Understanding Access Control

- Introduction to Access Control, on page 1
- Access Control Policy Default Action, on page 1
- Deep Inspection Using File and Intrusion Policies, on page 3
- Access Control Policy Inheritance, on page 7

Introduction to Access Control

Access control is a hierarchical policy-based feature that allows you to specify, inspect, and log (non-fast-pathed) network traffic.

Each managed device can be targeted by one access control policy. The data that the policy's *target devices* collect about your network traffic can be used to filter and control that traffic based on:

- simple, easily determined transport and network layer characteristics: source and destination, port, protocol, and so on
- the latest contextual information on the traffic, including characteristics such as reputation, risk, business relevance, application used, or URL visited
- realm, user, user group, or ISE attribute
- custom Security Group Tag (SGT)
- characteristics of encrypted traffic; you can also decrypt this traffic for further analysis
- whether unencrypted or decrypted traffic contains a prohibited file, detected malware, or intrusion attempt

Each type of traffic inspection and control occurs where it makes the most sense for maximum flexibility and performance. For example, reputation-based blocking uses simple source and destination data, so it can block prohibited traffic early in the process. In contrast, detecting and blocking intrusions and exploits is a last-line defense.

Access Control Policy Default Action

A newly created access control policy directs its target devices to handle all traffic using its *default action*.

In a simple access control policy, the default action specifies how target devices handle all traffic. In a more complex policy, the default action handles traffic that:

- is not trusted by Intelligent Application Bypass
- is not on a Security Intelligence Block list
- is not blocked by SSL inspection (encrypted traffic only)
- matches none of the rules in the policy (except Monitor rules, which match and log—but do not handle or inspect—traffic)

The access control policy default action can block or trust traffic without further inspection, or inspect traffic for intrusions and discovery data.

Note You **cannot** perform file or malware inspection on traffic handled by the default action. Logging for connections handled by the default action is initially disabled, though you can enable it.

If you are using policy inheritance, the default action for the lowest-level descendant determines final traffic handling. Although an access control policy can inherit its default action from its base policy, you cannot enforce this inheritance.

The following table describes the types of inspection you can perform on traffic handled by each default action.

Table 1: Access Control Policy Default Actions

Default Action	Effect on Traffic	Inspection Type and Policy
Access Control: Block All Traffic	block without further inspection	none
Access Control: Trust All Traffic	trust (allow to its final destination without further inspection)	none
Intrusion Prevention	allow, as long as it is passed by the intrusion policy you specify	intrusion, using the specified intrusion policy and associated variable set, and discovery, using the network discovery policy
Network Discovery Only	allow	discovery only, using the network discovery policy
Inherit from base policy	defined in base policy	defined in base policy

The following diagram illustrates the table.

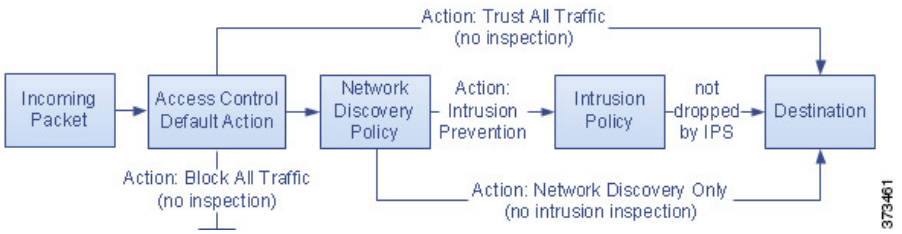

The following diagrams illustrate the **Block All Traffic** and **Trust All Traffic** default actions.

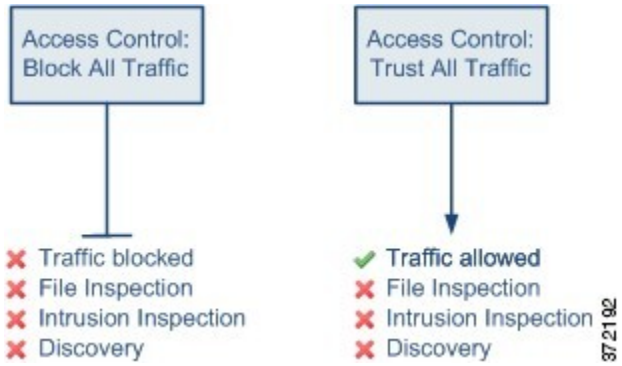

The following diagrams illustrate the **Intrusion Prevention** and **Network Discovery Only** default actions.

