ASA/PIX: Configure Active/Standby Failover in Transparent Mode

Contents

Introduction

The failover configuration requires two identical security appliances connected to each other through a dedicated failover link and, optionally, a stateful failover link. The health of the active interfaces and units is monitored to determine if specific failover conditions are met. If those conditions are met, failover occurs.

The security appliance supports two failover configurations:

- Active/Active Failover
- Active/Standby Failover

Each failover configuration has its own method to determine and perform failover. With Active/Active Failover, both units can pass network traffic. This lets you configure load balancing on your network. Active/Active Failover is only available on units that run in multiple context mode. With Active/Standby Failover, only one unit passes traffic while the other unit waits in a standby state. Active/Standby Failover is available on units that run in either single or multiple context mode. Both failover configurations support stateful or stateless (regular) failover.

A transparent firewall, is a Layer 2 firewall that acts like a *bump in the wire*, or a *stealth firewall*, and is not seen as a router hop to connected devices. The security appliance connects the same network on its inside and outside ports. Because the firewall is not a routed hop, you can easily introduce a transparent firewall into an existing network; it is unnecessary to readdress IP. You can set the adaptive security appliance to run in the default routed firewall mode or transparent firewall mode. When you change modes, the adaptive security appliance clears the configuration because many commands are not supported in both modes. If you already have a populated configuration, be sure to back up this configuration before you change the mode; you can use this backup configuration for reference when you create a new configuration. Refer to Transparent Firewall Configuration Example for more information on the configuration of the firewall appliance in Transparent mode.

This document focuses on how to configure an Active/Standby Failover in Transparent Mode on the ASA Security Appliance.

Note: VPN failover is not supported on units that run in multiple context mode. VPN failover is available for **Active/Standby Failover** configurations only.

Cisco recommends that you do not use the management interface for failover, especially for stateful failover in which the security appliance constantly sends the connection information from one security appliance to the other. The interface for failover must be at least of the same capacity

as the interfaces that pass regular traffic, and while the interfaces on the ASA 5540 are gigabit, the management interface is FastEthernet only. The management interface is designed for management traffic only and is specified as management0/0. But, you can use the **management-only** command in order to configure any interface to be a management-only interface. Also, for Management 0/0, you can disable management-only mode so the interface can pass through traffic just like any other interface. Refer to Cisco Security Appliance Command Reference, Version 8.0 for more information about the **management-only** command.

This configuration guide provides a sample configuration to include a brief introduction to the PIX/ASA 7.x Active/Standby technology. Refer to the ASA/PIX Command Reference Guide for a more in-depth sense of the theory based behind this technology.

Prerequisites

Requirements

Hardware Requirement

The two units in a failover configuration must have the same hardware configuration. They must be the same model, have the same number and types of interfaces, and the same amount of RAM.

Note: The two units do not need to have the same size Flash memory. If you use units with different Flash memory sizes in your failover configuration, make sure the unit with the smaller Flash memory has enough space to accommodate the software image files and the configuration files. If it does not, configuration synchronization from the unit with the larger Flash memory to the unit with the smaller Flash memory fails.

Software Requirement

The two units in a failover configuration must be in the operational modes (routed or transparent, single or multiple context). They must have the same major (first number) and minor (second number) software version, but you can use different versions of the software within an upgrade process; for example, you can upgrade one unit from Version 7.0(1) to Version 7.0(2) and have failover remain active. Cisco recommends that you upgrade both units to the same version to ensure long-term compatibility.

Refer to the Performing Zero Downtime Upgrades for Failover Pairs section of *Cisco Security Appliance Command Line Configuration Guide, Version 8.0* for more information on how to upgrade the software on a failover pair.

License Requirements

On the ASA security appliance platform, at least one of the units must have an **unrestricted (UR) license**.

Note: It might be necessary to upgrade the licenses on a failover pair in order to obtain additional features and benefits. Refer to License Key Upgrade on a Failover Pair for more information.

Note: The licensed features (such as SSL VPN peers or security contexts) on both security appliances that participate in failover must be identical.

Components Used

The information in this document is based on these software and hardware versions:

- ASA Security Appliance with 7.x version and later

The information in this document was created from the devices in a specific lab environment. All of the devices used in this document started with a cleared (default) configuration. If your network is live, make sure that you understand the potential impact of any command.

Related Products

This configuration can also be used with these hardware and software versions:

- PIX Security Appliance with 7.x version and later

Conventions

Refer to the Cisco Technical Tips Conventions for more information on document conventions.

Active/Standby Failover

This section describes Active/Standby Failover and includes these topics:

Active/Standby Failover Overview

Active/Standby Failover lets you use a standby security appliance to take over the functionality of a failed unit. When the active unit fails, it changes to the standby state while the standby unit changes to the active state. The unit that becomes active assumes the IP addresses, or, for a transparent firewall, the management IP address, and MAC addresses of the failed unit and begins to pass traffic. The unit that is now in standby state takes over the standby IP addresses and MAC addresses. Because network devices see no change in the MAC to IP address pairing, no ARP entries change or time out anywhere on the network.

Note: For multiple context mode, the security appliance can fail over the entire unit, which includes all contexts, but cannot fail over individual contexts separately.

Primary/Secondary Status and Active/Standby Status

The main differences between the two units in a failover pair are related to which unit is active and which unit is standby, namely which IP addresses to use and which unit is primary and actively passes traffic.

A few differences exist between the units based on which unit is primary, as specified in the configuration, and which unit is secondary:

- The primary unit always becomes the active unit if both units start up at the same time (and are of equal operational health).
- The primary unit MAC address is always coupled with the active IP addresses. The exception to this rule occurs when the secondary unit is active and cannot obtain the primary MAC address over the failover link. In this case, the secondary MAC address is used.

Device Initialization and Configuration Synchronization

Configuration synchronization occurs when one or both devices in the failover pair boot. Configurations are always synchronized from the active unit to the standby unit. When the standby unit completes its initial startup, it clears its running configuration, except for the failover commands that are needed to communicate with the active unit, and the active unit sends its entire configuration to the standby unit.

