

Configuring Security Contexts

This section contains the following topics:

- [Security Context Overview, page 9-1](#)
- [Enabling or Disabling Multiple Context Mode at the CLI, page 9-8](#)
- [Managing Memory for Rules, page 9-10](#)
- [Configuring Resource Classes, page 9-17](#)
- [Configuring Security Contexts, page 9-23](#)

Security Context Overview

You can partition a single FWSM into multiple virtual devices, known as security contexts. Each context has its own security policy, interfaces, and administrators. Multiple contexts are similar to having multiple standalone devices. Many features are supported in multiple context mode, including routing tables, firewall features, and management. Some features are not supported, including most dynamic routing protocols.

This section provides an overview of security contexts, and includes the following topics:

- [Common Uses for Security Contexts, page 9-1](#)
- [Unsupported Features, page 9-2](#)
- [Context Configuration Files, page 9-2](#)
- [How the FWSM Classifies Packets, page 9-3](#)
- [Sharing Interfaces Between Contexts, page 9-6](#)

Common Uses for Security Contexts

You might want to use multiple security contexts in the following situations:

- You are a service provider and want to sell security services to many customers. By enabling multiple security contexts on the FWSM, you can implement a cost-effective, space-saving solution that keeps all customer traffic separate and secure, and also eases configuration.
- You are a large enterprise or a college campus and want to keep departments completely separate.
- You are an enterprise that wants to provide distinct security policies to different departments.
- You have any network that requires more than one firewall.

Unsupported Features

Multiple context mode does not support the following features:

- Most dynamic routing protocols. BGP stub mode is supported.
Security contexts support only static routes or BGP stub mode. You cannot enable OSPF or RIP in multiple context mode.
- Multicast routing. Multicast bridging is supported.

Context Configuration Files

This section describes how the FWSM implements multiple context mode configurations and includes the following sections:

- [Context Configurations, page 9-2](#)
- [System Configuration, page 9-2](#)
- [Admin Context Configuration, page 9-2](#)

Context Configurations

The FWSM includes a configuration for each context that identifies the security policy, interfaces, and almost all the options you can configure on a standalone device. You can store context configurations on the internal Flash memory or the external Flash memory card, or you can download them from a TFTP, FTP, or HTTP(S) server.

System Configuration

The system administrator adds and manages contexts by configuring each context configuration location, allocated interfaces, and other context operating parameters in the system configuration, which, like a single mode configuration, is the startup configuration. The system configuration identifies basic settings for the FWSM. The system configuration does not include any network interfaces or network settings for itself; rather, when the system needs to access network resources (such as downloading the contexts from the server), it uses one of the contexts that is designated as the *admin context*. The system configuration does include a specialized failover interface for failover traffic only.

Admin Context Configuration

The admin context is just like any other context, except that when a user logs in to the admin context, then that user has system administrator rights and can access the system and all other contexts. The admin context is not restricted in any way, and can be used as a regular context. However, because logging into the admin context grants you administrator privileges over all contexts, you might need to restrict access to the admin context to appropriate users. The admin context must reside on Flash memory, and not remotely.

If your system is already in multiple context mode, or if you convert from single mode, the admin context is created automatically as a file on the internal Flash memory called `admin.cfg`. This context is named “admin.” If you do not want to use `admin.cfg` as the admin context, you can change the admin context.

How the FWSM Classifies Packets

Each packet that enters the FWSM must be classified, so that the FWSM can determine to which context to send a packet. The FWSM uses only one global MAC address across all interfaces. A single MAC address is usually not a problem unless multiple contexts want to share an interface. A router cannot direct packets to IP addresses on the same network if all IP addresses resolve to the same MAC address. Moreover, the bridging table of the switch would constantly change as the MAC address moves from one interface to another. The purpose of the security context classifier is to resolve this situation.

This section includes the following topics:

- [Valid Classifier Criteria, page 9-3](#)
- [Invalid Classifier Criteria, page 9-4](#)
- [Classification Examples, page 9-4](#)

Valid Classifier Criteria

If only one context is associated with the ingress interface, the FWSM classifies the packet into that context. In transparent firewall mode, unique interfaces for contexts are required, so this method is used to classify packets at all times.

If multiple contexts share an interface, then the classifier intercepts the packet and performs a destination IP address lookup. All other fields are ignored; only the destination IP address is used. To use the destination address for classification, the classifier must have knowledge about the subnets located behind each security context. The classifier relies on active NAT sessions to determine the subnets in each context. Active NAT sessions are created either by **static** commands, which create a permanent session, or by active dynamic NAT sessions.

For example, the classifier gains knowledge about subnets 10.10.10.0, 10.20.10.0 and 10.30.10.0 when the context administrators configure **static** commands in each context:

- Context A:

```
static (inside,shared) 10.10.10.0 10.10.10.0 netmask 255.255.255.0
```

- Context B:

```
static (inside,shared) 10.20.10.0 10.20.10.0 netmask 255.255.255.0
```

- Context C:

```
static (inside,shared) 10.30.10.0 10.30.10.0 netmask 255.255.255.0
```

If you use dynamic NAT, an active NAT session is created when the real host creates a connection through the shared interface. For traffic returning to the host, the active NAT session is used to classify the packet.