Tip The purpose of **Network Discovery Only** is to improve performance in a discovery-only deployment. Different configurations can disable discovery if you are only interested in intrusion detection and prevention.

Related Topics

Performance Considerations for Limited Deployments
Logging Connections with a Policy Default Action

Deep Inspection Using File and Intrusion Policies

Deep inspection uses intrusion and file policies as the last line of defense before traffic is allowed to its destination.

- *Intrusion policies* govern the system's intrusion prevention capabilities.

  For complete information, see Intrusion Detection and Prevention.

- *File policies* govern the system's file control and AMP for Networks capabilities.

  For complete information, see File Policies and Malware Protection.

Access control occurs before deep inspection; access control rules and the access control default action determine which traffic is inspected by intrusion and file policies.

By associating an intrusion or file policy with an access control rule, you are telling the system that before it passes traffic that matches the access control rule's conditions, you first want to inspect the traffic with an intrusion policy, a file policy, or both.

In an access control policy, you can associate one intrusion policy with each Allow and Interactive Block rule, as well as with the default action. Every unique **pair** of intrusion policy and variable set counts as one policy.

To associate intrusion and file policies with an access control rule, see:

- Access Control Rule Configuration to Perform Intrusion Prevention
- Configuring an Access Control Rule to Perform Malware Protection

Note By default, the system disables intrusion and file inspection of encrypted payloads. This helps reduce false positives and improve performance when an encrypted connection matches an access control rule that has intrusion and file inspection configured.

Related Topics

How Policies Examine Traffic For Intrusions
File Policies

Access Control Traffic Handling with Intrusion and File Policies

The following diagram shows the flow of traffic in an inline intrusion prevention and AMP for Networks deployment, as governed by an access control policy that contains four different types of access control rules and a default action.

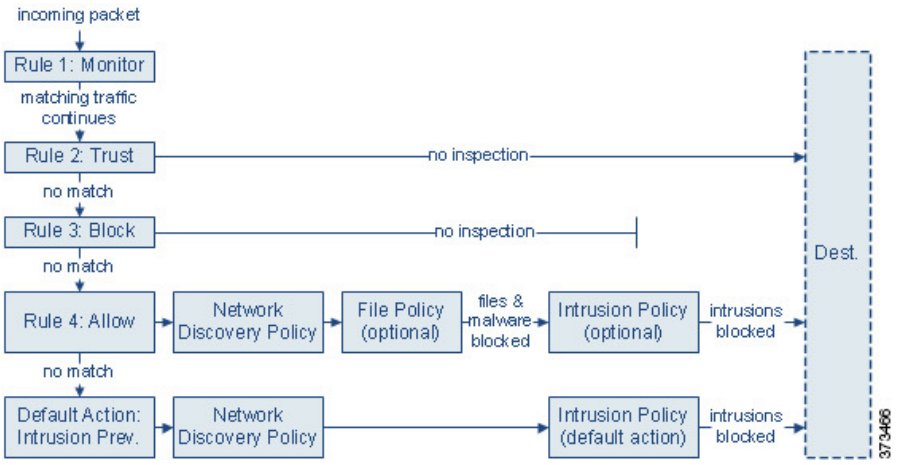

In the scenario above, the first three access control rules in the policy—Monitor, Trust, and Block—cannot inspect matching traffic. Monitor rules track and log but do not inspect network traffic, so the system continues to match traffic against additional rules to determine whether to permit or deny it. (However, see an important exception and caveat at Access Control Rule Monitor Action.) Trust and Block rules handle matching traffic without further inspection of any kind, while traffic that does not match continues to the next access control rule.

