

A10

ACOS 6.0.7

Configuring VRRP-A High Availability

March, 2025

© 2025 A10 Networks, Inc. All rights reserved.

Information in this document is subject to change without notice.

PATENT PROTECTION

A10 Networks, Inc. products are protected by patents in the U.S. and elsewhere. The following website is provided to satisfy the virtual patent marking provisions of various jurisdictions including the virtual patent marking provisions of the America Invents Act. A10 Networks, Inc. products, including all Thunder Series products, are protected by one or more of U.S. patents and patents pending listed at:

[a10-virtual-patent-marking](#).

TRADEMARKS

A10 Networks, Inc. trademarks are listed at: [a10-trademarks](#)

CONFIDENTIALITY

This document contains confidential materials proprietary to A10 Networks, Inc.. This document and information and ideas herein may not be disclosed, copied, reproduced or distributed to anyone outside A10 Networks, Inc. without prior written consent of A10 Networks, Inc.

DISCLAIMER

This document does not create any express or implied warranty about A10 Networks, Inc. or about its products or services, including but not limited to fitness for a particular use and non-infringement. A10 Networks, Inc. has made reasonable efforts to verify that the information contained herein is accurate, but A10 Networks, Inc. assumes no responsibility for its use. All information is provided "as-is." The product specifications and features described in this publication are based on the latest information available; however, specifications are subject to change without notice, and certain features may not be available upon initial product release. Contact A10 Networks, Inc. for current information regarding its products or services. A10 Networks, Inc. products and services are subject to A10 Networks, Inc. standard terms and conditions.

ENVIRONMENTAL CONSIDERATIONS

Some electronic components may possibly contain dangerous substances. For information on specific component types, please contact the manufacturer of that component. Always consult local authorities for regulations regarding proper disposal of electronic components in your area.

FURTHER INFORMATION

For additional information about A10 products, terms and conditions of delivery, and pricing, contact your nearest A10 Networks, Inc. location, which can be found by visiting www.a10networks.com.

Table of Contents

Overview of VRRP-A	6
VRRP-A Overview	7
VRRP-A Configuration	7
Basic VRRP-A Configuration	7
Multicast Heartbeat Destination Addresses	9
Unicast Heartbeat Destination Addresses	9
VRRP-A and Virtual Router IDs	10
Leader and Follower VRIDs	11
VRID Virtual MAC Addresses	11
VRRP-A, Virtual Router IDs, and Partitions	12
VRRP-A Active / Standby Device Selection	12
VRRP-A Active Device Selection Overview	13
Event Tracking for Weight and Priority	14
Priority Calculation	15
Track Event Delay	15
Preemption and VRRP-A Failover Behavior with Preemption Disabled	15
Preemption Delay	16
Active Device Selection Examples	16
VRRP-A Failover	20
Policy-Based Failover	20
Dynamic Priority Reduction	20
Force-Self-Standby	21
VRRP-A and Floating IP Addresses	22
VRRP-A and Configuration Synchronization	24
VRRP-A and Session Synchronization	25
VRRP-A Interfaces	26
Explicitly Configured VRRP-A Interfaces	26
Using the GUI to Configure a VRRP-A Interface	27

Using the CLI to Configure a VRRP-A Interface	28
Preferred Receive Interface for Session Synchronization	28
Manually Synchronizing Configurations of All Partitions Between ACOS Devices	29
Deploying VRRP-A	30
Basic VRRP-A Deployment	31
Force-Self-Standby Reload or Restart Persistence	31
Viewing VRRP-A Information	32
Viewing VRRP-A Config-Sync Status in the CLI	33
Viewing Peer Box Information	37
Configuring Preemption Delays	39
Use the GUI to Configure VRRP-A Preemption Delay	39
Use the CLI to Configure VRRP-A Preemption Delay	39
VRRP-A Configuration Examples	41
Simple Deployment	42
Deployment with Custom Priority Settings	42
Configuration of Tracking Options	44
Ethernet Interface Tracking	45
Trunk Tracking	45
VLAN Tracking	45
Gateway Tracking	46
Route Tracking	46
Preemption Delay	47
Increased VLANs for VRRP-A Failover Tracking	47
VLAN Tracking Configuration Example	48
Configuring a Leader and Follower VRID	49
Configuration Example 1	49
Configuration Example 2	50
Configuring VRRP-A VRID Lead Switching	51
Manual VRID Lead Switching Using Manual Configuration	52
Automatic VRID Lead Switching Using aVCS	53

VRRP-A Status in CLI Prompt	53
VRRP-A Policy-Based Failover Template	54
Template Rules and Partitions	55
Policy-Based Failover Template Rules for the Shared Partition	55
Policy-Based Failover Template Rules for L3V Partitions	55
Preemption and Levels	56
Precedence Causing a Failover	56
Higher Weight Scenario	57
Higher Priority Scenario	58
Equal Weight and Priority Scenario	59
Events Tracked for Weight via the Templates	60
Configuring VRRP-A Policy-Based Failovers	62
Using the GUI to Configure VRRP-A Policy-Based Failovers	63
Using the CLI to Create VRRP-A Policy-Based Failovers	64
Configuring Templates in the Shared Partition	64
Configuring Templates in L3V Partition	65
Verifying the Fail-Over Policy Template Configuration	65
VRRP-A Service Migration	68
Guidelines for Service Migration	69
Deploying Service Migration	69
Step 1: Configuring the Destination VRRP-A Pod	69
Step 2: Running the Migration Command	72
Step 3: Lowering the Destination Pod OSPF Cost Metric	73
Step 4: Running the Finish Migration Command	74
Glossary	76

Overview of VRRP-A

This chapter provides an overview of VRRP-A's high availability.

The following topics are covered:

VRRP-A Overview	7
VRRP-A Configuration	7
VRRP-A Active / Standby Device Selection	12
VRRP-A Failover	20
VRRP-A and Floating IP Addresses	22
VRRP-A and Configuration Synchronization	24
VRRP-A and Session Synchronization	25
VRRP-A Interfaces	26
Manually Synchronizing Configurations of All Partitions Between ACOS Devices	29

VRRP-A Overview

VRRP-A is the ACOS implementation of high availability that is completely different from the industry-standard implementation of Virtual Router Redundancy Protocol (VRRP). For purposes of operational familiarity, it borrows concepts from VRRP but is significantly different from VRRP. VRRP-A will not inter-operate with VRRP.

VRRP-A simplifies the configuration of multi-system redundancy and allows up to eight physical or virtual ACOS devices to serve as mutual backups for IP addresses.

VRRP-A inline mode is supported only with one VRRP-A VRID ID group.

VRRP-A can provide redundancy for the following IP resources:

- Virtual server IP addresses (VIPs)
- Floating IP addresses used as default gateways by downstream devices
- IPv6 NAT pools
- IPv4 NAT pools
- IPv4 static range lists and individual mappings for inside source NAT

VRRP-A Configuration


The following topics are covered:

Basic VRRP-A Configuration	7
VRRP-A and Virtual Router IDs	10
VRRP-A, Virtual Router IDs, and Partitions	12

Basic VRRP-A Configuration

[Figure 1](#) illustrates a basic VRRP-A configuration with two devices.

Figure 1 : Basic VRRP-A Configuration



This type of configuration is called Active-Standby mode, because only the active device processes traffic.

The elements in this illustration are described in [Table 1](#)

Table 1 : VRRP-A Basic Configuration Elements

Element	Description
Set ID	Unique identifier for a set of VRRP-A devices. All the devices in the set must be in the same Layer 2 broadcast domain.
Device ID	Unique identifier for each device within the VRRP-A set.
Heartbeat	The active ACOS device in the VRID periodically sends hello messages to the standby ACOS device and other backup devices. The hello messages indicate that the active device for the VRID is still operating. For more information, see: <ul style="list-style-type: none"> • Multicast Heartbeat Destination Addresses • Unicast Heartbeat Destination Addresses
Priority	Priority is used to help determine the order in which standby devices should become active devices. For more information, see VRRP-A Active / Standby Device Selection .

Multicast Heartbeat Destination Addresses

By default, VRRP-A uses an IP multicast address as the destination for VRRP-A heartbeat messages to peers. VRRP-A sends multicast hello messages to the following addresses:

- IPv4 multicast address 224.0.0.210, MAC address 01:00:5e:00:00:D2
- IPv6 link-local multicast address FF02::D2, MAC address 33:33:00:00:00:D2

If the standby device stops receiving hello messages from the active device, the operation for the VRID fails over to the standby device. The device to which the operation fails over becomes the new active device for the VRID.

To configure multicast heartbeats, use the `hello-interval` command. For more information about the command, see the *Command Line Interface Reference* document.

Unicast Heartbeat Destination Addresses

Instead of multicast heartbeat messages, you can configure VRRP-A to use unicast heartbeats instead for layer 3 redundancy across subnets using VRRP-A.

To do so, configure a VRRP-A peer group on each device in the VRRP-A set. For redundancy in the peer group on each ACOS device, you can configure multiple IP addresses for multiple data interfaces. You can configure a maximum of 16 IP addresses on the peer group.

NOTE: At least one unicast address must be configured on a data interface per peer ACOS device.

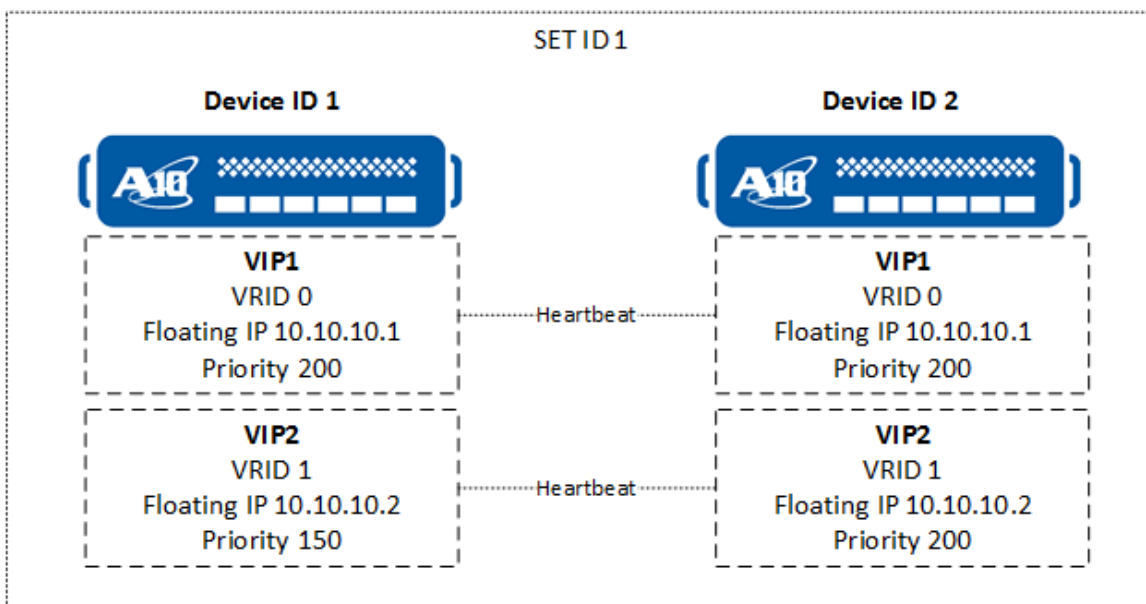
For reliability and redundancy, configure at least four IP Addresses in a peer group, with each IP address belonging to a different interface. This will enable you to allocate at least two interfaces per ACOS device. The peer group configuration on each VRRP-A device in the set should be the same, including the device's own IP address. When a device sends VRRP-A heartbeats to the members of a peer group, that device skips any IP addresses that belong to itself.

To configure peer groups for unicast heartbeat messages, use the `vrrp-a peer-group` command. For more information about the command, see the *Command Line Interface Reference* document.

VRRP-A and Virtual Router IDs

VRRP-A device ID, set ID, and priority can be configured per Virtual Router ID (VRID), as shown in [Figure 2](#).