The active unit is determined by these:

- If a unit boots and detects a peer already operative as active, it becomes the standby unit.
- If a unit boots and does not detect a peer, it becomes the active unit.
- If both units boot simultaneously, the primary unit becomes the active unit, and the secondary unit becomes the standby unit.

Note: If the secondary unit boots and does not detect the primary unit, it becomes the active unit. It uses its own MAC addresses for the active IP addresses. When the primary unit becomes available, the secondary unit changes the MAC addresses to those of the primary unit, which can cause an interruption in your network traffic. In order to avoid this, configure the failover pair with virtual MAC addresses. See the Configuring Active/Standby Failover section of this document for more information.

When the replication starts, the security appliance console on the active unit displays the message `Beginning configuration replication: Sending to mate`, and, when it is complete, the security appliance displays the message `End Configuration Replication to mate.` Within replication, commands entered on the active unit cannot replicate properly to the standby unit, and commands entered on the standby unit can be overwritten by the configuration that is replicated from the active unit. Do not enter commands on either unit in the failover pair within the configuration replication process. Dependent upon the size of the configuration, replication can take from a few seconds to several minutes.

From the secondary unit, you can observe the replication message as it synchronizes from the primary unit:

```
ASA> .

        Detected an Active mate
Beginning configuration replication from mate.
End configuration replication from mate.

ASA>
```

On the standby unit, the configuration exists only in running memory. In order to save the configuration to Flash memory after synchronization, enter these commands:

- For single context mode, enter the **copy running-config startup-config** command on the active unit. The command is replicated to the standby unit, which proceeds to write its configuration to Flash memory.
- For multiple context mode, enter the **copy running-config startup-config** command on the active unit from the system execution space and from within each context on disk. The command is replicated to the standby unit, which proceeds to write its configuration to Flash memory. Contexts with startup configurations on external servers are accessible from either unit over the network and do not need to be saved separately for each unit. Alternatively, you can copy the contexts on disk from the active unit to an external server, and then copy them to disk on the standby unit, where they become available when the unit reloads.

Command Replication

Command replication always flows from the active unit to the standby unit. As commands are entered on the active unit, they are sent across the failover link to the standby unit. You do not have to save the active configuration to Flash memory to replicate the commands.

Note: Changes made on the standby unit are not replicated to the active unit. If you enter a command on the standby unit, the security appliance displays the message `**** WARNING **** Configuration Replication is NOT performed from Standby unit to Active unit.` Configurations are no longer synchronized. This message is displayed even if you enter commands that do not affect the configuration.

If you enter the **write standby** command on the active unit, the standby unit clears its running configuration, except for the failover commands used to communicate with the active unit, and the active unit sends its entire configuration to the standby unit.

For multiple context mode, when you enter the **write standby** command in the system execution space, all contexts are replicated. If you enter the write standby command within a context, the command replicates only the context configuration.

Replicated commands are stored in the running configuration. In order to save the replicated commands to the Flash memory on the standby unit, enter these commands:

- For single context mode, enter the **copy running-config startup-config** command on the active unit. The command is replicated to the standby unit, which proceeds to write its configuration to Flash memory.
- For multiple context mode, enter the **copy running-config startup-config** command on the active unit from the system execution space and within each context on disk. The command is replicated to the standby unit, which proceeds to write its configuration to Flash memory. Contexts with startup configurations on external servers are accessible from either unit over the network and do not need to be saved separately for each unit. Alternatively, you can copy the contexts on disk from the active unit to an external server, and then copy them to disk on the standby unit.

Failover Triggers

The unit can fail if one of these events occurs:

- The unit has a hardware failure or a power failure.

- The unit has a software failure.
- Too many monitored interfaces fail.
- The **no failover active** command is entered on the active unit, or the **failover active** command is entered on the standby unit.

Failover Actions

In Active/Standby Failover, failover occurs on a unit basis. Even on systems that run in multiple context mode, you cannot failover individual or groups of contexts.

This table shows the failover action for each failure event. For each failure event, the table shows the failover policy (failover or no failover), the action taken by the active unit, the action taken by the standby unit, and any special notes about the failover condition and actions. The table shows the failover behavior.

Failure Event	Policy	Active Action	Standby Action	Notes
Active unit failed (power or hardware)	Failover	n/a	Become active; mark active as failed	No hello messages are received on any monitored interface or the failover link.
Formerly active unit recovers	No failover	Become standby	No action	None
Standby unit failed (power or hardware)	No failover	Mark standby as failed	n/a	When the standby unit is marked as failed, the active unit does not attempt to failover, even if the interface failure threshold is surpassed.
Failover link failed within operation	No failover	Mark failover interface as failed	Mark failover interface as failed	You must restore the failover link as soon as possible because the unit cannot failover to the standby unit while the failover link is down.
Failover link failed at startup	No failover	Mark failover interface as failed	Become active	If the failover link is down at startup, both units become active.

Stateful failover link failed	No failover	No action	No action	State information becomes out of date, and sessions are terminated if a failover occurs.
Interface failure on active unit above threshold	Failover	Mark active as failed	Become active	None
Interface failure on standby unit above threshold	No failover	No action	Mark standby as failed	When the standby unit is marked as failed, the active unit does not attempt to fail over even if the interface failure threshold is surpassed.

Regular and Stateful Failover

The security appliance supports two types of failover, regular and stateful. This section includes these topics:

Regular Failover

When a failover occurs, all active connections are dropped. Clients need to reestablish connections when the new active unit takes over.

Stateful Failover

When stateful failover is enabled, the active unit continually passes per-connection state information to the standby unit. After a failover occurs, the same connection information is available at the new active unit. Supported end-user applications are not required to reconnect to keep the same communication session.

