagram in Comparing System-Provided with Custom Policies, page 18-7.

To customize your intrusion prevention deployment, you can create multiple intrusion policies, each tailored to inspect traffic differently. Then, configure an access control policy with rules that specify which policy inspects which traffic. Access control rules can be simple or complex, matching and inspecting traffic using multiple criteria including security zone, network or geographical location, port, application, requested URL, or user. The scenario in Understanding How Policies Examine Traffic For Intrusions, page 18-2 shows a deployment where traffic is inspected by one of two intrusion policies.

The main function of intrusion policies is to manage which intrusion and preprocessor rules are enabled and how they are configured, as follows:

- Within each intrusion policy, you should verify that all rules applicable to your environment are enabled, and improve performance by disabling rules that are not applicable to your environment. In an inline deployment, you can specify which rules should drop or modify malicious packets. For more information, see Setting Rule States, page 27-19.

- You can modify existing rules and write new standard text rules as needed to catch new exploits or to enforce your security policies; see Understanding and Writing Intrusion Rules, page 30-1.

Other customizations you might make to an intrusion policy include:

- The sensitive data preprocessor detects sensitive data such as credit card numbers and Social Security numbers in ASCII text. Note that other preprocessors that detect specific threats (back orifice attacks, several portscan types, and rate-based attacks that attempt to overwhelm your network with excessive traffic) are configured in network analysis policies. For more information, see Detecting Specific Threats, page 28-1.

- Global thresholds cause the system to generate events based on how many times traffic matching an intrusion rule originates from or is targeted to a specific address or address range within a specified time period. This helps prevent the system from being overwhelmed with a large number of events. For more information, see Globally Limiting Intrusion Event Logging, page 29-1.

- Suppressing intrusion event notifications and setting thresholds for individual rules or entire intrusion policies can also can prevent the system from being overwhelmed with a large number of events. For more information, see Filtering Intrusion Event Notification Per Policy, page 27-20.

- In addition to intrusion events, you can enable logging to syslog facilities or send event data to an SNMP trap server. Per policy, you can specify intrusion event notification limits, set up intrusion event notification to external logging facilities, and configure external responses to intrusion events. For more information, see Configuring External Alerting for Intrusion Rules, page 39-1.

Limitations of Custom Policies

License: Protection

Because preprocessing and intrusion inspection are so closely related, you **must** be careful that your configuration allows the network analysis and intrusion policies processing and examining a single packet to complement each other.

By default, the system uses one network analysis policy to preprocess all traffic. The following diagram shows how a newly created access control policy in an inline, intrusion-prevention deployment initially handles traffic. The preprocessing and intrusion prevention phases are highlighted.

Notice how a default network analysis policy governs the preprocessing of all traffic handled by the access control policy. Initially, the system-provided Balanced Security and Connectivity network analysis policy is the default.

A simple way to tune preprocessing is to create and use a custom network analysis policy as the default, as summarized in Benefits of a Custom Network Analysis Policy, page 18-9. However, if you disable a preprocessor in a custom network analysis policy but the system needs to evaluate preprocessed packets against an enabled intrusion or preprocessor rule, the system automatically enables and uses the preprocessor although it remains disabled in the network analysis policy user interface.

Note In order to get the performance benefits of disabling a preprocessor, you **must** make sure that none of your intrusion policies have enabled rules that require that preprocessor.

An additional challenge arises if you use multiple custom network analysis policies. For advanced users with complex deployments, you can tailor preprocessing to specific security zones and networks by assigning custom network analysis policies to preprocess matching traffic. To accomplish this, you add custom *network analysis rules* to your access control policy. Each rule has an associated network analysis policy that governs the preprocessing of traffic that matches the rule.

Tip You configure network analysis rules as an advanced setting in an access control policy. Unlike other types of rules in the ASA FirePOWER module, network analysis rules invoke—rather than being contained by—network analysis policies.

The system matches packets to any configured network analysis rules in top-down order by rule number. Traffic that does not match any network analysis rule is preprocessed by the default network analysis policy. While this allows you a great deal of flexibility in preprocessing traffic, keep in mind that all packets, **regardless** of which network analysis policy preprocessed them, are subsequently matched to access control rules—and thus to potential inspection by intrusion policies—in their own process. In other words, preprocessing a packet with a particular network analysis policy does **not** guarantee that the packet will be examined with any particular intrusion policy. You **must** carefully configure your access control policy so it invokes the correct network analysis and intrusion policies to evaluate a particular packet.