Figure 2 : VRRP-A Configuration with VRIDs



The configuration example in [Figure 2](#) is called Active-Active mode because both devices can potentially process traffic on different VIPs.

A VRID is a logical container grouping functional configuration elements (for example, NAT pools, virtual servers, or floating IP addresses) together. Those elements, in turn, are picked up and processed by another device in the VRRP-A set in case of a failover.

By default, the shared partition and each L3V partition have its own VRID. The numerical value for this default VRID is 0; in previous releases, you could configure this VRID using the keyword `default`, which is no longer supported in Release 4.0.

NOTE: By default, a server load balancing virtual IP (VIP) is associated with VRID 0.

Admins with write privileges for the partition can assign a floating IP address to the partition's VRID. Generally, the floating IP address provides redundancy for the

default gateway IP address used by downstream devices. (For more information, see [VRRP-A and Floating IP Addresses](#).)

Parameters related to the selection of the active device and failover can be configured.

The total number of VRIDs that can be configured is dependent on the device model.

All devices in a VRRP-A cluster must use the same value for the IPv6 prefix length irrespective of the number of VRIDs configured. For more information about IPv6 prefix length, see *Command Line Interface Reference*.

NOTE:

The IPv6 prefix length must be modified only when the VRRP-A cluster is not in use. Else, unpredictable disruptions may occur to the traffic.

Leader and Follower VRIDs

Configure a VRID as a “leader” VRID and some others as “follower” VRIDs. The benefit of configuring the leader VRID is that when you are running the risk of consuming all your resources, one VRID can serve as the leader and others can be configured to operate as followers. This means that only on one VRID will you be able to configure all capabilities, such as configuring the failover policy template, the preempt mode, priority, and tracking options, while in the follower VRID, you can only configure Floating IP Addresses.

Any time the maximum number of VRIDs is exceeded for a device, the VRIDs are automatically configured as follower VRIDs or can be explicitly configured to be added as a follower VRID. To configure a leader or a follower VRID, refer to [Configuring a Leader and Follower VRID](#).

VRID Virtual MAC Addresses

VRRP-A assigns a virtual MAC address to each VRID. The VRRP-A virtual MAC addresses are numbered as follows:

```
021f.a00n.nnnn
```

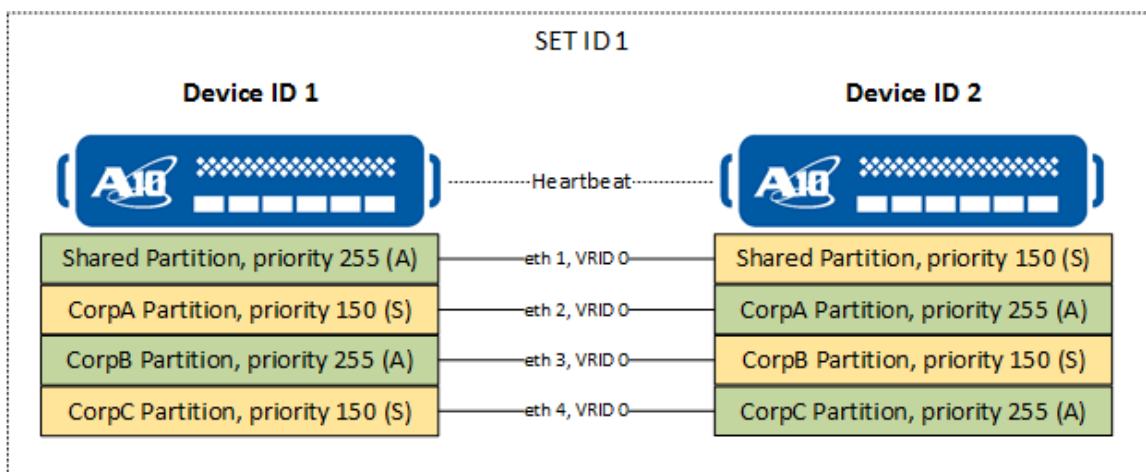
The `n.nnnn` portion of the address includes the partition ID, VRRP-A set ID, and VRID.

VRRP-A, Virtual Router IDs, and Partitions

VRRP-A is supported in the shared partition and L3V partitions. Layer 3 Virtualization allows each L3V partition to have its own VRID, independent from the VRIDs belonging to other partitions.

[Figure 3](#) shows an example configuration with VRID instances spanning L3V partitions across separate physical devices.

Figure 3 : VRRP-A Configuration with Partitions



In this example, a pair of ACOS devices are configured with L3V partitions. The VIPs and other IP resources in each partition are backed up by the partition's VRID. At any given time, one of the devices is the active device (A) for a VRID and the other device is the standby (S) for the VRID. For more information about how the active device is selected, see [VRRP-A Active / Standby Device Selection](#).

VRRP-A Active / Standby Device Selection

The following topics are covered:

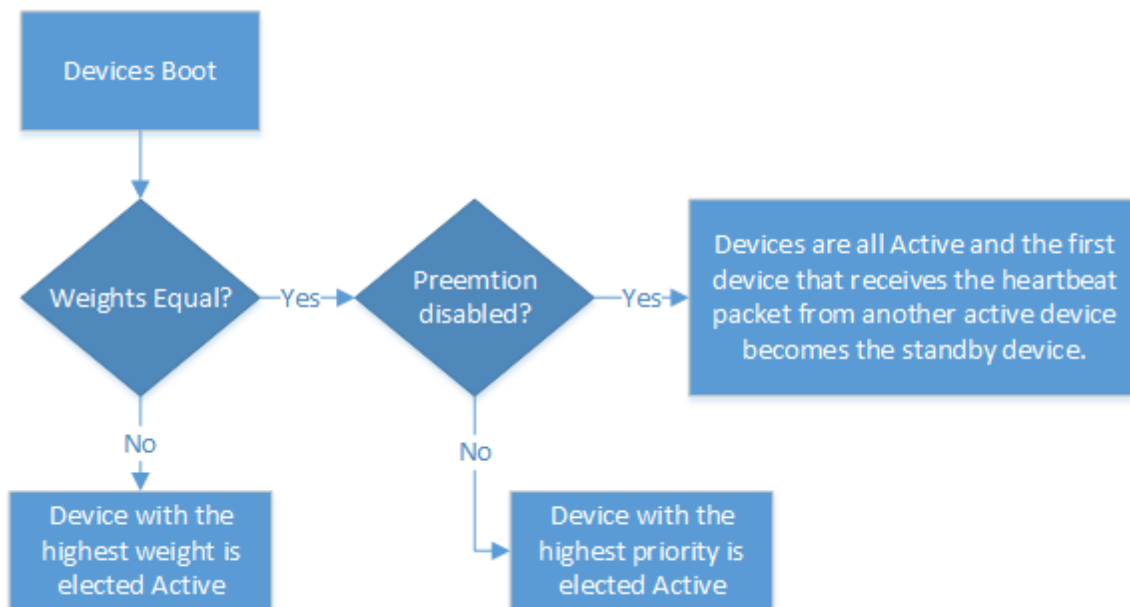
VRRP-A Active Device Selection Overview	13
Event Tracking for Weight and Priority	14
Priority Calculation	15
Track Event Delay	15
Preemption and VRRP-A Failover Behavior with Preemption Disabled	15

Preemption Delay	16
Active Device Selection Examples	16

VRRP-A Active Device Selection Overview

[Figure 4](#) shows the flowchart about how a device is selected as the Active device. This process is performed per VRID.

Figure 4 : VRRP-A Device Selection Flowchart



NOTE: VRRP-A active device selection is unrelated to and completely independent of, aVCS PU1 device selection in a virtual chassis.

The VRRP-A device selection process consists of two levels of failover considerations:

- The weight assigned to an event.
- Priority value assigned as part of the VRRP-A deployment.

While both weight and priority are factored into failover decisions, the weight of the VRID takes precedence over the VRID's priority.

Each VRID is allocated a fixed, total weight of 65534 and each event is allocated a corresponding configurable priority of 1-255. When the event occurs, the configured

weight for the event is deducted from the total weight assigned to the VRID. Some events (for example, when an interface goes down) are considered to be critical for a VRID and will reduce the VRID's weight. ACOS devices that are part of the same set periodically exchange their weight information with peer ACOS devices.

VRRP-A determines the Active or Standby status based on received weight information. Based on the newly computed weight, a failover will occur from an ACOS device of lower weight to another ACOS device of a higher weight. This concept is called preemption.

Preemption based on weight is always enabled. Preemption based on priority can be enabled or disabled.

Event Tracking for Weight and Priority

An ACOS device can be configured to track VRRP-A events for weight or priority or both. While both weight and priority can be configured to track the exact same events, such as the state of gateways, trunks, VLANs, interfaces, or routes, how they are accessed and assigned these values are different.

When you configure VRRP-A using the GUI or command line at the VRRP-A configuration level, you can specify the priority assigned to each event. When the event occurs, the value you assign to the event will be deducted from the priority you configure from 1-255 for the VRID of a given ACOS device.

However, if you specify a value for a failed event using a VRRP-A failover template, you are configuring the weight assigned to an event that will be deducted from a fixed total weight of 65534 assigned to every VRID for an ACOS device.

The fixed total weight for a VRID is not user configurable, unlike the priority assigned the VRID for an ACOS device which is configurable. A failed event will result in a deduction of the weight assigned to the event from the fixed total weight for the VRID. Since the weight of each VRID is tallied to figure out the weight and/or priority of an ACOS device, when multiple devices are configured with VRRP-A, this value is used to gauge the device that will serve as the Active or the Standby device. When the weight for all VRRP-A enabled devices are unequal, the device with the highest weight will serve as the Active device, the rest will be placed in a Standby state.

Priority Calculation

Every few seconds, VRRP-A recalculates the priority for each VRID on the active device for the VRIDs.

To calculate the priority, VRRP-A subtracts the priority values for all optional failover trigger events that have occurred from the configured priority value. The lowest value to which the priority can be reduced is 1.

Once the link or server health is restored, the failover trigger event is no longer occurring and the priority value is not subtracted during the next priority calculation.

Track Event Delay

To prevent unnecessary failovers caused by brief, temporary link or server health changes, VRRP-A uses a track event delay. The track event delay is the number of seconds a failover-causing event must persist before a failover occurs.

For example, if the track event delay is 5 seconds, and a tracked link goes down for 3 seconds, then comes back up, no failover is triggered. Failover is triggered only if the link stays down for at least 5 seconds.

The track event delay can be from 1-tenth second to 10 seconds. The default is 3 seconds.

Preemption and VRRP-A Failover Behavior with Preemption Disabled

Preemption allows failover to be triggered by manually changing the priority on one or more devices so that the active device no longer has the highest priority for the VRID.

Preemption is enabled by default. If preemption is disabled, failover is not triggered by manual priority changes. If disabling preemption, it is strongly recommended to configure a [VRRP-A Policy-Based Failover Template](#) to govern any failover events.

The VRRP-A failover behavior with preemption mode disabled is as follows. When a heartbeat link is lost, the VRRP pair becomes Active / Active. If preemption is

disabled, during the time that a heartbeat link is re-established, the VRRP device that received the heartbeat message first becomes the Standby device.

NOTE: CLI configuration for disabling preemption mode is `preempt-mode disable`.

Preemption can be configured on an individual VRID basis, by shared partition admins and private partition admins.

Preemption Delay

If VRRP-A preemption is enabled, failover can be triggered by VRRP-A configuration changes, in addition to network changes. By default, when preemption is enabled, failover can occur as soon as 3 seconds after an applicable configuration change takes effect.

This release enables you to configure a delay between VRRP-A configuration changes and the failover that occurs due to the changes. The default VRRP-A preemption delay is 6 seconds (60 units of 100 ms); the value can be customized to a value from 100 ms (1 unit of 100 ms) to 25.5 seconds (255 units of 100 ms).

NOTE: Preemption delay is only valid when VRRP-A is configured in multi-active environments, where a minimum of three active devices are configured.