The state information passed to the standby unit includes these:

- The NAT translation table
- The TCP connection states
- The UDP connection states
- The ARP table
- The Layer 2 bridge table (only when the Firewall runs in the **transparent firewall** mode)
- The HTTP connection states (if HTTP replication is enabled)
- The ISAKMP and IPSec SA table
- The GTP PDP connection database

The information that is not passed to the standby unit when stateful failover is enabled includes these:

- The HTTP connection table (unless HTTP replication is enabled)
- The user authentication (uauth) table
- The routing tables
- State information for security service modules

Note: If failover occurs within an active Cisco IP SoftPhone session, the call remains active because the call session state information is replicated to the standby unit. When the call is terminated, the IP SoftPhone client loses connection with the Cisco CallManager . This occurs because there is no session information for the CTIQBE hang-up message on the standby unit. When the IP SoftPhone client does not receive a response back from the Cisco CallManager within a certain time period, it considers the Cisco CallManager unreachable and unregisters itself.

LAN-Based Active/Standby Failover Configuration

Network Diagram

This document uses this network setup:

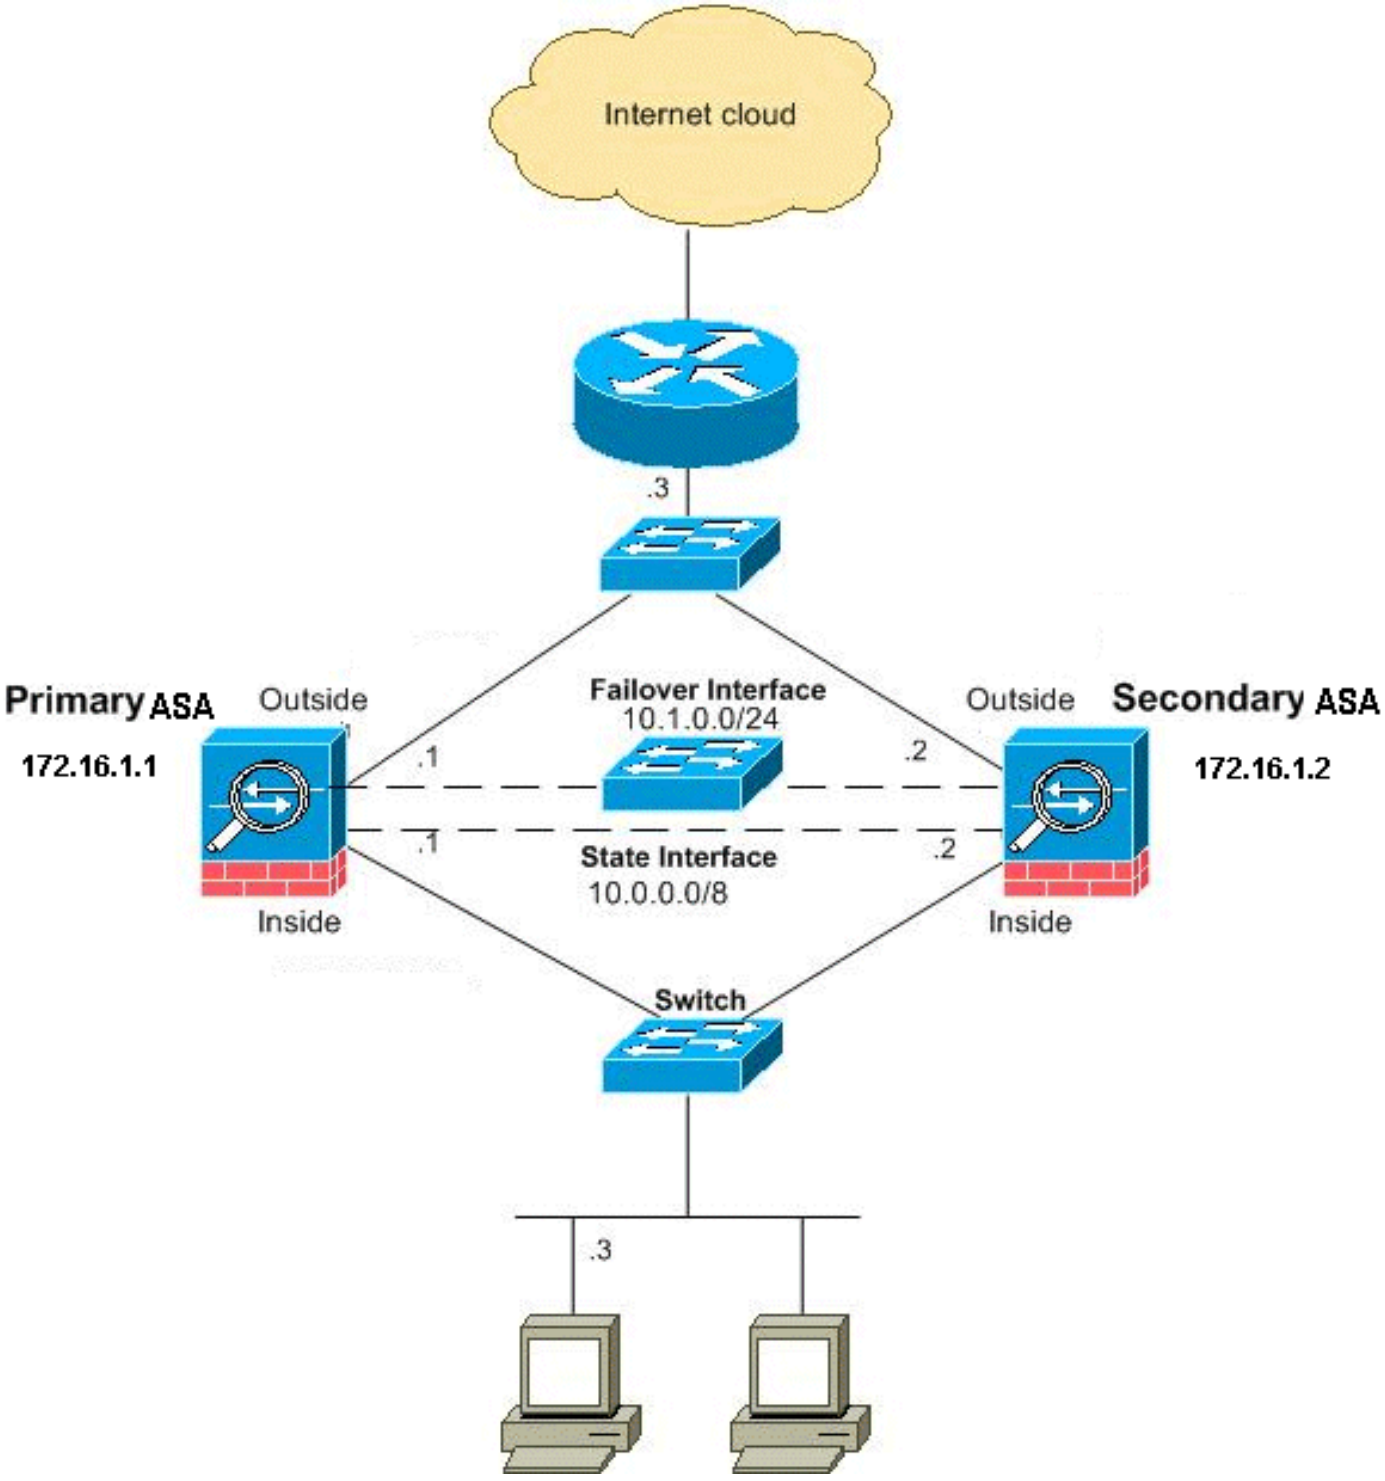

This section describes how to configure Active/Standby Failover in Transparent mode with an Ethernet failover link. When you configure LAN-based failover, you must bootstrap the secondary device to recognize the failover link before the secondary device can obtain the running configuration from the primary device.

Note: If you change from cable-based failover to LAN-based failover, you can skip many steps, such as the assignment of the active and standby IP addresses for each interface, that you completed for the cable-based failover configuration.

Primary Unit Configuration

Complete these steps in order to configure the primary unit in a LAN-based, Active/Standby Failover configuration. These steps provide the minimum configuration needed to enable failover

on the primary unit. For multiple context mode, all steps are performed in the system execution space unless otherwise noted.

In order to configure the primary unit in an Active/Standby Failover pair, complete these steps:

1. If you have not done so already, configure the active and standby IP addresses for the management interface (transparent mode). The standby IP address is used on the security appliance that is currently the standby unit. It must be in the same subnet as the active IP address.**Note:** Do not configure an IP address for the stateful failover link if you use a dedicated stateful failover interface. You use the **failover interface ip** command to configure a dedicated stateful failover interface in a later step.`hostname(config-if)#` **`ip address active_addr netmask standby standby_addr`** Unlike routed mode, which requires an IP address for each interface, a transparent firewall has an IP address assigned to the entire device. The security appliance uses this IP address as the source address for packets that originate on the security appliance, such as system messages or AAA communications. In the example, the IP address for the primary ASA is configured as shown below:`hostname(config)#` **`ip address 172.16.1.1 255.255.0.0 standby 172.16.1.2`** Here, 172.16.1.1 is used for the primary unit, and 172.16.1.2 assigns to the secondary (standby) unit.**Note:** In multiple context mode, you must configure the interface addresses from within each context. Use the **changeto context** command in order to switch between contexts. The command prompt changes to `hostname/context(config-if)#`, where context is the name of the current context.