To quickly identify possible overlaps between different contexts, a situation that leads to connectivity problems, enter the **show np 3 static** command in the system execution space.

**Note**

For management traffic destined for an interface, the interface IP address is used for classification.

Invalid Classifier Criteria

The following configurations are not used for packet classification:

- NAT exemption—The classifier does not use a NAT exemption configuration for classification purposes because NAT exemption does not identify the mapped (shared) interface.
- Routing table—The classifier does not use the routing table for classification. For example, if a context includes a static route that points to an external router as the next-hop to a subnet, and a different context includes a **static** command for the same subnet, then the classifier uses the **static** command to classify packets destined for that subnet and ignores the static route.

Classification Examples

Figure 9-1 shows multiple contexts sharing an outside interface, while the inside interfaces are unique, allowing overlapping IP addresses. The classifier assigns the packet to Context B because Context B includes the address translation that matches the destination address.

Figure 9-1 Packet Classification with a Shared Interface



Note that all new incoming traffic must be classified, even from inside networks. [Figure 9-2](#) shows a host on the Context B inside network accessing the Internet. The classifier assigns the packet to Context B because the ingress interface is VLAN 300, which is assigned to Context B.

Figure 9-2 *Incoming Traffic from Inside Networks*




Figure 9-3 shows a transparent firewall with a host on the Context B inside network accessing the Internet. The classifier assigns the packet to Context B because the ingress interface is VLAN 300, which is assigned to Context B.

Figure 9-3 *Transparent Firewall Contexts*



Sharing Interfaces Between Contexts

Routed Mode Only

The FWSM lets you share an interface between contexts. However, packet classification requirements might make sharing interfaces impractical. Because the classifier relies on active NAT sessions to classify the destination addresses to a context, the classifier is limited by how you can configure NAT. If you do not want to perform NAT, you must use unique interfaces.

**Note**

The FWSM does not support sharing the outside interface of one context with the inside interface of another context (known as cascading contexts). Traffic that is outbound from one context (from a higher to a lower security interface) can only enter another context as inbound traffic (lower to higher security); it cannot be outbound for both contexts, or inbound for both contexts.

This section includes the following topics:

- [NAT and Origination of Traffic, page 9-7](#)
- [Sharing an Outside Interface, page 9-7](#)
- [Sharing an Inside Interface, page 9-7](#)

NAT and Origination of Traffic

The type of NAT configured determines whether the traffic can originate on the shared interface or if it can only respond to an existing connection. When you use dynamic NAT, you cannot initiate a connection to the real addresses. Therefore, traffic from the shared interface must be in response to an existing connection. Static NAT, however, lets you initiate connections, so you can initiate connections on the shared interface.

Sharing an Outside Interface

When you have an outside shared interface (connected to the Internet, for example), the destination addresses on the inside are limited, and are known by the system administrator, so configuring NAT for those addresses is easy, even if you want to configure static NAT.

Sharing an Inside Interface

Configuring an inside shared interface poses a problem, however, if you want to allow communication between the shared interface and the Internet, where the destination addresses are unlimited. For example, if you want to allow inside hosts on the shared interface to initiate traffic to the Internet, then you need to configure static NAT statements for each Internet address. This requirement necessarily limits the kind of Internet access you can provide for users on an inside shared interface. (If you intend to statically translate addresses for Internet servers, then you also need to consider DNS entry addresses and how NAT affects them. For example, if a server sends a packet to `www.example.com`, then the DNS server needs to return the translated address. Your NAT configuration determines DNS entry management.)

Figure 9-4 shows two servers on an inside shared interface. One server sends a packet to the translated address of a web server, and the FWSM classifies the packet to go through Context C because it includes a static translation for the address. The other server sends the packet to the real untranslated address, and the packet is dropped because the FWSM cannot classify it.

Figure 9-4 *Originating Traffic on a Shared Interface*



Enabling or Disabling Multiple Context Mode at the CLI

Your FWSM might already be configured for multiple security contexts depending on how you ordered it from Cisco. If you are upgrading, however, you might need to convert from single mode to multiple mode by following the procedures in this section. ASDM does not support changing modes, so you need to change modes using the CLI.

This section contains the following topics:

- [Backing Up the Single Mode Configuration, page 9-9](#)
- [Enabling Multiple Context Mode, page 9-9](#)
- [Restoring Single Context Mode, page 9-9](#)

Backing Up the Single Mode Configuration

When you convert from single mode to multiple mode, the FWSM converts the running configuration into two files. The original startup configuration is not saved, so if it differs from the running configuration, you should back it up before proceeding.

Enabling Multiple Context Mode

The context mode (single or multiple) is not stored in the configuration file, even though it does endure reboots. If you need to copy your configuration to another device, set the mode on the new device to match using the **mode** command.

When you convert from single mode to multiple mode, the FWSM converts the running configuration into two files: a new startup configuration that comprises the system configuration, and `admin.cfg` that comprises the admin context (in the root directory of the internal Flash memory). The original running configuration is saved as `old_running.cfg` (in the root directory of the internal Flash memory). The original startup configuration is not saved. The FWSM automatically adds an entry for the admin context to the system configuration with the name “admin.”