In deployments where many sessions are synchronized, setting the preemption delay to a longer value can help ensure there is time for session synchronization to be completed before failover.

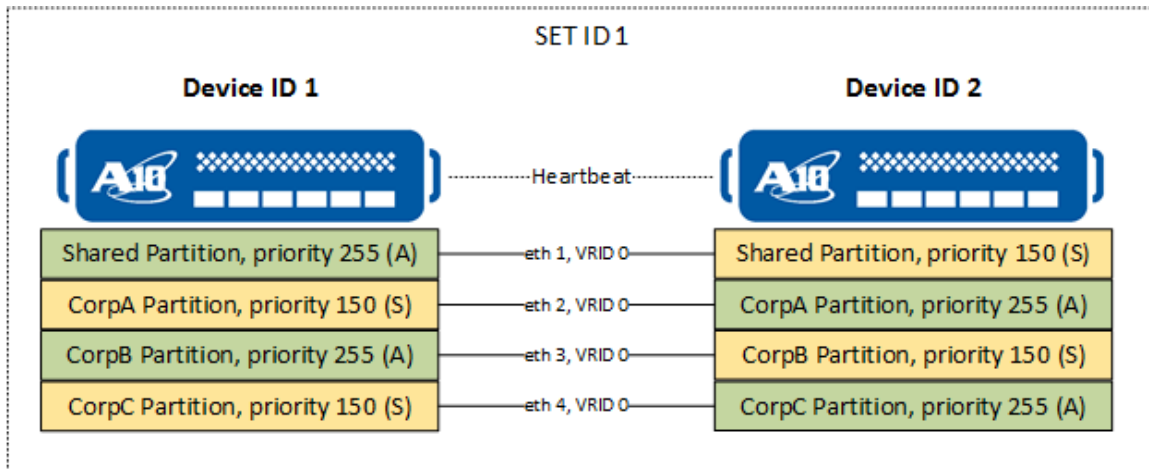
For more information about configuring preemption delays, see [Configuring Preemption Delays](#).

Active Device Selection Examples

One device in the VRRP-A set can be active at any given time for a given partition's VRID. The active device for a given VRID is selected based on VRRP-A priority. The device with the highest VRRP-A priority for a VRID becomes the active device for that

VRID. Each VRID has its own priority value, which can be 1-255. The default is 150. [Figure 5](#) shows an example.

Figure 5 : Selection of Active Device

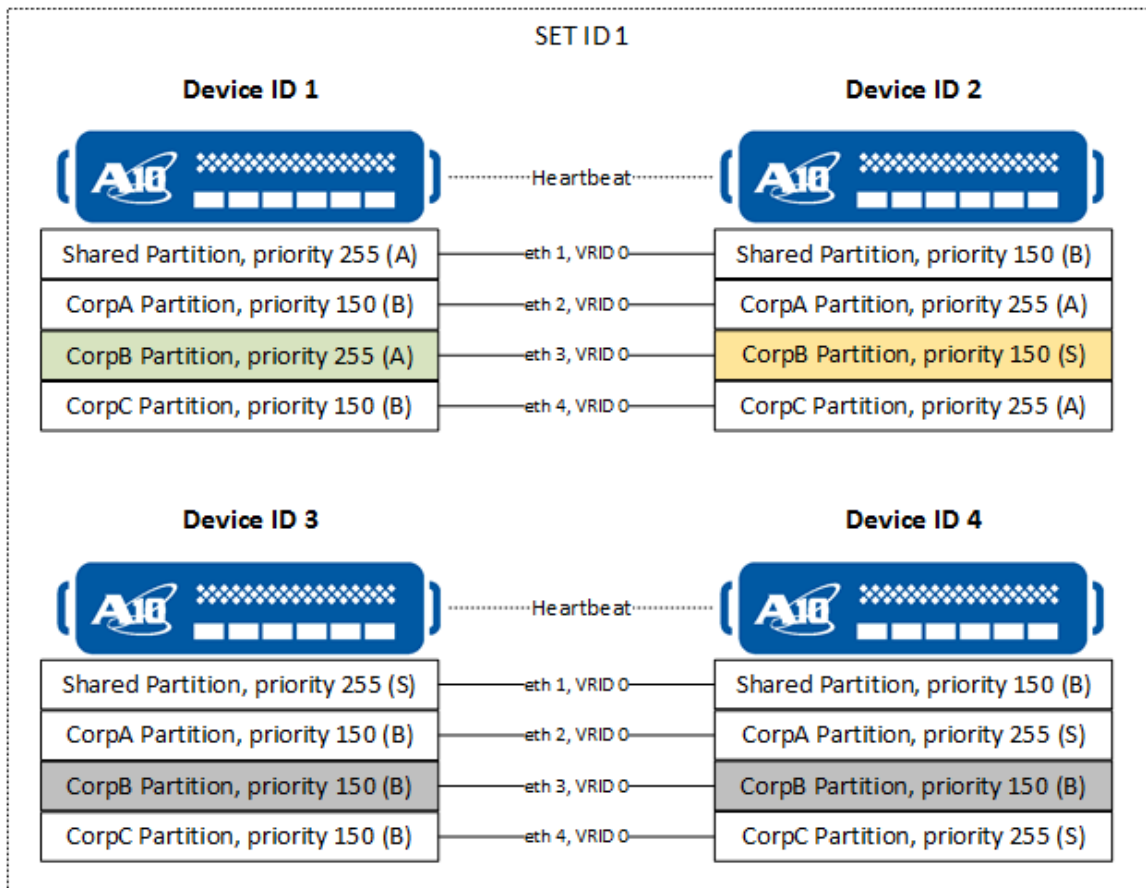


In this example, the priority of VRID 0 is set to 255 for the shared partition and partition CorpB on device 1, and partitions CorpA and CorpC on device 2. The active/standby state of each VRID on each device is based on these priority settings. In the illustration, "A" denotes active and "S" denotes standby.

If more than one device has the highest priority value, the device with the lowest device ID becomes the active device. For example, if the VRID priority is left set to the default value on all devices, device 1 becomes the active device for the VRID.

VRRP-A selects the device that has the second-highest priority for a VRID as the standby device for the VRID. In VRRP-A deployments on more than 2 devices, if more than one device has the second-highest priority value, the device with the lowest device ID becomes the standby device. [Figure 6](#) shows an example.

Figure 6 : Selection of Standby Device

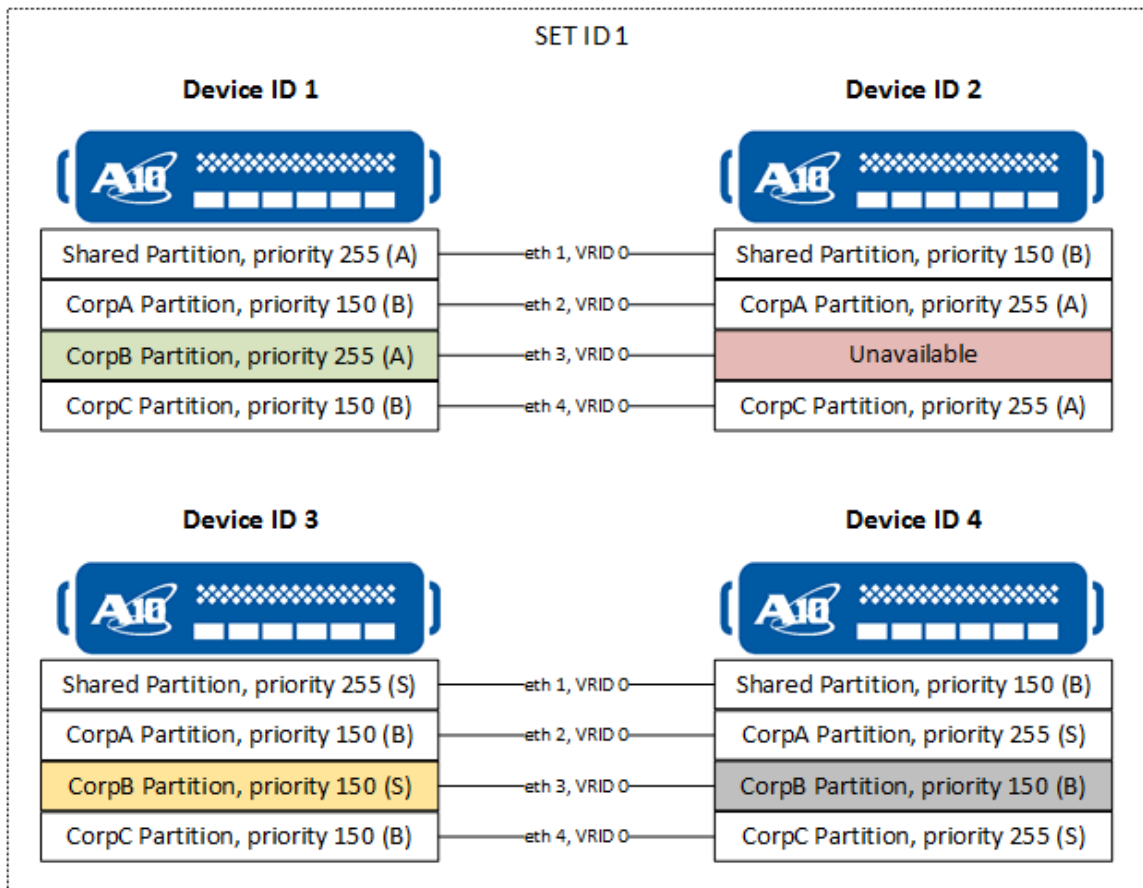


The priority for VRID 0 in partition CorpB is set to 255 on device 1 but is left to its default value on the other devices. Since device 2 is has the lowest device ID number among the remaining devices, device 2 becomes the standby for the VRID (denoted by “S” in the illustration). The other devices are backups (denoted by “B” in the illustration).

If session synchronization is enabled, sessions are copied from device 1 to device 2. If a failover occurs, device 2 becomes the active device and device 1 becomes the standby device. Sessions are not synchronized to the backups.

If the standby device becomes unavailable or its priority value is reduced below the priority value on another backup device, the other device becomes the new standby for the VRID, as shown in [Figure 7](#).

Figure 7 : Selection of New Standby Device



NOTE:

- In VRRP-A deployments of 3 or more devices, moving the active role to a backup other than the standby can cause loss of synchronized sessions, since sessions are synchronized only from the active device to the standby device.
- To avoid loss of sessions in this case, first change the priority on the backup so that it will assume the role of standby. Then, when VRRP-A fails over, the standby will have the synchronized sessions.

VRRP-A Failover

Failover of a VRID from the active ACOS device to the standby ACOS device can be triggered by any of the following events:

- The standby ACOS device stops receiving VRRP-A hello messages from the active ACOS device.
- The VRRP-A priority on the active device is dynamically reduced below the priority on the standby device. The priority can be dynamically reduced when a tracked default gateway, data port, or VLAN goes down, a tracked route is not in the data route table (for more information, see the tracking-options command in the *Command Line Interface Reference*), or a server bound to a service group in a VIP fails its health check.
- The VRRP-A priority on the active device is manually reduced below the priority on the standby device by an administrator, and preemption is enabled.
- The force-self-standby option is used on the active device by an administrator.

Policy-Based Failover

VRRP-A provides flexible event tracking and policy-based failover support through configurable templates. This feature allows policy-based failover even when VRRP-A preemption is disabled and the ACOS device typically would remain in an Active state, despite VRID priority changes. It allows for event-based failover once a tracked event occurs.

For more information, see [VRRP-A Policy-Based Failover Template](#).