2. (PIX security appliance platform only) Enable the LAN-based failover.`hostname(config)#` **`failover lan enable`**
3. Designate the unit as the primary unit.`hostname(config)#` **`failover lan unit primary`**
4. Define the failover interface.Specify the interface to be used as the failover interface.`hostname(config)#` **`failover lan interface`** *`if_name phy_if`*In this documentation, the "failover" (interface name for Ethernet0) is used for a failover interface.`hostname(config)#` **`failover lan interface failover`** *`Ethernet3`*The *if_name* argument assigns a name to the interface specified by the *phy_if* argument. The *phy_if* argument can be the physical port name, such as Ethernet1, or a previously created subinterface, such as Ethernet0/2.3.Assign the active and standby IP address to the failover link`hostname(config)#` **`failover interface ip`** *`if_name ip_addr`* **`mask standby`** *`ip_addr`*In this documentation, to configure the failover link, 10.1.0.1 is used for active, 10.1.0.2 for the standby unit, and "failover" is an interface name of Ethernet0.`hostname(config)#` **`failover interface ip failover 10.1.0.1 255.255.255.0 standby 10.1.0.2`** The standby IP address must be in the same subnet as the active IP address. You do not need to identify the standby address subnet mask.The failover link IP address and MAC address do not change at failover. The active IP address for the failover link always stays with the primary unit, while the standby IP address stays with the secondary unit.Enable the interface`hostname(config)#` **`interface`** *`phy_if`* `hostname(config-if)#` **`no shutdown`** In the example, Ethernet3 is used for failover:`hostname(config)#` **`interface ethernet3`** `hostname(config-if)#` **`no shutdown`**
5. (Optional) In order to enable stateful failover, configure the stateful failover link.Specify the interface to be used as the stateful failover link.`hostname(config)#` **`failover link`** *`if_name phy_if`*This example used "state" as an interface name for Ethernet2 to exchange the failover link state information:`hostname(config)#` **`failover link state Ethernet2`** **Note:** If the stateful failover link uses the failover link or a data interface, you only need to supply the *if_name* argument.The *if_name* argument assigns a logical name to the interface specified by the

phy_if argument. The *phy_if* argument can be the physical port name, such as Ethernet1, or a previously created subinterface, such as Ethernet0/2.3. This interface must not be used for any other purpose, except, optionally, as the failover link.Assign an active and standby IP address to the stateful failover link.**Note:** If the stateful failover link uses the failover link or data interface, skip this step. You have already defined the active and standby IP addresses for the interface.`hostname(config)#`**`failover interface`** *`ip if_name ip_addr`* **`mask standby`** *`ip_addr`*The 10.0.0.1 is used as an active and the 10.0.0.2 as a standby IP address for the stateful failover link in this example.`hostname(config)#`**`failover interface ip state 10.0.0.1 255.0.0.0 standby 10.0.0.2`** The standby IP address must be in the same subnet as the active IP address. You do not need to identify the standby address subnet mask.The stateful failover link IP address and MAC address do not change at failover unless they use a data interface. The active IP address always stays with the primary unit, while the standby IP address stays with the secondary unit.Enable the interface.**Note:** If the stateful failover link uses the failover link or data interface, skip this step. You have already enabled the interface.`hostname(config)#`**`interface phy_if`** `hostname(config-if)#`**`no shutdown`** **Note:** For example, in this scenario, Ethernet2 is used for the stateful failover link:`hostname(config)#`**`interface ethernet2`** `hostname(config-if)#`**`no shutdown`**