To enable multiple mode, enter the following command:

```
hostname(config)# mode multiple
```

You are prompted to reboot the FWSM.

Restoring Single Context Mode

If you convert from multiple mode to single mode, you might want to first copy a full startup configuration (if available) to the FWSM; the system configuration inherited from multiple mode is not a complete functioning configuration for a single mode device. For example, you can restore the old single-mode running configuration, if available, as the startup configuration. Because the system configuration does not have any network interfaces as part of its configuration, you must access the FWSM from a switch session to perform the copy.

To copy the old running configuration to the startup configuration and to change the mode to single mode, perform the following steps in the system execution space:

-
- Step 1** To copy the backup version of your original running configuration to the current startup configuration, enter the following command in the system execution space:

```
hostname(config)# copy old_running.cfg startup-config
```

- Step 2** To set the mode to single mode, enter the following command in the system execution space:

```
hostname(config)# mode single
```

The FWSM reboots.

Managing Memory for Rules

The FWSM supports a fixed number of rules for the entire system. In multiple context mode, the FWSM partitions the memory allocated to rule configuration, and assigns each context to a partition. This section describes how rule allocation works and how to manage memory partitions; it includes the following topics:

- [About Memory Partitions, page 9-10](#)
- [Default Rule Allocation, page 9-10](#)
- [Setting the Number of Memory Partitions, page 9-11](#)
- [Changing the Memory Partition Size, page 9-12](#)
- [Reallocating Rules Between Features for a Specific Memory Partition, page 9-15](#)

About Memory Partitions

In multiple context mode, the FWSM partitions the memory allocated to rule configuration, and assigns each context to a partition. By default, a context belongs to one of 12 partitions that offers a maximum number rules, including ACEs, AAA rules, and others. See the [“Default Rule Allocation”](#) section for a list of rule limits.

The FWSM assigns contexts to the partitions in the order they are loaded at startup. For example, if you have 12 contexts and the maximum number of rules is 14,103, each context is assigned to its own partition, and can use 14,103 rules. If you add one more context, then context number 1 and the new context number 13 are both assigned to partition 1, and can use 14,103 rules divided between them; the other 11 contexts continue to use 14,103 rules each. If you delete contexts, the partition membership does not shift, so you might have some unequal distribution until you reboot, at which time the contexts are evenly distributed.

**Note**

Rules are used up on a first come, first served basis, so one context might use more rules than another context.

You can manage memory partitions by manually assigning a context to a partition (see the [“Configuring Security Contexts”](#) section on page 9-23); reducing the number of partitions to better match the number of contexts you have (see the [“Setting the Number of Memory Partitions”](#) section on page 9-11); changing the size of a partition (see the [“Changing the Memory Partition Size”](#) section on page 9-12); and reallocating rules between features (see the [“Reallocating Rules Between Features for a Specific Memory Partition”](#) section on page 9-15).

Default Rule Allocation

[Table 9-1](#) lists the default number of rules for each feature type in multiple context mode, for the default 12 memory partitions.

**Note**

Some access lists use more memory than others. Depending on the type of access list, the actual limit the system can support will be less than the maximum. See the [“Maximum Number of ACEs”](#) section on page 12-5 for more information about ACEs and memory usage.

Table 9-1 Default Rule Allocation

Specification	Maximum per Partition (with 12 ¹ Partitions)
AAA Rules	1345
ACEs	14,801
established commands ²	96
Filter Rules	576
ICMP, Telnet, SSH, and HTTP Rules	384
Policy NAT ACEs ³	384
Inspect Rules	1537
Total Rules	19,219

1. Use the **show resource rule** command to view the default values for partitions other than 12.
2. Each **established** command creates a control and data rule, so this value is doubled in the Total Rules value.
3. This limit is lower than in release 2.3.

Setting the Number of Memory Partitions

This section describes how to set the number of partitions.

Guidelines

See the following guidelines for changing the number of partitions:



Caution

Failure to follow these guidelines might result in dropped access list configuration as well as other anomalies, including ACL tree corruption.

- The target partition and rule allocation settings must be carefully calculated, planned, and preferably tested in a non-production environment prior to making the change to ensure that all existing contexts and rules can be accommodated.
- When failover is used, both FWSMs need to be reloaded at the same time after making partition changes. Reloading both FWSMs causes an outage with no possibility for a zero-downtime reload. At no time should two FWSMs with a mismatched number of partitions or rule limits synchronize over failover.
- Changing the number of partitions requires you to reload the FWSM.
- When increasing the number of partitions, the default size of each partition is reduced.
- If you manually configured the partition sizes (see the [“Changing the Memory Partition Size” section on page 9-12](#)), the sizes you set might not be compatible with the new smaller partition sizes.
- If the current configured sizes do not fit into the new partitions, then the FWSM rejects the new memory partition configuration.
- The FWSM also checks the rule allocation (see the [“Reallocating Rules Between Features for a Specific Memory Partition” section on page 9-15](#)). If you manually allocated rules between features so that the total number of rules allocated is now greater than those available, then the FWSM rejects the new memory partition configuration. Similarly, if the absolute maximum number of rules for a feature is now exceeded, then the FWSM rejects the new memory partition configuration.