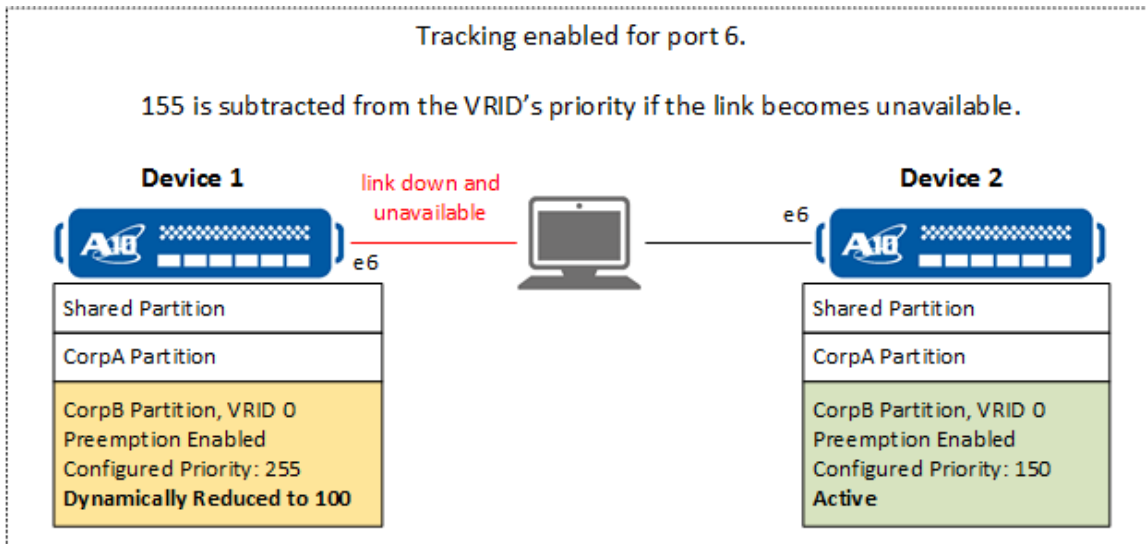
Dynamic Priority Reduction

Depending on the configuration, failover can be triggered even if the active device is still operational. For each VRID, each device begins with a statically configured priority for the VRID.

You can configure the priority of a VRID to be dynamically reduced based on changes in system or network conditions. For example, you can configure link tracking on an

Ethernet port. If the link goes down, the VRID priority on the device is reduced by the configured amount. [Figure 8](#) shows an example.

Figure 8 : Failover Based on Dynamic Priority Reduction



In this example, link tracking is enabled for Ethernet port 6 and configured to reduce the VRID priority by 155 if the link goes down. Device 1 has the higher configured priority value for VRID 0 in partition CorpB and is the active device. However, if the link on Ethernet port 6 goes down, the priority is dynamically reduced below the priority value on device 2. Device 2, therefore, becomes the active device for the VRID.

See [Events Tracked for Weight via the Templates](#) for a summary of events that can be tracked.

By default, none of the dynamic failover triggers are configured. They can be configured on an individual partition basis.

Force-Self-Standby

The force-self-standby option in `vrrp-a force-self-standby` provides a simple method to force a failover, without the need to change VRID priorities and use preemption. For more information about the command, see *Command Line Interface Reference*.

The option can be entered by shared partition admins and private partition admins. If entered in the shared partition, the option applies to all VRIDs on the device. If entered in a private partition, the option applies to all VRIDs within that partition but does not affect VRIDs in other partitions.

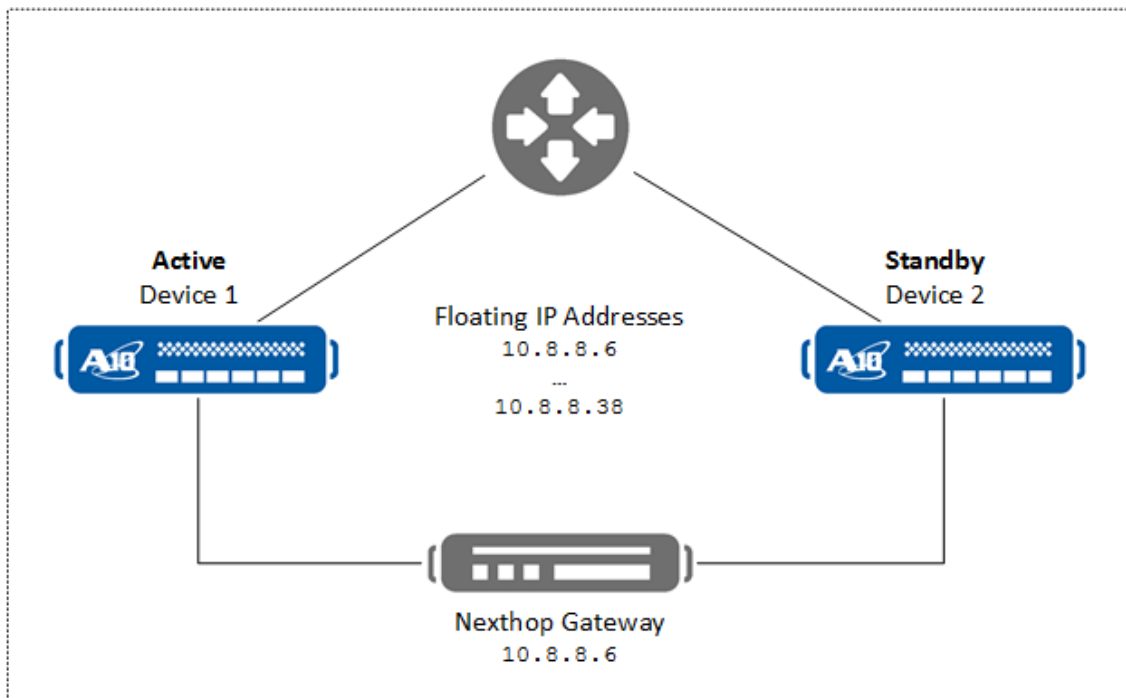
The option remains in effect until one of the other failover triggers occurs or the device is reloaded or rebooted. The option is not added to the configuration and does not persist across reloads or reboots.

VRRP-A and Floating IP Addresses

VRRP-A supports the use of floating IP addresses. Floating IP addresses can reside on any device in the VRRP-A set, and always reside on the device that is currently active. Because floating IP addresses are always reachable regardless of the VRRP-A states of individual devices, the floating IP addresses provide network stability.

In a typical VRRP-A deployment, floating IP addresses are configured for each of the ACOS device interfaces that are used as next-hop interfaces by other devices. [Figure 9](#) shows a simple example.

Figure 9 : Floating IP Address



In this example, a device uses ACOS device IP address 10.8.8.6 as its default gateway. This address is configured as a floating IP address on the ACOS device and always resides on the active VRRP-A device.

Because the address is configured as a floating IP address on the ACOS device, the address remains reachable by the client even if a VRRP-A failover occurs. For example, if VRRP-A fails over from device 1 to device 2, then the floating IP address also moves to device 2.

NOTE:

- A floating IP address cannot be the same as an address that already belongs to a device. For example, the IP address of an ACOS device interface cannot be a floating IP address.
 - A floating IP address can overlap with IP NAT pool IP address and virtual server IP. Smart NAT cannot be enabled when Floating IP and NAT IPs overlap.
 - In the case of CGNAT, floating IP and IP NAT pool cannot be the same IP address.
-

Advertisement of the Floating IP MAC Address After Failover

When a floating IP address moves from one ACOS device to another following failover, the MAC address associated with the floating IP address changes.

To help other devices find a floating IP address following failover, the new active ACOS device sends IPv4 gratuitous ARPs (for an IPv4 floating IP address) or ICMPv6 neighbor advertisements (for an IPv6 floating IP address). The other devices in the network learn the new MAC address from the gratuitous ARPs or neighbor advertisements.

VRRP-A and Configuration Synchronization

VRRP-A supports the following methods of configuration synchronization:

- Automated – Uses ACOS Virtual Chassis System (aVCS) to automatically synchronize the configuration. See “Automated Configuration Synchronization” in *Configuring ACOS Virtual Chassis Systems*.
- Manual – Use the `configure sync` CLI command. For more information, see the *Command Line Interface Reference*.

For more information about viewing the configuration synchronization status, see [Viewing VRRP-A Information](#)

VRRP-A and Session Synchronization

VRRP-A uses one active device and one standby device for a given VRID. If session synchronization (also called connection mirroring) is enabled, the active device backs up active sessions on the standby device.

Session synchronization applies primarily to Layer 4 sessions. Session synchronization does not apply to DNS sessions. Since these sessions are typically very short-lived, there is no benefit to synchronizing them. Likewise, session synchronization does not apply to NATted ICMP sessions or any static NAT sessions. Synchronization of these sessions is not needed since the newly active device will create a new flow for the session following a failover.

Session synchronization is disabled by default and can be enabled on individual virtual ports.

NOTE:

- For configuring VRRP-A for vThunder, the devices should have the same model and resources, such as CPU, Memory, and so on. The vThunder devices must be set to the same DPDK enable or disable mode. If the DPDK state of one device is 'enabled' and that of the other device is 'disabled,' then failover happens for both devices.
 - The session synchronization message is enhanced in ACOS 5.2.1-P4, but it is not understood by all previous ACOS 4.1.4 releases and ACOS 5.2.1-P3.
As a result, when upgrading ACOS, the session synchronization between ACOS 6.0.0 onwards and all ACOS 4.1.4 or ACOS 5.2.1-P3 earlier will not work.
-

To enable session synchronization, in the configuration of a virtual port for a virtual server, add the `ha-conn-mirror` CLI command. This is only allowable on certain port types, and whether the port type is configurable for session synchronization can be identified by doing a query for the `ha-conn-mirror` CLI command after the port type configuration.

```
slb virtual-server VS 192.0.2.0
  port 80 tcp
ha-conn-mirror
```

A specific port for session synchronization can be configured by using the `vrrp-a preferred-session-sync-port ethernet` CLI command.

For more information about the `ha-conn-mirror` and `vrrp-a preferred-session-sync-port ethernet` commands, see the *Command Line Interface Reference*.

VRRP-A Interfaces

VRRP-A sends hello messages and session synchronization messages on the ACOS device's VRRP-A interfaces.

Each ACOS device providing VRRP-A for a VRID must have at least one Up Ethernet interface that can reach the other ACOS devices. If an ACOS device cannot receive hello messages from the other devices for the VRID, the ACOS device's VRRP-A state becomes Active. In this case, there is more than one active VRRP-A device for the same VRID, which is invalid. One of the active VRRP-A devices is the ACOS device that does not have a connection to the other devices. The other active VRRP-A device is the ACOS device that can still reach the other ACOS devices (standby VRRP-A devices).

By default, no VRRP-A interfaces are explicitly configured. In this case, VRRP-A uses all Up Ethernet interfaces to send and listen for hello messages. For an interface that belongs to more than 1 VLAN, the ACOS device uses the lowest VLAN ID to which the interface belongs.

If a data interface has both IPv4 and IPv6 addresses, the primary IPv4 address is used.

Admins with write privileges for the shared partition can explicitly specify the ACOS device interfaces to use as VRRP-A interfaces. In this case, the specified interfaces are used as VRRP-A interfaces by the shared partition and all L3V partitions.

Explicitly Configured VRRP-A Interfaces

Optionally, you can explicitly configure individual interfaces to act as VRRP-A interfaces. In this case, the ACOS device uses only the configured VRRP-A interfaces for hello messages.

For individual L3V partitions, the interface requirement differs depending on whether VRRP-A interfaces are explicitly configured:

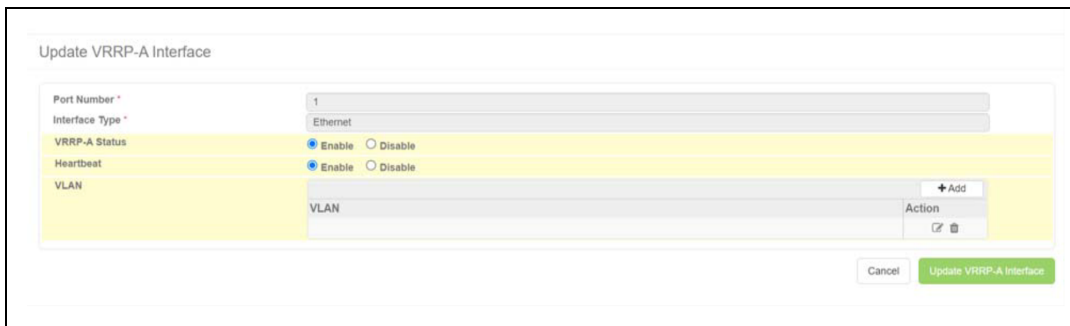
- If no VRRP-A interfaces are explicitly configured (the default situation), each L3V partition must have at least one Ethernet interface that can reach the other L3V partitions for the VRID.
- If at least one interface in the shared partition is explicitly configured as a VRRP-A interface, that interface is used for the shared partition's VRID hello messages and the hello messages of all the VRIDs in the L3V partitions.