6. Enable failover.`hostname(config)#`**`failover`** **Note:** Issue the **failover** command on the primary device first, and then issue it on the secondary device. After you issue the **failover** command on the secondary device, the secondary device immediately pulls the configuration from the primary device and sets itself as *standby*. The primary ASA stays up and passes traffic normally and marks itself as the *active* device. From that point on, whenever a failure occurs on the active device, the standby device comes up as active.
7. Save the system configuration to Flash memory.`hostname(config)#`**`copy running-config startup-config`**

Secondary Unit Configuration

The only configuration required on the secondary unit is for the failover interface. The secondary unit requires these commands to initially communicate with the primary unit. After the primary unit sends its configuration to the secondary unit, the only permanent difference between the two configurations is the **failover lan unit** command, which identifies each unit as primary or secondary.

For multiple context mode, all steps are performed in the system execution space unless noted otherwise.

In order to configure the secondary unit, complete these steps:

1. (PIX security appliance platform only) Enable LAN-based failover.`hostname(config)#`**`failover lan enable`**
2. Define the failover interface. Use the same settings that you used for the primary unit.Specify the interface to be used as the failover interface.`hostname(config)#`**`failover lan interface`** *`if_name phy_if`*In this documentation, Ethernet0 is used for a LAN failover interface.`hostname(config)#`**`failover lan interface failover`** *`Ethernet3`*The *if_name* argument assigns a name to the interface specified by the *phy_if* argument.Assign the active and standby IP address to the failover link.`hostname(config)#`**`failover interface ip`** *`if_name ip_addr`* **`mask standby`** *`ip_addr`*In this documentation, to configure the failover link, 10.1.0.1 is used for active, 10.1.0.2 for the standby unit, and "failover" is an interface name of

Ethernet0.`hostname(config)#`**`failover interface ip failover 10.1.0.1 255.255.255.0 standby 10.1.0.2`** **Note:** Enter this command exactly as you entered it on the primary unit when you configured the failover interface on the primary unit.Enable the interface.`hostname(config)#`**`interface`** *`phy_if`* `hostname(config-if)#`**`no shutdown`** For example, in this scenario, Ethernet0 is used for failover.`hostname(config)#`**`interface ethernet3`** `hostname(config-if)#`**`no shutdown`**

3. (Optional) Designate this unit as the secondary unit.`hostname(config)#`**`failover lan unit secondary`** **Note:** This step is optional because, by default, units are designated as secondary unless previously configured.
4. Enable failover.`hostname(config)#`**`failover`** **Note:** After you enable failover, the active unit sends the configuration in running memory to the standby unit. As the configuration synchronizes, the messages *Beginning configuration replication: Sending to mate* and *End Configuration Replication to mate* appear on the active unit console.
5. After the running configuration has completed replication, save the configuration to Flash memory.`hostname(config)#`**`copy running-config startup-config`**

Configurations

This document uses these configurations:

Primary ASA

```
ASA#show running-config ASA Version 7.2(3) ! !--- To set
the firewall mode to transparent mode, !--- use the
firewall transparent command !--- in global
configuration mode. firewall transparent hostname ASA
domain-name default.domain.invalid enable password
2KFQnbNIdI.2KYOU encrypted names ! interface Ethernet0
nameif failover description LAN Failover Interface !
interface Ethernet1 nameif inside security-level 100 !
interface Ethernet2 nameif outside security-level 0 !---
Configure no shutdown in the stateful failover interface
!--- of both Primary and secondary ASA. interface
Ethernet3 nameif state description STATE Failover
Interface ! interface Ethernet4 shutdown no nameif no
security-level no ip address ! interface Ethernet5
shutdown no nameif no security-level no ip address !
passwd 2KFQnbNIdI.2KYOU encrypted ftp mode passive dns
server-group DefaultDNS domain-name
default.domain.invalid access-list 100 extended permit
ip any any pager lines 24 mtu outside 1500 mtu inside
1500 !--- Assign the IP address to the Primary and !---
Seconday ASA Security Appliance. ip address 172.16.1.1
255.255.255.0 standby 172.16.1.2 failover failover lan
unit primary failover lan interface failover Ethernet0
failover lan enable failover key ****** failover link
state Ethernet3 failover interface ip failover 10.1.0.1
255.255.255.0 standby 10.1.0.2 failover interface ip
state 10.0.0.1 255.0.0.0 standby 10.0.0.2 asdm image
flash:/asdm-522.bin no asdm history enable arp timeout
14400 access-group 100 in interface outside route
outside 0.0.0.0 0.0.0.0 172.16.1.3 1 timeout xlate
3:00:00 timeout conn 1:00:00 half-closed 0:10:00 udp
0:02:00 icmp 0:00:02 timeout sunrpc 0:10:00 h323 0:05:00
h225 1:00:00 mgcp 0:05:00 mgcp-pat 0:05:00 timeout sip
0:30:00 sip_media 0:02:00 sip-invite 0:03:00 sip-
disconnect 0:02:00 timeout uauth 0:05:00 absolute no
snmp-server location no snmp-server contact snmp-server
```

```
enable traps snmp authentication linkup linkdown
coldstart telnet timeout 5 ssh timeout 5 console timeout
0 ! class-map inspection_default match default-
inspection-traffic ! ! policy-map type inspect dns
preset_dns_map parameters message-length maximum 512
policy-map global_policy class inspection_default
inspect dns preset_dns_map inspect ftp inspect h323 h225
inspect h323 ras inspect netbios inspect rsh inspect
rtsp inspect skinny inspect esmtp inspect sqlnet inspect
sunrpc inspect tftp inspect sip inspect xdmcp ! service-
policy global_policy global prompt hostname context
Cryptochecksum:d41d8cd98f00b204e9800998ecf8427e : end
```

Secondary ASA

```
ASA#show running-config ASA Version 7.2(3) ! hostname
ASA domain-name default.domain.invalid enable password
2KFQnbNIdI.2KYOU encrypted names ! failover failover lan
unit secondary failover lan interface failover Ethernet0
failover lan enable failover key ****** failover
interface ip failover 10.1.0.1 255.255.255.0 standby
10.1.0.2
```

Verify

Use of the show failover Command

This section describes the **show failover** command output. On each unit, you can verify the failover status with the **show failover** command.