The following diagram shows in focused detail how the network analysis policy (preprocessing) selection phase occurs before and separately from the intrusion prevention (rules) phase. For simplicity, the diagram eliminates the file/malware inspection phases. It also highlights the default network analysis and default-action intrusion policies.
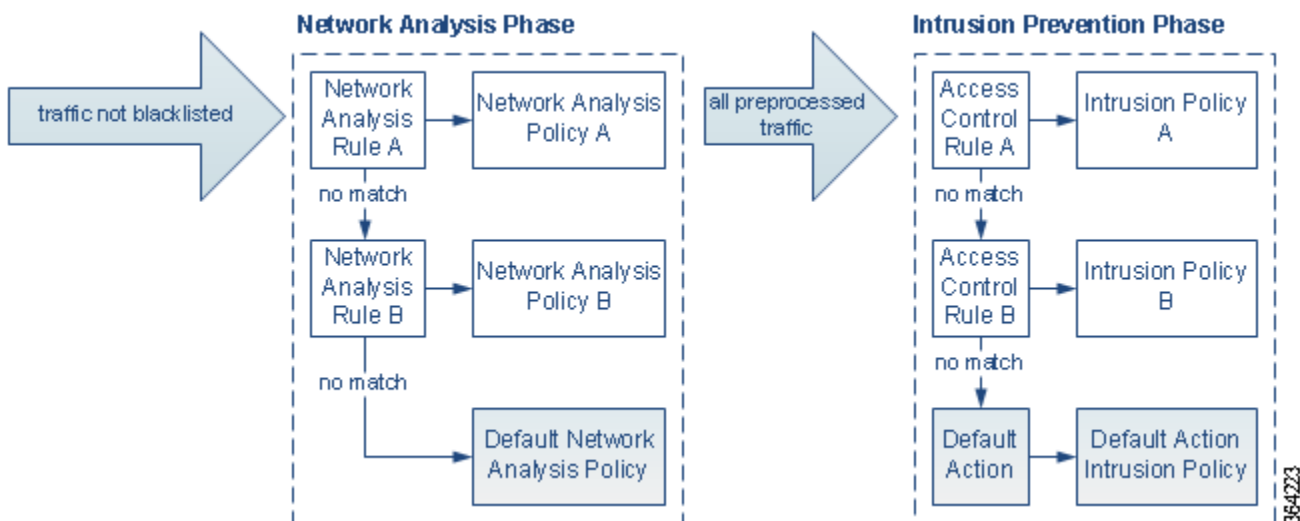
In this scenario, an access control policy is configured with two network analysis rules and a default network analysis policy:

- Network Analysis Rule A preprocesses matching traffic with Network Analysis Policy A. Later, you want this traffic to be inspected by Intrusion Policy A.

- Network Analysis Rule B preprocesses matching traffic with Network Analysis Policy B. Later, you want this traffic to be inspected by Intrusion Policy B.

- All remaining traffic is preprocessed with the default network analysis policy. Later, you want this traffic to be inspected by the intrusion policy associated with the access control policy's default action.

After the system preprocesses traffic, it can examine the traffic for intrusions. The diagram shows an access control policy with two access control rules and a default action:

- Access Control Rule A allows matching traffic. The traffic is then inspected by Intrusion Policy A.

- Access Control Rule B allows matching traffic. The traffic is then inspected by Intrusion Policy B.

- The access control policy's default action allows matching traffic. The traffic is then inspected by the default action's intrusion policy.

Each packet's handling is governed by a network analysis policy and intrusion policy pair, but the system does **not** coordinate the pair for you. Consider a scenario where you misconfigure your access control policy so that Network Analysis Rule A and Access Control Rule A do not process the same traffic. For example, you could intend the paired policies to govern the handling of traffic on a particular security zone, but you mistakenly use different zones in the two rules' conditions. This could cause traffic to be incorrectly preprocessed. For this reason, tailoring preprocessing using network analysis rules and custom policies is an **advanced** task.

Note that for a single connection, although the system selects a network analysis policy before an access control rule, some preprocessing (notably application layer preprocessing) occurs after access control rule selection. This does **not** affect how you configure preprocessing in custom network analysis policies.

Using the Navigation Panel

License: Protection

Network analysis and intrusion policies use similar user interfaces to edit and save changes to their configurations; see:

- Editing Network Analysis Policies, page 21-3
- Editing Intrusion Policies, page 26-4

A navigation panel appears on the left side of the user interface when you are editing either type of policy. The following graphic shows the navigation panel for the network analysis policy (left) and the intrusion policy (right).

A dividing line separates the navigation panel into links to policy settings you can configure with (below) or without (above) direct interaction with policy layers. To navigate to any settings page, click its name in the navigation panel. Dark shading of an item in the navigation panel highlights your current settings page. For example, in the illustration above the Policy Information page would be displayed to the right of the navigation panel.

Policy Information

The Policy Information page provides configuration options for commonly used settings. As shown in the illustration for the network analysis policy panel above, a policy change icon (⚠) appears next to **Policy Information** in the navigation panel when the policy contains unsaved changes. The icon disappears when you save your changes.