Using the GUI to Configure a VRRP-A Interface

To use the GUI to configure a VRRP-A interface:

1. Hover over System in the menu bar, then select VRRP-A.
2. Select the VRRP-A Interface tab, then select the interface type you want to enable from the drop-down list.
3. For the interface, you want to enable VRRP-A, click Edit in the Actions column.
4. On the Update, VRRP-A Interface page, select the options you want to configure for the VRRP-A interface.

You can enable or disable VRRP-A, or set the interface type and state.



The screenshot shows the 'Update VRRP-A Interface' configuration page. It features several input fields and radio buttons:

- Port Number ***: A text input field containing the value '1'.
- Interface Type ***: A dropdown menu with 'Ethernet' selected.
- VRRP-A Status**: Radio buttons for 'Enable' (selected) and 'Disable'.
- Heartbeat**: Radio buttons for 'Enable' (selected) and 'Disable'.
- VLAN**: A text input field containing 'VLAN' and an 'Add' button.
- Action**: A dropdown menu with 'Edit' selected.

At the bottom right, there are two buttons: 'Cancel' and 'Update VRRP-A Interface'.

For specific information about the fields on this screen, refer to the online help.

Using the CLI to Configure a VRRP-A Interface

Use the `vrrp-a interface` command to configure a VRRP-A interface. For example, to configure ethernet interface 2 as the VRRP-A interface:

```
ACOS(config)# vrrp-a interface ethernet 2
ACOS(config-ethernet:2)#
```

Then, you can configure the interface type and state.

For more information, see the `vrrp-a interface` command in the *Command Line Interface Reference*.

Preferred Receive Interface for Session Synchronization

By default, VRRP-A on a backup device automatically selects the Ethernet interface on which to receive synchronized sessions. The ACOS device provides an option that enables you to specify the Ethernet interface(s) on which it is preferred to receive synchronized sessions.

Using the GUI to Configure the Session Synchronization Interface for VRRP-A

1. Hover over System in the menu bar, then select VRRP-A.
2. Click the Settings tab and select Global; this should be the active page by default.
3. Near the bottom of the page, expand the Session Sync Port category to display the configuration area for the feature.
4. Select “Ethernet” from the Interface Type drop-down list.
5. Specify the interface number in the Interface field.
6. If the interface belongs to more than one VLAN, specify the VLAN ID in the VLAN field.
7. Click Add.
8. Repeat for additional interfaces.
9. Click OK.

Session Sync Port

Session Sync Port

Ethernet

Add

Interface Type	Interface	VLAN	
Ethernet	2	1	✕

OK

Using the CLI to Configure the Session Synchronization Interface for VRRP-A

To specify the Ethernet interface on which to receive synchronized sessions, use the `vrrp-a preferred-session-sync-port ethernet` command on the device that you want to receive the session information. For more information about the command, see *Command Line Interface Reference*.

For example, if you want to receive session information on Ethernet port 2 on device ACOS-2:

```
ACOS-2 (config) # vrrp-a preferred-session-sync-port ethernet 2
```

Manually Synchronizing Configurations of All Partitions Between ACOS Devices

The `configure sync` command is used to manually synchronize the running-config and startup-config of all the partitions such as Shared Partition, L3V Partition, and Service Partition from one ACOS device to another ACOS device.

This feature is supported for the ACOS devices that are deployed in VRRP-A or non-VRRP-A environments.

For more information about configuration options, see `configure sync` in the *Command Line Interface Reference* and *System Configuration and Administration Guide*.

Deploying VRRP-A

This chapter provides examples of VRRP-A deployment.

The following topics are covered:

Basic VRRP-A Deployment	31
Force-Self-Standby Reload or Restart Persistence	31
Viewing VRRP-A Information	32
Configuring Preemption Delays	39
VRRP-A Configuration Examples	41
Configuring a Leader and Follower VRID	49
Configuring VRRP-A VRID Lead Switching	51
VRRP-A Status in CLI Prompt	53

Basic VRRP-A Deployment

Basic VRRP-A deployment requires only the following steps on each ACOS device:

1. Specify the VRRP-A device ID and set ID, then enable VRRP-A using the `enable` command:

```
ACOS(config)# vrrp-a common
ACOS(config-common)# device-id 5
ACOS(config-common)# set-id 10
ACOS(config-common)# enable
ACOS-Active(config-common)# exit
ACOS-Active(config)#
```

You can specify a device ID from 1-8, and a set ID from 1-15. Note that the prompt is updated to reflect whether the device is a standby device or the active device. Ensure that each VRRP-A cluster in a Layer 2 broadcast domain has unique set IDs.

NOTE: If aVCS is configured, its device ID is the same as the VRRP-A device ID. If Scaleout is configured and enabled, VRRP-A should NOT be enabled.
Device-IDs for VRRP-A enabled configurations must be in the range of 1-8.

2. As needed, configure floating IP addresses for individual VRIDs.

The following example configures a floating IP address (192.168.9.9) for VRID 13:

```
ACOS-Active(config)# vrrp-a vrid 13
ACOS-Active(config-vrid:13)# floating-ip 192.168.9.9.
```

Force-Self-Standby Reload or Restart Persistence

If the force-self-standby option is selected via the CLI, this feature will allow a single partition, multiple partitions, or the device to remain in a forced self-standby state even though the device or partition goes through a reload or device restart.

To place VRRP-A VRID 2 into a forced-self-standby state, issue the following command:

```
ACOS(config)# vrrp-a force-self-standby vrid 2 enable
```

To place VRID 2 in a self-standby state even after a restart or reload, use the following command:

```
ACOS(config)# vrrp-a force-self-standby-persistent vrid 2  
Doing so may cause some of the VRID(s) to become inactive in the whole  
vrrp-a set.  
WARNING: Please confirm that you want to perform this operation.  
Doing so may cause some of the VRID(s) to become inactive in the whole  
vrrp-a set; make sure you verify that an active device will be available  
after performing this operation. [yes/no]: yes
```

If you issue the `vrrp-a force-self-standby` command, the partition or device will not be placed in a force-self-standby state after the device reloads or restart completes. You must configure self-standby with `vrrp-a force-self-standby-persistent` to have your device retain its self-standby state after a future reload or restart operation. For more information about the commands, see *Command Line Interface Reference*.

To place all VRIDs in partition (partA) in the self-standby state even after a restart or reload, execute the following command in partition “partA:”

```
ACOS[partA](config)# vrrp-a force-self-standby enable  
Doing so may cause some of the VRID(s) to become inactive in the whole  
vrrp-a set.  
WARNING: Please confirm that you want to perform this operation.  
Doing so may cause some of the VRID(s) to become inactive in the whole  
vrrp-a set; make sure you verify that an active device will be available  
after performing this operation. [yes/no]: yes
```

Viewing VRRP-A Information

To view VRRP-A configuration information and statistics, use the VRRP-A Show Commands. For more information about the commands, see *Command Line Interface Reference*.

To view VRRP-A synchronization status information, see [Viewing VRRP-A Config-Sync Status in the CLI](#).

Viewing VRRP-A Config-Sync Status in the CLI

For config-sync status, use the `show config-sync` command.

In the following examples, ACOS1 is the local device and ACOS2 and ACOS3 are the peers.

Config-Sync Status for New and Reloaded/Rebooted Devices

In the case of a new device with no config-sync configuration, or when a device is reloaded or rebooted, the sync status of the running-config will always be reset and the sync status will be “N/A”. For example:

```
ACOS1(config)# show config-sync detail
Partition Name      Sync Status for running-config and startup-config
-----
-----
shared              (running-config)N/A
shared              (startup-config)N/A
shared              (running-config)N/A
shared              (startup-config)N/A
```

This means that if you already have config-sync configured and the device is rebooted, you will have to reconfigure the config-sync settings in the running-config. The sync status of the startup-config persists through any reload or reboot operation; and will not be changed by the reload or reboot process.

Config-Sync Status After HA Sync

The following commands sync both the running-config and startup-config for the shared partition from the local device to the peer, then view the config-sync status:

```
ACOS1(config)# configure sync all partition shared 192.168.216.202
Running config will be sync'ed, but it has been changed.
Do you want to save before sync'ing? [y/n]:y
Running-config saved.
User name []? admin
Password []?
```

```

ACOS1(config)#
ACOS1(config)# show config-sync
Partition Name      Sync Status for running-config and startup-config
-----
shared              (running-config) sync to 192.168.216.201 at 04:30:36 IST
Fri Mar 5 2021
shared              (startup-config) sync to 192.168.216.201 at 04:30:36 IST
Fri Mar 5 2021
shared              (running-config) sync to 192.168.216.203 at 04:30:55 IST
Fri Mar 5 2021
shared              (startup-config) sync to 192.168.216.203 at 04:30:55 IST
Fri Mar 5 2021

```

The output from the peer device (ACOS2) shows that both the running-config and startup-config are synced from ACOS1:

```

ACOS2(config)# show config-sync
Partition Name      Sync Status for running-config and startup-config
-----
shared              (running-config) is synced from ip 192.168.216.201 at
20:31:54 IST Wed May 18 2016
shared              (startup-config) is synced from ip 192.168.216.201 at
20:32:13 IST Wed May 18 2016

```

Config-Sync Status After Changes are Made in the Running-Config

If any configuration changes are made in the running-config on ACOS1, the status will change. Suppose we make a configuration change on ACOS1 (for example, configure a VRRP-A failover policy template), then view the config-sync status again:

```

ACOS1(config)# vrrp-a fail-over-policy-template temp
ACOS1(config-failover-policy)# show config-sync
Partition Name      Sync Status for running-config and startup-config
-----
shared              (running-config) not synced to 192.168.216.201 because
it's changed at 05:54:56 IST Fri Mar 5 2021

```

```
shared          (startup-config) sync to 192.168.216.201 at 04:30:36 GMT
Fri Mar 5 2021
shared          (running-config) not synced to 192.168.216.203 because
it's changed at 05:54:56 IST Fri Mar 5 2021
shared          (startup-config) sync to 192.168.216.203 at 04:30:55 IST
Fri Mar 5 2021
```

The status is changed to “not synced” because of the change made in the running-config. After running the `write memory` command, the status appears as shown below:

```
ACOS1(config)# write memory
Building configuration...
Write configuration to primary default startup-config
[OK]
ACOS1(config)# show config-sync
Partition Name   Sync Status for running-config and startup-config
-----
shared          (running-config) not synced to 192.168.216.201 because
it's changed at 05:54:56 GMT Fri Mar 5 2021
shared          (startup-config) not synced to 192.168.216.201 because
write memory at 05:58:12 GMT Fri Mar 5 2021
shared          (running-config) not synced to 192.168.216.203 because
it's changed at 05:54:56 GMT Fri Mar 5 2021
shared          (startup-config) not synced to 192.168.216.203 because
write memory at 05:58:12 GMT Fri Mar 5 2021
```

The sync status of the startup-config is now “not synced” because of the change.

Config-Sync Status When Synced From the Peer Device

Suppose that later on, we sync the running-config from ACOS2 to ACOS1. Below is the `detail` output from the local device (ACOS1) after this happens:

```
ACOS1(config)# show config-sync detail
Partition Name   Sync Status for running-config and startup-config
-----
shared          (running-config) sync to ip 192.168.216.202 at 08:19:37
IST Fri May 13 2016
shared          (startup-config)N/A
```

```
shared          (running-config) is synced from ip 192.168.216.202 at
18:40:39 IST Mon May 16 2016
shared          (startup-config)N/A
```

Viewing Detailed Config-Sync Status

The **detail** option shows when the “sync to” status changed to the “is synced from” status, or vice versa.