Primary ASA

```
ASA#show failover Failover On Cable status: N/A - LAN-based failover enabled Failover
unit Primary Failover LAN Interface: failover Ethernet0 (up) Unit Poll frequency 200
milliseconds, holdtime 800 milliseconds Interface Poll frequency 5 seconds, holdtime
25 seconds Interface Policy 1 Monitored Interfaces 2 of 250 maximum Version: Ours
7.2(3), Mate 7.2(3) Last Failover at: 00:08:03 UTC Jan 1 1993 This host: Primary -
Active Active time: 1820 (sec) Interface inside (172.16.1.1): Normal Interface
outside (172.16.1.1): Normal Other host: Secondary - Standby Ready Active time: 0
(sec) Interface inside (172.16.1.2): Normal Interface outside (172.16.1.2): Normal
Stateful Failover Logical Update Statistics Link : state Ethernet3 (up) Stateful Obj
xmit xerr rcv rerr General 185 0 183 0 sys cmd 183 0 183 0 up time 0 0 0 0 RPC
services 0 0 0 0 TCP conn 0 0 0 0 UDP conn 0 0 0 0 ARP tbl 0 0 0 0 L2BRIDGE Tbl 2 0 0
0 Xlate_Timeout 0 0 0 0 Logical Update Queue Information Cur Max Total Recv Q: 0 1
7012 Xmit Q: 0 1 185
```

Secondary ASA

```
ASA(config)#show failover Failover On Cable status: N/A - LAN-based failover enabled
Failover unit Secondary Failover LAN Interface: failover Ethernet0 (up) Unit Poll
frequency 200 milliseconds, holdtime 800 milliseconds Interface Poll frequency 5
seconds, holdtime 25 seconds Interface Policy 1 Monitored Interfaces 2 of 250 maximum
Version: Ours 7.2(3), Mate 7.2(3) Last Failover at: 16:39:12 UTC Aug 9 2009 This
host: Secondary - Standby Ready Active time: 0 (sec) Interface inside (172.16.1.2):
Normal Interface outside (172.16.1.2): Normal Other host: Primary - Active Active
time: 1871 (sec) Interface inside (172.16.1.1): Normal Interface outside
(172.16.1.1): Normal Stateful Failover Logical Update Statistics Link : state
Ethernet3 (up) Stateful Obj xmit xerr rcv rerr General 183 0 183 0 sys cmd 183 0 183
0 up time 0 0 0 0 RPC services 0 0 0 0 TCP conn 0 0 0 0 UDP conn 0 0 0 0 ARP tbl 0 0
0 0 L2BRIDGE Tbl 0 0 0 0 Xlate_Timeout 0 0 0 0 Logical Update Queue Information Cur
Max Total Recv Q: 0 1 7043 Xmit Q: 0 1 183
```

Use the **show failover state** command to verify the state.

Primary ASA

```
ASA#show failover state State Last Failure Reason Date/Time This host - Primary
Active None Other host - Secondary Standby Ready Comm Failure 00:02:36 UTC Jan 1 1993
====Configuration State=== Sync Done ====Communication State=== Mac set
```

Secondary unit

```
ASA#show failover state State Last Failure Reason Date/Time This host - Secondary
Standby Ready None Other host - Primary Active None ====Configuration State=== Sync
Done - STANDBY ====Communication State=== Mac set
```

In order to verify the IP addresses of the failover unit, use the **show failover interface** command.

Primary unit

```
ASA#show failover interface interface failover Ethernet0 System IP Address: 10.1.0.1
255.255.255.0 My IP Address : 10.1.0.1 Other IP Address : 10.1.0.2 interface state
Ethernet3 System IP Address: 10.0.0.1 255.255.255.0 My IP Address : 10.0.0.1 Other IP
Address : 10.0.0.2
```

Secondary unit

```
ASA#show failover interface interface failover Ethernet0 System IP Address: 10.1.0.1
255.255.255.0 My IP Address : 10.1.0.2 Other IP Address : 10.1.0.1 interface state
Ethernet3 System IP Address: 10.0.0.1 255.255.255.0 My IP Address : 10.0.0.2 Other IP
Address : 10.0.0.1
```

View of Monitored Interfaces

In order to view the status of monitored interfaces: In single context mode, enter the **show monitor-interface** command in global configuration mode. In multiple context mode, enter the **show monitor-interface** within a context.

Primary ASA

```
ASA(config)#show monitor-interface This host: Primary - Active Interface inside
(172.16.1.1): Normal Interface outside (172.16.1.1): Normal Other host: Secondary -
Standby Ready Interface inside (172.16.1.2): Normal Interface outside (172.16.1.2):
Normal
```

Secondary ASA

```
ASA(config)#show monitor-interface This host: Secondary - Standby Ready Interface
inside (172.16.1.2): Normal Interface outside (172.16.1.2): Normal Other host:
Primary - Active Interface inside (172.16.1.1): Normal Interface outside
(172.16.1.1): Normal
```

Note: If you do not enter a failover IP address, the **show failover** command displays 0.0.0.0 for the IP address and interface monitoring remains in a *waiting* state. Refer to the show failover section of the *Cisco Security Appliance Command Reference, Version 7.2* for more information about the different failover states.

Display of the Failover Commands in the Running Configuration

In order to view the failover commands in the running configuration, enter this command:

```
hostname(config)#show running-config failover
```

All of the failover commands are displayed. On units that run in multiple context mode, enter the **show running-config failover** command in the system execution space. Enter the **show running-config all failover** command in order to display the failover commands in the running configuration and include commands for which you have not changed the default value.