Detailed Steps

To change the number of memory partitions, perform the following steps:

- Step 1** (Optional) To view the current mapping of contexts to memory partitions, enter the following command from **Tools > Command Line Interface**:

```
show resource acl-partition
```

The following sample output shows that 2 memory partitions are configured:

```
Total number of configured partitions = 2
Partition #0
  Mode                :exclusive
  List of Contexts    :bandn, borders
  Number of contexts  :2(RefCount:2)
  Number of rules     :0(Max:53087)
Partition #1
  Mode                :non-exclusive
  List of Contexts    :admin, momandpopA, momandpopB, momandpopC
                    :momandpopD
  Number of contexts  :5(RefCount:5)
  Number of rules     :6(Max:53087)
```

For information about exclusive and non-exclusive partitions, see the [“Configuring Security Contexts” section on page 9-23](#).

- Step 2** In the System execution space, from the Configuration > Device Management > Resource Allocation > Local Rules pane, set the number of partitions in the Number of ACL Partitions field, between 1 and 12.



Note If you assign a context to a partition, the partition numbering starts with 0. So if you have 12 partitions, the partition numbers are 0 through 11. See the [“Configuring Security Contexts” section on page 9-23](#) to assign contexts to partitions.

- Step 3** Click **Apply**.

- Step 4** To reload the FWSM so the change can take effect, choose **Tools > System Reload**.

If you are using failover, wait a few seconds before reloading the standby unit as well; the standby unit does not reload automatically, and the memory partitions must match on both units. Traffic loss can occur because both units are down at the same time.



Note If you add a secondary unit at a later date, then after the new secondary unit synchronizes the configuration, immediately reload the secondary unit so that the memory partitions are the same. During the initial synchronization, the configuration might not fit properly in the secondary unit memory partitions, but after reloading, and another configuration synchronization, the secondary unit will be operational.

Changing the Memory Partition Size

The FWSM lets you set the memory size of each partition.

Guidelines

See the following guidelines before you change partition sizes:



Caution

Failure to follow these guidelines might result in dropped access list configuration as well as other anomalies, including ACL tree corruption.

- The target partition and rule allocation settings must be carefully calculated, planned, and preferably tested in a non-production environment prior to making the change to ensure that all existing contexts and rules can be accommodated.
- When failover is used, both FWSMs need to be reloaded at the same time after making partition changes. Reloading both FWSMs causes an outage with no possibility for a zero-downtime reload. At no time should two FWSMs with a mismatched number of partitions or rule limits synchronize over failover.
- Changing the partition sizes requires you to reload the FWSM.
- Change the number of partitions before you set the partition sizes; changing the number of partitions affects the overall number of rules per partition. If you increase the number of partitions, for example, then the number of rules available per partition will be smaller. Therefore, your partition size configuration might be invalid, and you might need to reconfigure all your partition sizes. Changing the number of partitions requires you to reload the FWSM before you change the partition sizes; then changing the partition sizes requires a second reload.
- Allocate contexts to specific partitions before you set the partition sizes (see the [“Configuring Security Contexts”](#) section on page 9-23). If you plan all your partition sizes based on the contexts currently assigned to a partition, but you did not specifically allocate the contexts, then you run the risk of context assignments shifting after a reload (for example if you add or subtract contexts).
- Reduce the size of partition(s) before increasing the size of other partition(s). The FWSM rejects any increases in size if there is not free space available.
- If the existing number of ACEs does not fit into the new partition size, then the resizing is rejected.
- In addition to the memory partitions to which the FWSM assigns contexts, the FWSM uses a backup tree partition to process changes to rules so traffic can continue to use the old configuration until the new configuration is ready. The backup tree must be as large as the largest partition. Therefore, some memory is automatically assigned to the backup tree in tandem with the largest partition; so be sure to include the backup tree in your calculations.
- If you reduce the size of a partition, the FWSM checks the rule allocation (see the [“Reallocating Rules Between Features for a Specific Memory Partition”](#) section on page 9-15). If you manually allocated rules between features so that the total number of rules allocated is now greater than those available, then the FWSM rejects the resizing of the partition. Similarly, if the absolute maximum number of rules for a feature is now exceeded, then the FWSM rejects the resizing of the partition.

Detailed Steps

To set the size of the memory partitions, perform the following steps:

- Step 1** From the System execution space, go to the Configuration > Device Management > Resource Allocation > Local Rules pane.

The Current Configured Size column shows the number of rules that currently assigned to each partition. The backup tree partition (not configurable) is used to process changes to rules so traffic can continue to use the old configuration until the new configuration is ready.

The total number of rules available displays at the bottom of the table.

- Step 2** Calculate how you want to resize the partitions. Note that the backup tree must be as large as the largest partition. Therefore, some memory is automatically assigned to the backup tree in tandem with the largest partition; so be sure to include the backup tree in your calculations.

For example, if you reduce the sizes of partitions 0 through 5 to 15,000, then after applying your change, ASDM shows that you have 25,314 rules to reallocate to other partitions.