In this example, the configuration is synced from ACOS2 to ACOS1:

```
ACOS2(config)# configure sync all partition shared 192.168.216.201
User name []? admin
Password []?
ACOS2(config)#
```

The **show config-sync** command on ACOS1 will now show the “is synced from” status as the most recent sync status:

```
ACOS1(config)# show config-sync
Partition Name   Sync Status for running-config and startup-config
-----
shared          (running-config) is synced from ip 192.168.216.202 at
04:31:00 IST Fri Mar 5 2021
shared          (startup-config) is synced from ip 192.168.216.202 at
04:31:02 IST Fri Mar 5 2021
```

The **detail** option can be used to view information about when the sync status was changed:

```
ACOS1(config)# show config-sync detail
Partition Name   Sync Status for running-config and startup-config
-----
shared          (running-config) not synced because it's changed at
20:36:50 IST Wed May 18 2016
shared          (startup-config) not synced because write memory at
20:37:50 IST Wed May 18 2016
shared          (running-config) is synced from ip 192.168.216.202 at
20:41:45 IST Wed May 18 2016
```

```
shared          (startup-config) is synced from ip 192.168.216.202 at
20:41:59 IST Wed May 18 2016
```

The above output shows the previous sync status of both the running-config and startup-config on ACOS1 before the sync from ACOS2 to ACOS1.

For more information, see `show config-sync` in the *Command Line Interface Reference*.

Viewing Peer Box Information

The VRRP-A session sync works across different platforms and AX platform boxes. Now user uses `vrrp- a hello` packet TLV to get the peer box's information, such as blades number, version and platform model, CPU number of each blade. The number of boards and number of CPUs of each board could be different between the active box and standby box. This scenario is any platform to any platform. It includes variable chassis or non-chassis platforms, such as active is a non-chassis box, standby is a chassis platform box, or vice versa. This also includes cases with different CPU numbers on each box.

NOTE: On the active box, in the VRRP-A's API, a source port value is automatically filled in the UDP header so that session gets synced to the expected target board and target CPU.

Steps to get target board, target CPU information, and sync a session:

1. The system API gives information about the board number and CPU number of each board of the standby box. The board number and CPU number are provided by the standby box system API, and VRRP-a exchanges this information among each peer by VRRP-a hello packets. VRRP-a hello API return with two parameter "`target_blade_index`" and "`target_cpu_index`".
2. The values "`target_blades_number`" and "`target_cpu_number`" can be obtained by calling API. This platform information belongs to the target standby box. Different application teams will implement parts for their session type.
3. On the VRRP-a infra side, after the information mentioned above is fetched, the corresponding value is filled in the UDP header.

For example:

```
Blade 1 : #define HA_SYNC_UDP_PORT_BOARD1      0x5700
Blade 2 : #define HA_SYNC_UDP_PORT_BOARD2      0xD700
Blade 3 : #define HA_SYNC_UDP_PORT_BOARD3      0xD800
Blade 4 : #define HA_SYNC_UDP_PORT_BOARD4      0xD900
```

4. The VRRP-a hello packet TLV mechanism is used to get the peer box's blade number, version, and CPU number of each blade.

For example:

- a. The TLV after the header:

```
struct _vrrp_tlv {
UINT8 type;
        UINT8 subtype;
        UINT16 len;
        UINT8 data[1];
} __attribute__((__packed__));
```

- b. Data structure for the platform contains the following information:

```
new Type VRRP_TYPE_PLATFORM_INFO
enum {
    VRRP_TYPE_VRID = 1,
    VRRP_TYPE_ACTIVE_PEER,
    VRRP_TYPE_PBF_INFO,
    VRRP_TYPE_SERIAL_NUM,
    VRRP_TYPE_INLINE,
    VRRP_TYPE_PLATFORM_INFO
};
```

- c. Data structure contains the following information:

- struct vrrp_cpu_board_info
- struct _vrrp_platform_info_msg_entry
- struct vrrp_cpu_board_info

Configuring Preemption Delays

The following topics are covered:

- [Use the GUI to Configure VRRP-A Preemption Delay](#)39
- [Use the CLI to Configure VRRP-A Preemption Delay](#)39

Use the GUI to Configure VRRP-A Preemption Delay

1. Navigate to **System > VRRP-A > Global**.
2. Edit the value in the Preemption Delay field, to a value from 1-255. The default VRRP-A preemption delay is 6 seconds (60 units of 100 ms). You can globally set the delay to a value from 100 ms (1 unit of 100 ms) to 25.5 seconds (255 units of 100 ms).
3. Click **OK**.

Use the CLI to Configure VRRP-A Preemption Delay

To configure the VRRP-A preemption delay, use the `preemption-delay` command at the VRRP-A common configuration level of the CLI. For more information about the command, see *Command Line Interface Reference*.

In the example below, the `show vrrp-a` command reveals three devices in Active-Active mode:

```
ACOS# show vrrp-a
vrid 0
Unit          State          Weight          Priority
1 (Local)     Active         65534           200
              became Active at: Aug 5 05:09:20 2015
              for 0 Day, 5 Hour,10 min
2 (Peer)      Standby        65534           180
3 (Peer)      Standby        65534           160
vrid 1
Unit          State          Weight          Priority
1 (Local)     Standby        65534           160
              became Standby at: Aug 4 06:23:39 2015
```

```

                                for 1 Day, 3 Hour, 56 min
2 (Peer)      Active           65534           200
3 (Peer)      Standby          65534           180          *
vrid 2
Unit          State            Weight           Priority
1 (Local)     Standby          65534           150
              became Standby at: Aug 4 06:23:38 2015
                                for 1 Day, 3 Hour, 56 min
2 (Peer)      Standby          65534           160          *
3 (Peer)      Active           65534           200
vrid that is running: 0 1 2

```

To see how preemption delay works, suppose we track route 110.0.0.0 255.255.255.0:

```

ACOS# show run | sec vrrp-a vrid 0

vrrp-a vrid 0
  blade-parameters
    priority 200
  tracking-options
    route 110.0.0.0 255.255.255.0 priority-cost 255

vrrp-a vrid 0
  blade-parameters
    priority 200
  tracking-options
    route 110.0.0.0 255.255.255.0 priority-cost 255
  blade-parameters
    priority 180
  tracking-options
    route 110.0.0.0 255.255.255.0 priority-cost 255
  blade-parameters
    priority 160
  tracking-options
    route 110.0.0.0 255.255.255.0 priority-cost 255

```

An interface becoming unavailable would also cause the route to disappear and then reappear. Without preemption delay, the following sequence of events would occur (all times given are approximate and will depend on your exact configuration):

- After the interface goes down, the route disappears (less than 1 second)
- The ACOS device notices the missing route and lowers the device priority for the VRID (less than 1 second).
- Failover occurs (less than 1 second)
- The route is restored (several seconds, depending on your specific configuration)
- The devices notice the route and restore the original device priority for the VRID (less than 1 second)
- Preemption happens 8 seconds after the original device priority is restored

In the same scenario, with preemption delay configured for 25.5 seconds:

```
ACOS# show run | sec vrrp-a common
vrrp-a common
  device-id 1
  set-id 1
  enable
  preemption-delay 255
```

Below is the timing of the same sequence of events:

- After the interface goes down, the route disappears (less than 1 second).
- The ACOS device notices the missing route and lowers the device priority for the VRID (less than 1 second).
- Failover occurs (less than 1 second).
- The route is restored (several seconds, depending on your specific configuration).
- The devices notice the route and restore the original device priority for the VRID (less than 1 second).
- Preemption happens 25 seconds after the original device priority is restored.

VRRP-A Configuration Examples

The following sections show configuration examples for VRRP-A. The first example shows a very basic deployment using the minimum required commands. The second example includes priority configuration changes and tracking.

Simple Deployment

The following commands deploy VRRP-A on a set of 2 ACOS devices.

Commands on ACOS-1

Enter the following commands on ACOS-1:

```
ACOS-1(config)# vrrp-a common
ACOS-1(config-common)# device-id 1
ACOS-1(config-common)# set-id 1
ACOS-1(config-common)# enable
```

Commands on ACOS-2

Enter the following commands on ACOS-2:

```
ACOS-2(config)# vrrp-a common
ACOS-2(config-common)# device-id 2
ACOS-2(config-common)# set-id 1
ACOS-2(config-common)# enable
```

Deployment with Custom Priority Settings

The following commands configure VRRP-A on 2 ACOS devices, with the priority settings shown in [Selection of Active Device](#).

- On the device ACOS-1, the priority value is set to 255 on the shared partition's VRID 0 and partition CorpB's VRID 0. The priority value is left set to its default (150) for CorpA and CorpC; the priority does not need to be set since this is the default value.
- On the device ACOS-2, the priority value is set to 255 on the VRID 0s in partitions CorpA and CorpC. The priority value is left set to its default (150) for the shared partition and CorpB; the priority does not need to be set since this is the default value.

These commands also configure a floating IP address for each VRID. [Selection of Active Device](#) does not include the floating IP addresses. (For more on floating IP addresses, see [VRRP-A and Floating IP Addresses](#).)

Commands on ACOS-1

```
ACOS-1 (config) # vrrp-a common
ACOS-1 (config-common) # device-id 1
ACOS-1 (config-common) # set-id 1
ACOS-1 (config-common) # enable
ACOS-1 (config-common) # exit
ACOS-1 (config) # vrrp-a vrid 0
ACOS-1 (config-vrid:0) # floating-ip 10.10.10.2
ACOS-1 (config-vrid:0) # blade-parameters
ACOS-1 (config-vrid:0-blade-parameters) # priority 255
ACOS-1 (config-vrid:0-blade-parameters) # exit
ACOS-1 (config-vrid:0) # exit
ACOS-1 (config) # active-partition CorpB
Current active partition: CorpB
ACOS-1 [CorpB] (config) # vrrp-a vrid 0
ACOS-1 [CorpB] (config-vrid:0) # floating-ip 30.30.30.2
ACOS-1 [CorpB] (config-vrid:0) # blade-parameters
ACOS-1 [CorpB] (config-vrid:0-blade-parameters) # priority 255
ACOS-1 [CorpB] (config-vrid:0-blade-parameters) # exit
ACOS-1 [CorpB] (config-vrid:0) # exit
ACOS-1 [CorpB] (config) #
ACOS-1 (config) # active-partition CorpA
Current active partition: CorpA
ACOS-1 [CorpA] (config) # vrrp-a vrid 0
ACOS-1 [CorpA] (config-vrid:0) # floating-ip 20.20.20.2
ACOS-1 [CorpA] (config-vrid:0) # exit
ACOS-1 [CorpA] (config) #
ACOS-1 (config) # active-partition CorpC
Current active partition: CorpC
ACOS-1 [CorpC] (config) # vrrp-a vrid 0
ACOS-1 [CorpC] (config-vrid:0) # floating-ip 40.40.40.2
ACOS-1 [CorpC] (config-vrid:0) # exit
ACOS-1 [CorpC] (config) #
```

Commands on ACOS-2

```
ACOS-2 (config) # vrrp-a common
ACOS-2 (config-common) # device-id 2
ACOS-2 (config-common) # set-id 1
ACOS-2 (config-common) # enable
```

```
ACOS-2 (config-common) # exit
ACOS-2 # active-partition CorpA
Current active partition: CorpA
ACOS-2 [CorpA] # config
ACOS-2 [CorpA] (config) # vrrp-a vrid 0
ACOS-2 [CorpA] (config-vrid:0) # floating-ip 20.20.20.2
ACOS-2 [CorpA] (config-vrid:0) # blade-parameters
ACOS-2 [CorpA] (config-vrid:0-blade-parameters) # priority 255
ACOS-2 [CorpA] (config-vrid:0-blade-parameters) # exit
ACOS-2 [CorpA] (config-vrid:0) # exit
ACOS-2 [CorpA] (config) # active-partition CorpC
Currently active partition: CorpC
ACOS-2 [CorpC] (config) # vrrp-a vrid 0
ACOS-2 [CorpC] (config-vrid:0) # floating-ip 40.40.40.2
ACOS-2 [CorpC] (config-vrid:0) # blade-parameters
ACOS-2 [CorpC] (config-vrid:0-blade-parameters) # priority 255
ACOS-2 [CorpC] (config-vrid:0-blade-parameters) # exit
ACOS-2 [CorpC] (config-vrid:0) # exit
ACOS-2 [CorpC] (config) # active-partition CorpB
Current active partition: CorpB
ACOS-1 [CorpB] (config) # vrrp-a vrid 0
ACOS-1 [CorpB] (config-vrid:0) # floating-ip 30.30.30.2
ACOS-1 [CorpB] (config-vrid:0) # exit
ACOS-1 [CorpB] (config) #
```

Configuration of Tracking Options

The following commands configure VRRP-A tracking options. For simplicity, commands for only a single device are shown.