Failover Functionality Tests

Complete these steps in order order to test failover functionality:

1. Test that your active unit or failover group passes traffic as expected with FTP (for example) to send a file between hosts on different interfaces.
2. Force a failover to the standby unit with this command:For Active/Standby Failover, enter this command on the active unit:`hostname(config)#no failover active`
3. Use FTP to send another file between the same two hosts.
4. If the test was not successful, enter the **show failover command** in order to check the failover status.
5. When you are finished, you can restore the unit or failover group to active status with this command:For Active/Standby Failover, enter this command on the active unit:`hostname(config)#failover active`

Forced Failover

In order to force the standby unit to become active, enter one of these commands:

Enter this command on the standby unit:

```
hostname#failover active
```

Enter this command on the active unit:

```
hostname#no failover active
```

Disabled Failover

In order to disable failover, enter this command:

```
hostname(config)#no failover
```

If you disable failover on an Active/Standby pair, it causes the active and standby state of each unit to be maintained until you restart. For example, the standby unit remains in standby mode so that both units do not start to pass traffic. In order to make the standby unit active (even with failover disabled), see the Forcing Failover section.

If you disable failover on an Active/Active pair, it causes the failover groups to remain in the active state on whichever unit they are currently active on, no matter which unit they are configured to prefer. The **no failover** command can be entered in the system execution space.

Restoration of a Failed Unit

In order to restore a failed unit to an unfailed state, enter this command:

```
hostname(config)#failover reset
```

If you restore a failed unit to an unfailed state, it does not automatically make it active; restored units or groups remain in the standby state until made active by failover (forced or natural). An exception is a failover group configured with the preempt command. If previously active, a failover group becomes active if it is configured with the preempt command and if the unit on which it failed is its preferred unit.

Troubleshoot

When a failover occurs, both security appliances send out system messages. This section includes these topics

Failover Monitoring

This example demonstrates what happens when failover has not started to monitor the network interfaces. Failover does not start to monitor the network interfaces until it has heard the second `hello` packet from the other unit on that interface. This takes about 30 seconds. If the unit is attached to a network switch that runs Spanning Tree Protocol (STP), this takes twice the `forward delay` time configured in the switch, which is typically configured as 15 seconds, plus this 30 second delay. This is because at ASA bootup and immediately after a failover event, the network switch detects a temporary bridge loop. Upon detection of this loop, it stops to forward packets on these interfaces for the `forward delay` time. It then enters the `listen` mode for an additional `forward delay` time, within which time the switch listens for bridge loops but does not forward traffic or forward failover `hello` packets. After twice the forward delay time (30 seconds), traffic flow resumes. Each ASA remains in a `waiting` mode until it hears 30 seconds worth of `hello` packets from the other unit. Within the time that the ASA passes traffic, it does not fail the other unit based on not hearing the `hello` packets. All other failover monitoring still occurs, that is, Power, Interface Loss of Link, and Failover Cable `hello`.

For failover, Cisco strongly recommends that customers enable portfast on all switch ports that connect to ASA interfaces. In addition, channeling and trunking must be disabled on these ports. If the interface of the ASA goes down within failover, the switch does not have to wait 30 seconds while the port transitions from a state of listening to learning to forwarding.

```
Failover On
Cable status: Normal
Reconnect timeout 0:00:00
```

```
This host: Primary - Active
Active time: 6930 (sec)
Interface inside (172.16.1.1): Normal (Waiting)
Interface outside (172.16.1.1): Normal (Waiting)
Other host: Secondary - Standby
Active time: 15 (sec)
Interface inside (172.16.1.2): Normal (Waiting)
Interface outside (172.16.1.2): Normal (Waiting)
```

In summary, check these steps in order to narrow down the failover problems:

- Check the network cables connected to the interface in the waiting/failed state and, if it is possible, replace them.
- If there is a switch connected between the two units, verify that the networks connected to the interface in the waiting/failed state function correctly.
- Check the switch port connected to the interface in the waiting/failed state and, if it is possible, use the another FE port on the switch..
- Check that you have enabled port fast and disabled both trunking and channeling on the switch ports that are connected to the interface.

Unit Failure

In this example, failover has detected a failure. Note that Interface 1 on the primary unit is the source of the failure. The units are back in `waiting` mode because of the failure. The failed unit has removed itself from the network (interfaces are down) and no longer sends `hello` packets on the network. The active unit remains in a `waiting` state until the failed unit is replaced and failover communications starts again.

```
Failover On
Cable status: Normal
Reconnect timeout 0:00:00
This host: Primary - Standby (Failed)
Active time: 7140 (sec)
Interface inside (172.16.1.2): Normal (Waiting)
Interface outside (172.16.1.2): Failed (Waiting)
Other host: Secondary - Active
Active time: 30 (sec)
Interface inside (172.16.1.1): Normal (Waiting)
Interface outside (172.16.1.1): Normal (Waiting)
```

LU allocate connection failed

A memory problem might exist if you receive this error message:

```
LU allocate connection failed
```

This issue is documented in Cisco bug ID CSCte80027 (registered customers only) . In order to resolve this issue, upgrade your firewall to a software version in which this bug is fixed. Some of the ASA software versions under which this bug got fixed are 8.2(4), 8.3(2), 8.4(2).

Failover System Messages

The security appliance issues a number of system messages related to failover at priority level 2, which indicates a critical condition. In order to view these messages, refer to the Cisco Security

Appliance Logging Configuration and System Log Messages to enable logging and to see descriptions of the system messages.

Note: Within switchover, failover logically shuts down and then brings up interfaces, which generates syslog **411001** and **411002** messages. This is normal activity.

Debug Messages

In order to see debug messages, enter the **debug fover** command. Refer to the Cisco Security Appliance Command Reference for more information.

Note: Because debugging output is assigned high priority in the CPU process, it can drastically affect system performance. For this reason, use the **debug fover** commands only to troubleshoot specific problems or within troubleshooting sessions with Cisco technical support staff.