Partition	Bootup Partition Size	Configured Size	Difference
0	19219	15000	4219
1	19219	15000	4219
2	19219	15000	4219
3	19219	15000	4219
4	19219	15000	4219
5	19219	15000	4219
			Total Reduced: 25314

If you want to distribute the rules evenly across the other 6 partitions plus the backup tree, then you can add 3616 rules to each (with 2 left over). Remember that the backup tree must be as large as the largest partition, so you must consider the backup tree in your calculations. For example, if you want to make partition 6 have 24,001 rules, then you can allocate the rules like this:

Partition	Bootup Partition Size	Configured Size	Difference
6	19219	24001	4782
Backup Tree	19219	24001	4782
7	19219	22369	3150
8	19219	22369	3150
9	19219	22369	3150
10	19219	22369	3150
11	19219	22369	3150
			Total Increased: 25314

- Step 3** You must reduce the size of partitions before you can increase any sizes. For a partition you want to reduce, double-click the partition row, or click **Edit**.

The Resource Rule dialog box appears.

- Step 4** In the Partition Size field, enter a new number.

By default, the rule allocations are adjusted to the default settings for the new size.

- a. If you previously reallocated rules to be specific values (as opposed to “default”) then you must reduce the total number of rules to fit the new partition size either by entering new numbers or by choosing “default” for all features. See the [“Reallocating Rules Between Features for a Specific Memory Partition” section on page 9-15](#) for more information.
- b. Click **OK**.

- Step 5** To reduce the size of other partitions, repeat Steps 3 and 4.
- Step 6** Click **Apply** to apply all size reductions.
- Step 7** To increase a partition, double-click the partition row, or click **Edit**.
The Resource Rule dialog box appears.
- Step 8** In the Partition Size field, enter a new number, and click **OK**.
By default, the rule allocations are adjusted to the default settings for the new size.
You cannot increase the total number of rules per feature until after you reload the FWSM. After you reload, you can change the rule allocation according to the “[Reallocating Rules Between Features for a Specific Memory Partition](#)” section on page 9-15.
- Step 9** To increase the size of other partitions, repeat Steps 7 and 8.
- Step 10** Click **Apply** to apply all size increases.
- Step 11** To reload the FWSM so the change can take effect, choose **Tools > System Reload**.
If you are using failover, wait a few seconds before reloading the standby unit as well; the standby unit does not reload automatically, and the memory partitions must match on both units. Traffic loss can occur because both units are down at the same time.



Note If you add a secondary unit at a later date, then after the new secondary unit synchronizes the configuration, immediately reload the secondary unit so that the memory partitions are the same. During the initial synchronization, the configuration might not fit properly in the secondary unit memory partitions, but after reloading, and another configuration synchronization, the secondary unit will be operational.

Reallocating Rules Between Features for a Specific Memory Partition

To set the rule allocation globally for all partitions, see the “[Reallocating Rules Between Features](#)” section on page A-8. Setting the rule allocation for a specific partition overrides the global setting. For information about setting the partition size, see the “[Changing the Memory Partition Size](#)” section on page 9-12.

Guidelines



Caution

Failure to follow these guidelines might result in dropped access list configuration as well as other anomalies, including ACL tree corruption.

- The target partition and rule allocation settings must be carefully calculated, planned, and preferably tested in a non-production environment prior to making the change to ensure that all existing contexts and rules can be accommodated.
- When failover is used, both FWSMs need to be reloaded at the same time after making partition changes. Reloading both FWSMs causes an outage with no possibility for a zero-downtime reload. At no time should two FWSMs with a mismatched number of partitions or rule limits synchronize over failover.

Detailed Steps

To reallocate rules for a given partition, perform the following steps:

- Step 1** To view the total number of rules available per partition, the default values, current rule allocation, and the absolute maximum number of rules you can allocate per feature, enter the following command using **Tools > Command Line Interface**:

```
show resource rule partition [number]
```

For example, the following display shows the maximum rules as 19219 for partition 0 (this is an example only, and might differ from the actual number of rules for your system):

```
show resource rule partition 0
```

Result of the command: "show resource rule partition 0"

CLS Rule	Default Limit	Configured Limit	Absolute Max
Policy NAT	384	384	833
ACL	14801	14801	14801
Filter	576	576	1152
Fixup	1537	1537	3074
Est Ctl	96	96	96
Est Data	96	96	96
AAA	1345	1345	2690
Console	384	384	768
Total	19219	19219	

```
Partition Limit - Configured Limit = Available to allocate
19219 - 19219 = 0
```



Note If you increase the size of a partition but have not yet reloaded, the maximum number of rules remains at the old smaller size. You have to reload to see the increased limits. If you decrease the size of a partition but have not yet reloaded, the new smaller number of rules is reflected right away.

- Step 2** To view the number of rules currently being used so you can plan your reallocation, enter the following command at the Command Line Interface tool.

```
show np 3 acl count partition_number
```

For example, the following is sample output from the **show np 3 acl count 0** command, and shows the number of inspections (Fixup Rule) close to the maximum of 9216. You might choose to reallocate some access list rules (ACL Rule) to inspections.

```
show np 3 acl count 0
```

Result of the command: "show np 3 acl count 0"

```
----- CLS Rule Current Counts -----
CLS Filter Rule Count      :          0
CLS Fixup Rule Count       :        9001
CLS Est Ctl Rule Count     :           4
CLS AAA Rule Count         :           15
CLS Est Data Rule Count    :           4
CLS Console Rule Count     :          16
```

```

CLS Policy NAT Rule Count      :           0
CLS ACL Rule Count             :        30500
CLS ACL Uncommitted Add       :           0
CLS ACL Uncommitted Del       :           0
...