For more information about these commands, see *Command Line Interface Reference*.

NOTE: When finished configuring specific tracking options, use the `exit` or `end` CLI commands to return to the global configuration mode.

Example of using the `exit` command to reach the global configuration mode:

```
ACOS (config-vrid:1-blade-parameters) # exit
ACOS (config-vrid:1) # exit
```

```
ACOS(config)#
```

Ethernet Interface Tracking

The following command configures tracking of Ethernet port 1. If the port's link goes down, 100 is subtracted from the VRID's priority value.

```
ACOS(config)# vrrp-a vrid 1  
ACOS(config-vrid:1)# blade-parameters  
ACOS(config-vrid:1-blade-parameters)# tracking-options  
ACOS(config-vrid:1-blade-parameters-track...)# interface ethernet 1  
priority-cost 100
```

Trunk Tracking

The following commands configure the ACOS device to track the status of a trunk (but not ports within the trunk). The VRRP-A protocol will reduce the priority of the device by 100 if the trunk goes down.

```
ACOS(config)# vrrp-a vrid 1  
ACOS(config-vrid:1)# blade-parameters  
ACOS(config-vrid:1-blade-parameters)# tracking-options  
ACOS(config-vrid:1-blade-parameters-track...)# trunk 1 priority-cost 100
```

The following commands configure the ACOS device to track the status of a trunk, as well as the status of ports within the trunk. If trunk 2 goes down, the priority of the associated VRID is reduced by 150. If a port within the trunk goes down, the priority of the associated VRID is reduced by 40.

```
ACOS(config)# vrrp-a vrid 2  
ACOS(config-vrid:2)# blade-parameters  
ACOS(config-vrid:2-blade-parameters)# tracking-options  
ACOS(config-vrid:2-blade-parameters-track...)# trunk 2 priority-cost 150  
per-port-pri 40
```

VLAN Tracking

The following command configures tracking of VLAN 69. If the VLAN is quiet for more than 30 seconds, 50 is subtracted from the VRID priority.

```
ACOS(config)# vrrp-a vrid 1  
ACOS(config-vrid:1)# blade-parameters  
ACOS(config-vrid:1-blade-parameters)# tracking-options
```

```
ACOS(config-vrid:1-blade-parameters-track...)# vlan 69 timeout 30  
priority-cost 50
```

For more information about enhanced VLAN tracking capabilities, see [Increased VLANs for VRRP-A Failover Tracking](#).

Gateway Tracking

The following commands configure tracking of gateway 10.10.10.1. If the gateway stops responding to pings, 100 is subtracted from the VRID's priority value.

```
ACOS(config)# health monitor gatewayhml  
ACOS(config-health:monitor)# method icmp  
ACOS(config-health:monitor)# exit  
ACOS(config)# slb server gateway1 10.10.10.1  
ACOS(config-real server)# health-check gatewayhml  
ACOS(config-real server)# exit  
ACOS(config)# vrrp-a vrid 0  
ACOS(config-vrid:0)# blade-parameters  
ACOS(config-vrid:0-blade-parameters)# tracking-options  
ACOS(config-vrid:0-blade-parameters-track...)# gateway 10.10.10.1  
priority-cost 100
```

Route Tracking

The following command configures tracking of routes to destination network 3000::/64. If the IPv6 route table does not contain any routes to the destination, 105 is subtracted from the VRID priority.

```
ACOS(config)# vrrp-a vrid 0  
ACOS(config-vrid:0)# blade-parameters  
ACOS(config-vrid:0-blade-parameters)# tracking-options  
ACOS(config-vrid:0-blade-parameters-track...)# route 3000::/64 priority-  
cost 105
```

The following commands configure tracking of routes to destination network 5000::/64. If the IPv6 route table does not contain any routes to the destination, 80 is subtracted from the VRID priority.

```
ACOS(config)# vrrp-a vrid 1  
ACOS(config-vrid:1)# blade-parameters  
ACOS(config-vrid:1-blade-parameters)# tracking-options
```

```
ACOS(config-vrid:1-blade-parameters-track...)# route 5000::/64 priority-  
cost 80
```

Preemption Delay

The following commands configure the desired preemption delay value of 200 milliseconds:

```
ACOS(config)# vrrp-a common  
ACOS(config-common)# preemption-delay 200
```

For more information about preemption delay, see [Preemption Delay](#).

Increased VLANs for VRRP-A Failover Tracking

The VRRP-A failover tracking feature supports an increased number of VLANs. The weight (as assigned using failover policy templates—for details refer to [VRRP-A Policy-Based Failover Template](#)) or priority assigned to the VLAN tracked event will be used in calculations that will impact VRRP-A failover decisions.

If VRRP-A stops detecting traffic on a VLAN, VRRP-A reduces the priority for the VRID. The priority value to subtract can be specified individually for each VLAN.

To configure VLAN tracking use the following command:

```
vlan vlan-id timeout timeout-value priority-cost value
```

ACOS provides two ways to track VLANs for VRRP-A:

- Using the priority tracking option—The tracking option performs failover by decreasing a VRID's priority value.
- Using failover policy templates—The failover policy template performs failover by decreasing a VRID's weight value.

You can track a total of 64 different VLANs in three possible ways:

- Configure all 64 VLANs that can be tracked using tracking options.
- Configure all 64 VLANs that can be tracked using a single failover policy template or split over multiple templates.
- Configure a mix of both options mentioned above. That is configure 32 VLANs using the tracking option and configure the remaining 32 VLANs using a failover policy template.

NOTE: Tracking options and the failover policy template can track the same VLANs. For example, you configure 64 VLANs using the tracking option. You can also configure tracking of 64 of the *same* VLANs using a fail-over-policy template. In essence, you are still only tracking 64 VLANs.

VLAN Tracking Configuration Example

After you have enabled VRRP-A on the devices, set the VRRP-A VRID, and then set VLAN tracking options for 11 existing VLANs based on their priority cost:

```
vrrp-a vrid 31
  blade-parameters
    tracking-options
      vlan 1000 timeout 30 priority-cost 1
      vlan 1001 timeout 30 priority-cost 1
      vlan 1002 timeout 30 priority-cost 1
      vlan 1003 timeout 30 priority-cost 1
      vlan 1004 timeout 30 priority-cost 1
      vlan 1005 timeout 30 priority-cost 1
      vlan 1006 timeout 30 priority-cost 1
      vlan 1007 timeout 30 priority-cost 1
      vlan 1008 timeout 30 priority-cost 1
      vlan 1009 timeout 30 priority-cost 1
      vlan 1010 timeout 30 priority-cost 1
```

The following example shows how 11 VLANs of differing timeout values and weights are being tracked in a failover template called test.

```
!
vrrp-a fail-over-policy-template test
  vlan 222 timeout 50 weight 100
  vlan 223 timeout 30 weight 200
  vlan 224 timeout 40 weight 160
  vlan 225 timeout 80 weight 105
  vlan 226 timeout 65 weight 118
  vlan 227 timeout 43 weight 145
  vlan 228 timeout 67 weight 120
  vlan 229 timeout 59 weight 156
  vlan 230 timeout 33 weight 66
  vlan 231 timeout 82 weight 90
  vlan 232 timeout 98 weight 75
```

```
vlan 233 timeout 72 weight 64
!
```

Configuring a Leader and Follower VRID

Configuration Example 1

To configure a leader VRID, do the following:

1. In the shared partition, configure a vrid-lead:

```
ACOS(config) #vrrp-a vrid-lead lead1
ACOS(config-vrid-lead:lead1)# partition shared vrid 1
```

2. In the L3V partition, configure the VRID you wish to designate as the follower.

- a. To do so, enter the partition:

```
ACOS# active-partition p1
```

- b. In the currently active partition, p1, enter the following command:

```
ACOS[p1](config)# vrrp-a vrid 1
ACOS[p1](config-vrid:1)# follow vrid-lead lead1
```

- c. Use the [show vrrp-a](#) command to view your configuration:

```
ACOS[p1](config-vrid-follow)# show vrrp-a
vrid 0
Unit          State          Weight          Priority
1 (Local)     Active         65534           150
              became Active at: May 20 12:04:05 2013
              for 0 Day,22 Hour,54 min
2 (Peer)      Standby        65534           150
*
vrid 1 (follow vrid-lead lead1)
Unit          State          Weight          Priority
1 (Local)     Active         65534           150
              became Active at: May 20 12:04:05 2013
              for 0 Day,22 Hour,54 min
```

```

2 (Peer)          Standby          65534          150
*
vrid that is running: 0 1

```

3. By default, the default-vrid-lead exists in the shared partition:

```

vrrp-a vrid-lead default-vrid-lead
partition shared vrid 1

```

Configuration Example 2

To configure a leader in the shared partition, issue the following commands:

1. Define a VRID called **myname** as the leader:

```

ACOS(config)# vrrp-a vrid-lead myname
ACOS(config-vrid-lead:myname)# partition p1 vrid 0

```

2. Use the following show command to display your VRID leader configuration:

```

ACOS(config)# show vrrp-a vrid-lead
vrrp-a vrid-lead default-vrid-lead
partition shared vrid 0
vrrp-a vrid-lead myname
partition p1 vrid 0

```

3. In the L3V partition, configure the VRID to follow the lead VRID:

```

ACOS[p1](config)# vrrp-a vrid 3
ACOS[p1](config-vrid:3)# follow vrid-lead myname

```

4. Observe your configuration using the following command:

```

ACOS[p1](config)# show running-config | sec vrrp-a
vrrp-a vrid 3
follow vrid-lead myname

```

5. As the Admin, if you want to change the VRID's role from being a follower to standalone VRID again, issue the following command:

```

ACOS[p1](config)# vrrp-a vrid 3
ACOS[p1](config-vrid: 3)# no follow vrid-lead default-vrid-lead
To view your configuration, use the show run command:
ACOS[p1](config)# show running-config | vrrp-a

```

```
!
vrrp-a vrid 3
!
```

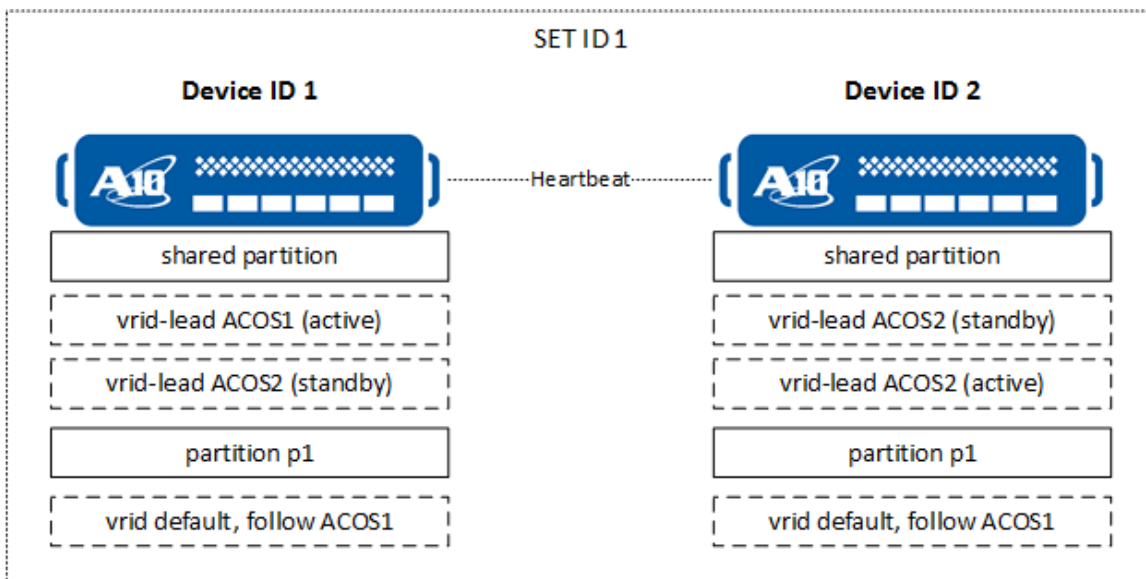
If you try to remove a VRID that is configured as a lead VRID and that has other VRIDs that follow it, you will not be allowed to remove it. You must remove the VRIDs that are configured as followers before you remove the lead VRID.