SNMP

In order to receive SNMP syslog traps for failover, configure the SNMP agent to send SNMP traps to SNMP management stations, define a syslog host, and compile the Cisco syslog MIB into your SNMP management station. Refer to the **snmp-server** and **logging** commands in the Cisco Security Appliance Command Reference for more information.

Failover Polltime

In order to specify the failover unit poll and hold times, use the **failover polltime** command in global configuration mode.

The `failover polltime unit msec [time]` polls hello messages in order to represent the time interval in order to check the existence of the standby unit.

Similarly, the `failover holdtime unit msec [time]` represents the setting a time period during which a unit must receive a hello message on the failover link, after which the peer unit is declared failed.

In order to specify the data interface poll and hold times in an Active/Standby failover configuration, use the **failover polltime interface** command in global configuration mode. In order to restore the default poll and hold times, use the **no** form of this command.

```
failover polltime interface [msec] time [holdtime time]
```

Use the **failover polltime interface** command in order to change the frequency at which hello packets are sent out on data interfaces. This command is available for Active/Standby failover only. For Active/Active failover, use the **polltime interface** command in the failover group configuration mode instead of the **failover polltime interface** command.

You cannot enter a *holdtime* value that is less than 5 times the interface poll time. With a faster poll time, the security appliance can detect failure and trigger failover faster. However, faster detection can cause unnecessary switchovers when the network is temporarily congested. Interface testing begins when a hello packet is not heard on the interface for over half the hold time.

You can include both failover polltime unit and failover polltime interface commands in the

configuration.

This example sets the interface poll time frequency to 500 milliseconds and the hold time to 5 seconds:

```
hostname(config)#failover polltime interface msec 500 holdtime 5
```

Refer to the failover polltime section of the *Cisco Security Appliance Command Reference, Version 7.2* for more information.

Export Certificate/Private Key in Failover Configuration

The primary device automatically replicates the private key/certificate to the secondary unit. Issue the **write memory** command in the active unit in order to replicate the configuration, which includes the certificate/private key, to the standby unit. All the keys/certificates on the standby unit are erased and repopulated by the active unit configuration.

Note: You must not manually import the certificates, keys, and trust points from the active device and then export to the standby device.

WARNING: Failover message decryption failure.

Error message:

```
Failover message decryption failure. Please make sure both units have the
same failover shared key and crypto license or system is not out of memory
```

This problem occurs due to failover key configuration. In order to resolve this issue, remove the failover key, and configure the new shared key.

Problem: Failover is always flapping after configuring the transparent Active/Standby multiple mode failover

Failover is steady when the inside interfaces of both ASAs are directly connected and outside interfaces of both ASA are directly connected. But failover is flapping when a switch is used in between.

Solution: Disable the BPDU on the ASA interfaces in order to resolve this issue.

ASA Modules Failover

If Advanced Inspection and Prevention Security Services Module (AIP-SSM) or Content Security and Control Security Services Module (CSC-SSM) are used in active and standby units, then it operates independently of the ASA in terms of failover. **Modules must be configured manually in active and standby units, the failover does not replicate the module configuration.**

In terms of failover, both ASA units that have AIP-SSM or CSC-SSM modules must be of the same hardware type. For example, if the primary unit have the ASA-SSM-10 module, the secondary unit must have the ASA-SSM-10 module.

Failover message block alloc failed

Error Message `%PIX|ASA-3-105010: (Primary) Failover message block alloc failed`

Explanation: Block memory was depleted. This is a transient message, and the security appliance should recover. *Primary* can also be listed as *Secondary* for the secondary unit.

Recommended Action: Use the **show blocks** command in order to monitor the current block memory.

AIP Module Failover Problem

If you have two ASAs in a failover configuration and each has an AIP-SSM, you must manually replicate the configuration of the AIP-SSMs. Only the configuration of the ASA is replicated by the failover mechanism. The AIP-SSM is not included in the failover.

First, the AIP-SSM operates independently of the ASA in terms of failover. For failover, all that is needed from an ASA perspective is that the AIP modules be of the same hardware type. Beyond that, as with any other portion of failover, the configuration of the ASA between the active and standby must be in sync.

As for the set up of the AIPs, they are effectively independent sensors. There is no failover between the two, and they have no awareness of each other. They can run independent versions of code. That is, they do not have to match, and the ASA does not care about the version of code on the AIP with respect to failover.

ASDM initiates a connection to the AIP through the management interface IP that you configured on the AIP. In other words, it connects to the sensor typically through HTTPS, which depends on how you set up the sensor.

You could have a failover of the ASA independent of the IPS (AIP) modules. You are still connected to the same one because you connect to its management IP. In order to connect to the other AIP, you must reconnect to its manangement IP to configure it and access it.

Refer to ASA: Send Network Traffic from the ASA to the AIP SSM Configuration Example for more information and sample configurations on how to send network traffic that passes through the Cisco ASA 5500 Series Adaptive Security Appliance (ASA) to the Advanced Inspection and Prevention Security Services Module (AIP-SSM) (IPS)

Known Issues

When you attempt to access the ASDM on the secondary ASA with version 8.x software and ASDM version 6.x for failover configuration, this error is received:

```
Error: The name on the security certificate is invalid or does not match the name of the site
```

In the certificate, the Issuer and the Subject Name is the IP address of the *active* unit, not the IP address of the *standby* unit.

In ASA version 8.x, the internal (ASDM) certificate is replicated from the active unit to the standby unit, which causes the error message. But, if the same firewall runs on version 7.x code with 5.x ASDM and you try to access ASDM, you receive this regular security warning:

```
The security certificate has a valid name matching the name of the page you are trying to view
```

When you check the certificate, the issuer and the subject name is the IP address of the standby unit.

Related Information

- **Cisco ASA 5500 Series Adaptive Security Appliances**
- **Cisco PIX Firewall Software**
- **Firewall Services Module (FWSM) Failover Configuration**
- **FWSM Failover Troubleshooting**
- **How Failover Works on the Cisco Secure PIX Firewall**
- **Technical Support & Documentation - Cisco Systems**