```



Note The **established** command creates two types of rules, control and data. Both of these types are shown in the display, but you allocate both rules by setting the number of **established** commands; you do not set each rule separately.

Step 3 To reallocate rules between features, go to System > Configuration > Device Management > Resource Allocation > Local Rules, double-click the partition row you want to edit, or click **Edit**.

The Resource Rule dialog box appears.

Step 4 For each rule type, enter a new number or choose **default** or **max** from the drop-down list.

The default and max options set the number of rules to the default or maximum setting. The default and max numbers depend on the number of partitions you set, and each partition size. If you choose the max option, be sure to set other features lower to accommodate this value. You cannot leave the other feature values at default if you set one or more features to max; you need to manually decrease the value of one or more features to accommodate the increase in the feature you set to max.

Step 5 Click **Apply**.

These settings take effect immediately when you Apply; you do not need to reload the FWSM.

Configuring Resource Classes

By default, all security contexts have unlimited access to the resources of the FWSM, except where maximum limits per context are enforced. However, if you find that one or more contexts use too many resources, and they cause other contexts to be denied connections, for example, then you can configure resource management to limit the use of resources per context. The FWSM manages resources by assigning contexts to resource classes. Each context uses the resource limits set by the class.



Note The FWSM does not limit the bandwidth per context; however, the switch containing the FWSM can limit bandwidth per VLAN. See the switch documentation for more information.

This section contains the following topics:

- [Classes and Class Members Overview, page 9-17](#)
- [Adding a Resource Class, page 9-20](#)

Classes and Class Members Overview

The FWSM manages resources by assigning contexts to resource classes. Each context uses the resource limits set by the class. This section includes the following topics:

- [Resource Limits Overview, page 9-18](#)
- [Default Class Overview, page 9-19](#)

- [Class Members Overview, page 9-20](#)

Resource Limits Overview

When you create a class, the FWSM does not set aside a portion of the resources for each context assigned to the class; rather, the FWSM sets the maximum limit for a context. If you oversubscribe resources, or allow some resources to be unlimited, a few contexts can “use up” those resources, potentially affecting service to other contexts.

You can set the limit for all resources together as a percentage of the total available for the device. Also, you can set the limit for individual resources as a percentage or as an absolute value.

You can oversubscribe the FWSM by assigning more than 100 percent of the resources across all contexts. For example, you can set the Bronze class to limit connections to 20 percent per context, and then assign 10 contexts to the class for a total of 200 percent. If contexts concurrently use more than the system limit, then each context gets less than the 20 percent you intended. (See the following figure.)



The FWSM lets you assign unlimited access to one or more resources in a class, instead of a percentage or absolute number. When a resource is unlimited, contexts can use as much of the resource as the system has available. For example, Context A, B, and C are in the Silver Class, which limits each class member to 1 percent of the system inspections per second, for a total of 3 percent; but the three contexts are currently only using 2 percent combined. Gold Class has unlimited access to inspections. The contexts in the Gold Class can use more than the 97 percent of “unassigned” inspections; they can also use the 1 percent of inspections not currently in use by Context A, B, and C, even if that means that Context A, B, and C are unable to reach their 3 percent combined limit. (See the following figure.) Setting unlimited access is similar to oversubscribing the FWSM, except that you have less control over how much you oversubscribe the system.



Default Class Overview

All contexts belong to the default class if they are not assigned to another class; you do not have to actively assign a context to the default class.

If a context belongs to a class other than the default class, those class settings always override the default class settings. However, if the other class has any settings that are not defined, then the member context uses the default class for those limits. For example, if you create a class with a 2 percent limit for all concurrent connections, but no other limits, then all other limits are inherited from the default class. Conversely, if you create a class with a 2 percent limit for *all* resources, the class uses no settings from the default class.

By default, the default class provides unlimited access to resources for all contexts, except for the following limits, which are by default set to the maximum allowed per context:

- Telnet sessions—5 sessions.
- SSH sessions—5 sessions.
- IPSec sessions—5 sessions.
- MAC addresses—65,535 entries.

The following figure shows the relationship between the default class and other classes. Contexts A and C belong to classes with some limits set; other limits are inherited from the default class. Context B inherits no limits from default because all limits are set in its class, the Gold class. Context D was not assigned to a class, and is by default a member of the default class.



Class Members Overview

To use the settings of a class, assign the context to the class when you define the context. All contexts belong to the default class if they are not assigned to another class; you do not have to actively assign a context to default. You can only assign a context to one resource class. The exception to this rule is that limits that are undefined in the member class are inherited from the default class; so in effect, a context could be a member of default plus another class.