The fourth and final rule in the policy, an Allow rule, invokes various other policies to inspect and handle matching traffic, in the following order:

- **Discovery: Network Discovery Policy**—First, the network discovery policy inspects traffic for discovery data. Discovery is passive analysis and does not affect the flow of traffic. Although you do not explicitly enable discovery, you can enhance or disable it. However, allowing traffic does not automatically guarantee discovery data collection. The system performs discovery only for connections involving IP addresses that are explicitly monitored by your network discovery policy.
- **AMP for Networks and File Control: File Policy**—After traffic is inspected by discovery, the system can inspect it for prohibited files and malware. AMP for Networks detects and optionally blocks malware in many types of files, including PDFs, Microsoft Office documents, and others. If your organization wants to block not only the transmission of malware files, but all files of a specific type (regardless of whether the files contain malware), *file control* allows you to monitor network traffic for transmissions of specific file types, then either block or allow the file.
- **Intrusion Prevention: Intrusion Policy**—After file inspection, the system can inspect traffic for intrusions and exploits. An intrusion policy examines decoded packets for attacks based on patterns, and can block or alter malicious traffic. Intrusion policies are paired with *variable sets*, which allow you to use named values to accurately reflect your network environment.
- **Destination**—Traffic that passes all the checks described above passes to its destination.

An Interactive Block rule (not shown in the diagram) has the same inspection options as an Allow rule. This is so you can inspect traffic for malicious content when a user bypasses a blocked website by clicking through a warning page.

Traffic that does not match any access control rules in the policy with an action other than Monitor is handled by the default action. In this scenario, the default action is an Intrusion Prevention action, which allows traffic to its final destination as long as it is passed by the intrusion policy you specify. In a different deployment, you might have a default action that trusts or blocks all traffic without further inspection. Note that the system can inspect traffic allowed by the default action for discovery data and intrusions, but not prohibited files or malware. You **cannot** associate a file policy with the access control default action.

Note Sometimes, when a connection is analyzed by an access control policy, the system must process the first few packets in that connection, **allowing them to pass**, before it can decide which access control rule (if any) will handle the traffic. However, so these packets do not reach their destination uninspected, you can specify an intrusion policy (in the Advanced settings for the access control policy) to inspect these packets and generate intrusion events.

File and Intrusion Inspection Order

In your access control policy, you can associate multiple Allow and Interactive Block rules with different intrusion and file policies to match inspection profiles to various types of traffic.

Note Traffic allowed by an Intrusion Prevention or Network Discovery Only default action can be inspected for discovery data and intrusions, but cannot be inspected for prohibited files or malware. You **cannot** associate a file policy with the access control default action.

You do not have to perform both file and intrusion inspection in the same rule. For a connection matching an Allow or Interactive Block rule:

- without a file policy, traffic flow is determined by the intrusion policy
- without an intrusion policy, traffic flow is determined by the file policy
- without either, allowed traffic is inspected by network discovery only

Tip The system does not perform any kind of inspection on trusted traffic. Although configuring an Allow rule with neither an intrusion nor file policy passes traffic like a Trust rule, Allow rules let you perform discovery on matching traffic.

The diagram below illustrates the types of inspection you can perform on traffic that meets the conditions of either an Allow or user-bypassed Interactive Block access control rule. For simplicity, the diagram displays traffic flow for situations where both (or neither) an intrusion and a file policy are associated with a single access control rule.

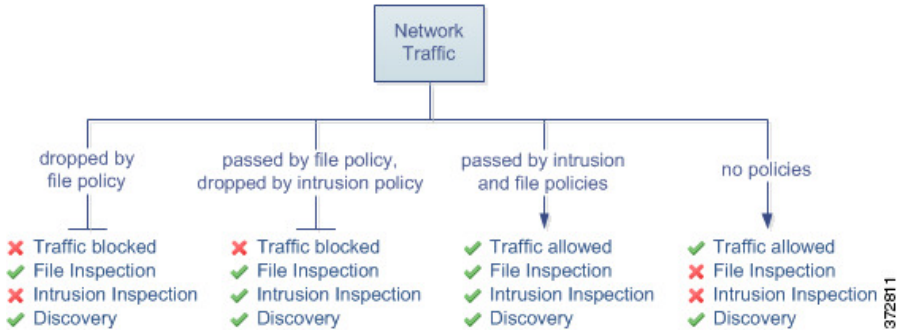

For any single connection handled by an access control rule, file inspection occurs before intrusion inspection. That is, the system does not inspect files blocked by a file policy for intrusions. Within file inspection, simple blocking by type takes precedence over malware inspection and blocking.