Rules (intrusion policy only)

The Rules page in an intrusion policy allows you to configure rule states and other settings for shared object rules, standard text rules, and preprocessor rules. For more information, see Tuning Intrusion Policies Using Rules, page 27-1.

Settings (network analysis policy) and Advanced Settings (intrusion policy)

The Settings page in a network analysis policy allows you to enable or disable preprocessors and access preprocessor configuration pages. Expanding the **Settings** link displays sublinks to individual configuration pages for all enabled preprocessors in the policy. For more information, see Configuring Preprocessors in a Network Analysis Policy, page 21-6.

The Advanced Settings page in an intrusion policy allows you to enable or disable advanced settings and access configuration pages for those advanced settings. Expanding the **Advanced Settings** link displays sublinks to individual configuration pages for all enabled advanced settings in the policy. For more information, see Configuring Advanced Settings in an Intrusion Policy, page 26-6.

Policy Layers

The Policy Layers page displays a summary of the layers that comprise your network analysis or intrusion policy. Expanding the Policy Layers link displays sublinks to summary pages for the layers in your policy. Expanding each layer sublink displays further sublinks to the configuration pages for all rules, preprocessors, or advanced settings that are enabled in the layer. For more information, see Using Layers in a Network Analysis or Intrusion Policy, page 19-1.

Resolving Conflicts and Committing Policy Changes

License: Protection

When you edit a network analysis or intrusion policy, you must save (or *commit*) your changes before the system recognizes them.

Note   After you save, you must apply a network analysis or intrusion policy for your changes to take effect. If you apply a policy without saving, the system uses the most recently saved configuration. Although you can reapply an intrusion policy independently, network analysis policies are applied with their parent access control policy.

Resolving Editing Conflicts

The Network Analysis Policy page and Intrusion Policy page display whether each policy has unsaved changes; see Editing Network Analysis Policies, page 21-3 and Editing Intrusion Policies, page 26-4.

Cisco recommends that only one person edit a policy at a time. If you are editing the same network analysis or intrusion policy via multiple user interface instances as the same user, and you save your changes for one instance, you cannot save your changes for the other instance.

Resolving Configuration Dependencies

To perform their particular analysis, many preprocessors and intrusion rules require that traffic first be decoded or preprocessed in a certain way, or have other dependencies. When you save a network analysis or intrusion policy, the system either automatically enables required settings, or warns you that disabled settings will have no effect on traffic, as follows:

- You cannot save an intrusion policy if you added an SNMP rule alert but did not configure SNMP alerting. You must either configure SNMP alerting or disable the rule alert, then save again.

- You cannot save an intrusion policy if it includes enabled sensitive data rules but you have not enabled the sensitive data preprocessor. You must either allow the system to enable the preprocessor and save the policy, or disable the rules and save again.

- If you disable a required preprocessor in a network analysis policy, you can still save the policy. However, the system automatically uses the disabled preprocessor with its current settings, even though the preprocessor remains disabled in the user interface. For more information, see Limitations of Custom Policies, page 18-11.

- If you disable inline mode in a network analysis policy but enable the Inline Normalization preprocessor, you can still save the policy. However, the system warns you that normalization settings will be ignored. Disabling inline mode also causes the system to ignore other settings that allow preprocessors to modify or block traffic, including checksum verification and rate-based attack prevention. For more information, see Allowing Preprocessors to Affect Traffic in Inline Deployments, page 21-5 and Normalizing Inline Traffic, page 24-6.

Committing, Discarding, and Caching Policy Changes

While editing a network analysis or intrusion policy, if you exit the policy editor without saving your changes, the system caches those changes. Your changes are cached even when you log out of the system or experience a system crash. The system cache can store unsaved changes for one network analysis and one intrusion policy; you must commit or discard your changes before editing another policy of the same type. The system discards the cached changes when you edit another policy without saving your changes to the first policy, or when you import an intrusion rule update.

You can commit or discard policy changes on the Policy Information page of either the network analysis or intrusion policy editor; see Editing Network Analysis Policies, page 21-3 and Editing Intrusion Policies, page 26-4.

The following table summarizes how to save or discard changes to a network analysis or intrusion policy.

Table 18-1   Committing Changes to a Network Analysis or Intrusion Policy

To...	On the Policy Information page, you can...
save changes to the policy	click **Commit Changes**. Optionally, enter a comment; click **OK** to continue committing.
discard all unsaved changes	click **Discard Changes**, then click **OK** to discard your changes and go to the Intrusion Policy page. If you do not want to discard your changes, click **Cancel** to return to the Policy Information page.
exit the policy, but cache changes	select any menu or other path to another page. On exiting, click **Leave page** when prompted, or click **Stay on page** to remain in the advanced editor.