Adding a Resource Class

To add a resource class, perform the following steps:

-
- Step 1** If you are not already in the System configuration mode, in the Device List pane, double-click **System** under the active device IP address.
 - Step 2** On the Context Management > Resource Class pane, click **Add**.
The Add Resource Class dialog box appears.
 - Step 3** In the Resource Class field, enter a class name up to 20 characters in length.
 - Step 4** In the Count Limited Resources area, set the concurrent limits for resources.

For resources that do not have a system limit, you cannot set the percentage; you can only set an absolute value. If you do not set a limit, the limit is inherited from the default class. If the default class does not set a limit, then the resource is unlimited, or the system limit if available.

You can set one or more of the following limits:

- **Hosts**—Sets the limit for concurrent hosts that can connect through the FWSM. Select the check box to enable this limit. If you set the limit to 0, it is unlimited.
- **Telnet**—Sets the limit for concurrent Telnet sessions. Select the check box to enable this limit. You can set the limit as a percentage by entering any integer greater than 1 and selecting **Percent** from the list. You can assign more than 100 percent if you want to oversubscribe the device. Or you can set the limit as an absolute value by entering an integer between 1 and 5 and selecting **Absolute** from the list. The system has a maximum of 100 sessions divided between all contexts.
- **ASDM Sessions**—Sets the limit for concurrent ASDM sessions. Select the check box to enable this limit. You can set the limit as a percentage by entering any integer greater than 1 and selecting **Percent** from the list. You can assign more than 100 percent if you want to oversubscribe the device. Or you can set the limit as an absolute value by entering an integer between 1 and 5 and selecting **Absolute** from the list. The system has a maximum of 80 sessions divided between all contexts. ASDM sessions use two HTTPS connections: one for monitoring that is always present, and one for making configuration changes that is present only when you make changes. For example, the system limit of 32 ASDM sessions represents a limit of 64 HTTPS sessions, divided between all contexts.
- **Connections**—Sets the limit for concurrent TCP or UDP connections between any two hosts, including connections between one host and multiple other hosts. Select the check box to enable this limit. You can set the limit as a percentage by entering any integer greater than 1 and selecting **Percent** from the list. You can assign more than 100 percent if you want to oversubscribe the device. Or you can set the limit as an absolute value by entering an integer between 0 (system limit) and the system limit for your model, and selecting **Absolute** from the list. See the *Cisco ASDM Release Notes* for the connection limit for your model.
- **Xlates**—Sets the limit for address translations. Select the check box to enable this limit. If you set the limit to 0, it is unlimited.
- **SSH**—Sets the limit for SSH sessions. Select the check box to enable this limit. You can set the limit as a percentage by entering any integer greater than 1 and selecting **Percent** from the list. You can assign more than 100 percent if you want to oversubscribe the device. Or you can set the limit as an absolute value by entering an integer between 1 and 5 and selecting **Absolute** from the list. The system has a maximum of 100 sessions divided between all contexts.
- **MAC Entries**—(Transparent mode only) Sets the limit for MAC address entries in the MAC address table. Select the check box to enable this limit. You can set the limit as a percentage by entering any integer greater than 1 and selecting **Percent** from the list. You can assign more than 100 percent if you want to oversubscribe the device. Or you can set the limit as an absolute value by entering an integer between 0 (system limit) and 65535 and selecting **Absolute** from the list.

Step 5 In the Rate Limited Resources area, set the rate limit for resources.

If you do not set a limit, the limit is inherited from the default class. If the default class does not set a limit, then it is unlimited by default.

You can set one or more of the following limits:

- **Conns/sec**—Sets the limit for connections per second. Select the check box to enable this limit. If you set the limit to 0, it is unlimited.
- **Syslogs/sec**—Sets the limit for system log messages per second. Select the check box to enable this limit. If you set the limit to 0, it is unlimited.
- **Inspects/sec**—Sets the limit for application inspections per second. Select the check box to enable this limit. If you set the limit to 0, it is unlimited.

Step 6 Click **OK**.

Monitoring Context Resource Usage

To monitor resource usage of all contexts from the system execution space, perform the following steps:

-
- Step 1** If you are not already in the System mode, in the Device List pane, double-click **System** under the active device IP address.
- Step 2** Click the **Monitoring** button on the toolbar.
- Step 3** Click **Context Resource Usage**.

Click each resource type to view the resource usage for all contexts:

- **ASDM**—Shows the usage of ASDM connections.
 - Context—Shows the name of each context.
 - Existing Connections (#)—Shows the number of existing connections.
 - Existing Connections (%)—Shows the connections used by this context as a percentage of the total number of connections used by all contexts.
 - Peak Connections (#)—Shows the peak number of connections since the statistics were last cleared, either using the **clear resource usage** command or because the device rebooted.
- **Telnet**—Shows the usage of Telnet connections.
 - Context—Shows the name of each context.
 - Existing Connections (#)—Shows the number of existing connections.
 - Existing Connections (%)—Shows the connections used by this context as a percentage of the total number of connections used by all contexts.
 - Peak Connections (#)—Shows the peak number of connections since the statistics were last cleared, either using the **clear resource usage** command or because the device rebooted.
- **SSH**—Shows the usage of SSH connections.
 - Context—Shows the name of each context.
 - Existing Connections (#)—Shows the number of existing connections.
 - Existing Connections (%)—Shows the connections used by this context as a percentage of the total number of connections used by all contexts.
 - Peak Connections (#)—Shows the peak number of connections since the statistics were last cleared, either using the **clear resource usage** command or because the device rebooted.
- **Xlates**—Shows the usage of network address translations.
 - Context—Shows the name of each context.
 - Xlates (#)—Shows the number of current xlates.
 - Xlates (%)—Shows the xlates used by this context as a percentage of the total number of xlates used by all contexts.
 - Peak (#)—Shows the peak number of xlates since the statistics were last cleared, either using the **clear resource usage** command or because the device rebooted.
- **NATs**—Shows the number of NAT rules.
 - Context—Shows the name of each context.
 - NATs (#)—Shows the current number of NAT rules.