For example, consider a scenario where you normally want to allow certain network traffic as defined in an access control rule. However, as a precaution, you want to block the download of executable files, examine downloaded PDFs for malware and block any instances you find, and perform intrusion inspection on the traffic.

You create an access control policy with a rule that matches the characteristics of the traffic you want to provisionally allow, and associate it with both an intrusion policy and a file policy. The file policy blocks the download of all executables, and also inspects and blocks PDFs containing malware:

- First, the system blocks the download of all executables, based on simple type matching specified in the file policy. Because they are immediately blocked, these files are subject to neither malware nor intrusion inspection.
- Next, the system performs malware cloud lookups for PDFs downloaded to a host on your network. Any PDFs with a malware disposition are blocked, and are not subject to intrusion inspection.
- Finally, the system uses the intrusion policy associated with the access control rule to inspect any remaining traffic, including files not blocked by the file policy.

Note Until a file is detected and blocked in a session, packets from the session may be subject to intrusion inspection.

Access Control Policy Inheritance

Especially useful in multidomain deployments, you can nest access control policies, where each policy inherits the rules and settings from an ancestor (or *base*) policy. You can enforce this inheritance, or allow lower-level policies to override their ancestors.

Access control uses a hierarchical policy-based implementation. Just as you create a domain hierarchy, you can create a corresponding hierarchy of access control policies. A *descendant*, or *child*, access control policy inherits rules and settings from its direct *parent*, or base, policy. That base policy may have its own parent policy from which it inherits rules and settings, and so on.

An access control policy's rules are nested between its parent policy's Mandatory and Default rule sections. This implementation enforces Mandatory rules from ancestor policies, but allows the current policy to write rules that preempt Default rules from ancestor policies.

You can lock the following settings to enforce them in all descendant policies. Descendant policies can override unlocked settings.

- Security Intelligence — connections that are allowed or blocked based on the latest reputation intelligence for IP addresses, URLs, and domain names.
- HTTP Response pages — Displaying a custom or system-provided response page when you block a user's website request.
- Advanced settings — Specifying associated subpolicies, network analysis settings, performance settings, and other general options.

When using policy inheritance, the default action for the lowest-level descendant determines final traffic handling. Although an access control policy can inherit its default action from an ancestor policy, you cannot enforce this inheritance.

Policy Inheritance and Multitenancy

Access control's hierarchical policy-based implementation complements multitenancy.

In a typical multidomain deployment, access control policy hierarchy corresponds to domain structure, and you apply the lowest-level access control policy to managed devices. This implementation allows selective access control enforcement at a higher domain level, while lower-level domain administrators can tailor deployment-specific settings. (You must use roles, not policy inheritance and enforcement alone, to restrict administrators in descendant domains.)

For example, as a Global domain administrator for your organization, you can create an access control policy at the Global level. You can then require that all your devices, which are divided into subdomain by function, use that Global-level policy as a base policy.

When subdomain administrators log into the Firepower Management Center to configure access control, they can deploy the Global-level policy as-is. Or, they can create and deploy a descendant access control policy within the boundaries of the Global-level policy.

Note Although the most useful implementation of access control inheritance and enforcement complements multitenancy, you can create a hierarchy of access control policies within a single domain. You can also assign and deploy access control policies at any level.

Related Topics

Managing Access Control Policy Inheritance
Blocking Traffic with Security Intelligence
HTTP Response Pages and Interactive Blocking
Access Control Policy Advanced Settings