- NATs (%)—Shows the NAT rules used by this context as a percentage of the total number of NAT rules used by all contexts.
- Peak NATs (#)—Shows the peak number of NAT rules since the statistics were last cleared, either using the **clear resource usage** command or because the device rebooted.
- **Syslogs**—Shows the rate of system log messages.
 - Context—Shows the name of each context.
 - Syslog Rate (#/sec)—Shows the current rate of system log messages.
 - Syslog Rate (%)—Shows the system log messages generated by this context as a percentage of the total number of system log messages generated by all contexts.
 - Peak Syslog Rate (#/sec)—Shows the peak rate of system log messages since the statistics were last cleared, either using the **clear resource usage** command or because the device rebooted.

Step 4 Click **Refresh** to refresh the view.

Configuring Security Contexts

To add a security context, perform the following steps:

-
- Step 1** If you are not already in the System configuration mode, in the Device List pane, double-click **System** under the active device IP address.
- Step 2** On the Context Management > Security Contexts pane, click **Add**.
The Add Context dialog box appears.
- Step 3** In the Security Context field, enter the context name as a string up to 32 characters long.
This name is case sensitive, so you can have two contexts named “customerA” and “CustomerA,” for example. “System” or “Null” (in upper or lower case letters) are reserved names, and cannot be used.
- Step 4** In the Interface Allocation area, click the **Add** button to assign an interface to the context.
- Step 5** From the Interfaces > Physical Interface drop-down list, choose an interface.
You can assign the main interface, in which case you leave the subinterface ID blank, or you can assign a subinterface or a range of subinterfaces associated with this interface. In transparent firewall mode, only interfaces that have not been allocated to other contexts are shown. If the main interface was already assigned to another context, then you must choose a subinterface.
- Step 6** (Optional) In the Interfaces > Subinterface Range (optional) drop-down list, choose a subinterface ID.
For a range of subinterface IDs, choose the ending ID in the second drop-down list, if available.
In transparent firewall mode, only subinterfaces that have not been allocated to other contexts are shown.
- Step 7** (Optional) In the Aliased Names area, check **Use Aliased Name in Context** to set an aliased name for this interface to be used in the context configuration instead of the interface ID.
- a. In the Name field, sets the aliased name.
An aliased name must start with a letter, end with a letter, and have as interior characters only letters, digits, or an underscore. This field lets you specify a name that ends with a letter or underscore; to add an optional digit after the name, set the digit in the Range field.
 - b. (Optional) In the Range field, set the numeric suffix for the aliased name.

If you have a range of subinterfaces, you can enter a range of digits to be appended to the name.

- Step 8** (Optional) To enable context users to see physical interface properties even if you set an aliased name, check **Show Hardware Properties in Context**.
- Step 9** Click **OK** to return to the Add Context dialog box.
- Step 10** (Optional) To assign this context to a resource class, choose a class name from the Resource Assignment > Resource Class drop-down list.

You can add or edit a resource class directly from this area. See the [“Configuring Resource Classes” section on page 9-17](#) for more information.

- Step 11** (Optional) To map a context to a specific memory partition, choose a partition number from the Resource Assignment > ACL Partition drop-down list.

See the [“Setting the Number of Memory Partitions” section on page 9-11](#) to configure the number of memory partitions.

When you assign a context to a partition, then the partition becomes *exclusive*. An exclusive partition only includes contexts that you specifically assign to it. Partitions that do not have contexts specifically assigned to them are non-exclusive and contexts are allocated to them in a round-robin fashion.



Note If you assign contexts to all partitions, then they are all exclusive. However, if you later add a context that is not assigned to a partition, then contexts are allocated to exclusive partitions in a round-robin fashion, and the first best-fit exclusive partition available is used for the allocation of the new context. However, if none of the exclusive partitions can accommodate the rules of the new context, then it is assigned to partition 0 by default, even though partition 0 also cannot accommodate the context rules. The context rules will not load completely, so you need to manually adjust the way contexts are assigned to make room.

- Step 12** To set the context configuration location, identify the URL by choosing a file system type from the Config URL drop-down list and entering a path in the field.

For example, the combined URL for FTP has the following format:

```
ftp://server.example.com/configs/admin.cfg
```

- Step 13** (Optional) For external filesystems, set the username and password by clicking **Login**.
- Step 14** (Optional) To set the failover group for active/active failover, choose the group name in the Failover Group drop-down list.
- Step 15** (Optional) Add a description in the Description field.
-