Configuring VRRP-A VRID Lead Switching

This section describes how to switch lead VRIDs in a VRRP-A configuration without affecting traffic when doing so.

The example in this section utilizes the topology shown in [Figure 10](#):

Figure 10 : VRRP-A VRID Lead Switching Example Topology



Suppose you want to change the default VRID (VRID 0) on partition p1 to follow ACOS2. Using the `vrrp-a vrid 0 follow vrid-lead-ACOS2` command could accomplish this, however, there may be a disruption of the traffic on partition p1's VRID because of the time interval required for manual configuration on both sides. During this time interval, the VRID 0 could be double active or double standby on both devices.

This section provides alternative solutions that do not affect traffic:

- [Manual VRID Lead Switching Using Manual Configuration](#)
- [Automatic VRID Lead Switching Using aVCS](#)

Manual VRID Lead Switching Using Manual Configuration

To manually configure VRID lead switching without affecting traffic:

1. Use either an existing VRID whose up or down status is the same as the VRID you are configuring, or create a new VRID. In the steps below, the example is creating a new VRID called `vrid 5`. The following shows ACOS1, do the same on ACOS2.

```
ACOS1-Active(config)# vrrp-a vrid 5
Doing so may cause some of the VRID(s) to become inactive in the whole
vrrp-a set.
WARNING: Please confirm that you want to perform this operation.
Doing so may cause some of the VRID(s) to become inactive in the whole
vrrp-a set; make sure you verify that an active device will be
available after performing this operation. [yes/no]: yes
```

2. Ensure that the new VRID's active device is the same as the current VRID (in this example, the active device is ACOS1). You can accomplish this by manually setting the dynamic priority of the new VRID on ACOS2 to a lower number than on ACOS1, or by using the `vrrp-a force-self-standby vrid vrid#` command for the new VRID on ACOS2.

```
ACOS2(config)# vrrp-a force-self-standby vrid 5 enable
```

3. Create the VRID lead that you will use in the next step.

```
ACOS1-Active(config)# vrrp-a vrid-lead lead1
ACOS1-Active((config-vrid-lead:lead1)# partition shared vrid 5
```

4. In partition p1 on both devices, use the `follow vrid-lead vrid_lead_template` command to make the traffic in the partition follow the new VRID. Since the new VRID has the same VRRP-A status as the existing lead VRID, traffic is not impacted during this step. The following shows ACOS1, do the same on ACOS2.

```
ACOS1(config)# active-partition p1
Current active partition: p1
ACOS1-Active[p1](config)#vrrp-a vrid 0
ACOS1-Active[p1](config-vrid:5)#follow vrid-lead lead1
ACOS1-Active[p1](config-vrid:5)#exit
```

5. Change the dynamic priority of the new VRID on ACOS1, or the `vrrp-a force-self-standby vrid vrid# disable` command for the new VRID on ACOS2, thus causing the new VRID to become active on ACOS2. This switch is done by message negotiation among peer devices, so no traffic is lost in this step.

```
ACOS2(config)# vrrp-a force-self-standby vrid 5 disable
```

Automatic VRID Lead Switching Using aVCS

aVCS uses almost real-time config sync across devices in the same virtual cluster. This means when the VRID lead has switched the configurations on both devices are updated almost simultaneously so that traffic is preserved.

VRRP-A Status in CLI Prompt

The CLI prompt can be customized to identify the devices that are running VRRP-A configuration.

For information about how to do this, see “VRRP-A/aVCS Status in Command Prompt” in the *Command Line Interface Reference*.

VRRP-A Policy-Based Failover Template

VRRP-A provides flexible event tracking and policy-based failover support via a template. This feature allows policy-based failover even when VRRP-A preemption is disabled and the ACOS device typically would remain in an Active state, despite VRID priority changes. It allows for event-based failover once a tracked event occurs.

The following topics are covered:

Template Rules and Partitions	55
Preemption and Levels	56
Precedence Causing a Failover	56
Events Tracked for Weight via the Templates	60
Configuring VRRP-A Policy-Based Failovers	62

Template Rules and Partitions

The following topics are covered:

Policy-Based Failover Template Rules for the Shared Partition55
Policy-Based Failover Template Rules for L3V Partitions 55

Policy-Based Failover Template Rules for the Shared Partition

When creating templates in the shared partition, adhere to the following rules:

- You may create multiple templates, however, only one template can be associated with a particular VRID. For example, if you create two templates called “template1” and “template2,” you can only associate either template1 or template2 with VRID1. Assuming template1 and template2 have different events configured and you want template1 (that tracks for VLANs and gateways) to also address the events that template2 tracks (for interfaces and routes), edit template1 to include the events covered by template2. You cannot associate both templates to VRID1 to have it track all four events (VLANs, gateways, interfaces, and routes).
- Create templates with unique names, however, the second time you try to create a template with the same name, you will enter the template module and can edit the events listed in the template.
- You can create a maximum of 32 templates in the shared partition.
- You can associate the same template to multiple VRIDs. For example, VRID1 and VRID23 can both be attached to template1.

Policy-Based Failover Template Rules for L3V Partitions

When creating policy-based failover templates in an L3V partition, adhere to the following rules:

- You may create multiple templates, however, only one template can be associated with a particular VRID. For example, if you create two templates called “template1” and “template2,” you can only associate either template1 or template2 with VRID1. Assuming template1 and template2 have different events configured and you want

template1 (that tracks for VLANs and gateways) to also address the events that template2 tracks (for interfaces and routes), edit template1 to include the events covered by template2. You cannot associate both templates to VRID1 to have it track all four events (VLANs, gateways, interfaces, and routes).

- Create templates with unique names, however, the second time you try to create a template with the same name, you will enter the template module and can edit the events listed in the template.
- To associate a template with a private partition, you must switch to an active partition.
- Since each private partition can have its own template, and templates cannot be shared across multiple partitions, template names do not need to be unique. For example, template1 can be the template for private partition1 and template1 can be the template for private partition2. They can be named the same but can track different events for different VRIDs.
- You can create a maximum of 32 templates in the private partition.

Preemption and Levels

Preemption is always enabled on Level 1 (weight) and can be configured to be as enabled or disabled in Level 2 (priority). With preemption always being enabled on Level 1, the higher weight of an ACOS device will always place it in an Active state and the remaining ACOS devices in a Standby state. Level 1 weight assignments for events will result in the deduction of that weight from the VRID when an event occurs. When the weight is reduced for a particular VRID, it may result in a failover to another device if its total weight is lower than that of its peer devices. Since preemption works at the VRID level, any VRID with higher weight can take over for another VRID in a different ACOS device of a lower weight. When the original Active device is operational again, its weight is subject to change, and if it has the highest weight amongst its peers, it will resume its role and take over as the Active device.

Precedence Causing a Failover

The VRRP-A policy failover template provides a flexible way of defining events that will trigger failover of an ACOS device from Active to Standby state or vice versa. The

following briefly summarizes the order in which VRRP-A failover is determined:

- When the ACOS devices in a VRRP-A configuration have unequal weight, the one with the higher weight will be the Active device. Refer to [Higher Weight Scenario](#).
- If all ACOS devices have equal weight, the ACOS device with the higher priority will be the Active device. Refer to [Higher Priority Scenario](#).
- If the two ACOS devices have equal weight and equal priority, the device with the lower Device ID will be the Active device. Refer to [Equal Weight and Priority Scenario](#).

Higher Weight Scenario

Use the `show vrrp-a` command to view the Active/Standby status of the ACOS devices: Note that Unit 1 is the Active device and its weight is “65534” for VRID1:

```
ACOS(config)# show vrrp-a
vrid 0
Unit          State          Weight          Priority
2 (Local)     Standby        65534           150
              became Standby at: May 20 12:10:05 2013
              for 0 Day,22 Hour,54 min
1 (Peer)      Active         65534           150
vrid 1
Unit          State          Weight          Priority
2 (Local)     Standby        65534           150
              became Standby at: May 20 12:20:05 2013
              for 0 Day,22 Hour,54 min
1 (Peer)      Active         65534           200
```

After the event results in a reduction of the weight, Unit 1 now has a weight of “65334” for VRID1, and failover was triggered that placed Unit 1 in a Standby state:

```
ACOS(config)# show vrrp-a
vrid 0
Unit          State          Weight          Priority
2 (Local)     Standby        65534           150
              became Standby at: May 20 12:10:05 2013
              for 0 Day,22 Hour,54 min
1 (Peer)      Active         65534           150
```

```
vrid 1
Unit          State          Weight          Priority
2 (Local)    Active         65534           150
              became Active at:  May 20 12:10:05 2013
              for 0 Day,22 Hour,54 min
1 (Peer)     Standby        65334           200
```

Higher Priority Scenario

Use the `show vrrp-a` command to view the Active/Standby status of the ACOS devices: Note for VRID1 that Unit 1 is the Active device and its priority is “200:”

```
ACOS(config)# show vrrp-a
vrid 0
Unit          State          Weight          Priority
2 (Local)    Standby        65534           150
vrid 1          became Standby at:  May 20 12:10:05 2013
              for 0 Day,22 Hour,54 min
1 (Peer)     Active         65534           150

Unit          State          Weight          Priority
2 (Local)    Standby        65534           150
              became Standby at:  May 20 12:10:05 2013
              for 0 Day,22 Hour,54 min
1 (Peer)     Active         65534           200
```

After the event results in a reduction of the priority, Unit 1 now has a priority of “1” for VRID1, and failover was triggered that placed Unit 1 in a Standby state:

```
ACOS(config)# show vrrp-a
vrid 0
Unit          State          Weight          Priority
2 (Local)    Standby        65534           150
              became Standby at:  May 20 12:15:05 2013
              for 0 Day,22 Hour,54 min
1 (Peer)     Active         65534           150
vrid 1
Unit          State          Weight          Priority
2 (Local)    Active         65534           150
              became Active at:  May 20 12:15:05 2013
```

```

                                for 0 Day,22 Hour,54 min
1 (Peer)      Standby      65534      1      *

```

Equal Weight and Priority Scenario

Use the `show vrrp-a` command to view the Active/Standby status of the ACOS devices: Note that VRID1 has an unequal weight of “65534” and an equal priority of “150:”

```

ACOS(config)# show vrrp-a
vrid 0
Unit          State          Weight          Priority
2 (Local)     Standby        65534           150             *
              became      Standby at:    May 20 12:10:05 2013
              for 0 Day,22 Hour,54 min
1 (Peer)     Active         65534           150
vrid 1
Unit          State          Weight          Priority
2 (Local)     Active         65534           150
              became      Active at:     May 20 12:20:05 2013
              for 0 Day,22 Hour,54 min
1 (Peer)     Standby        65334           150             *

```

After the event results in an increase in weight, Unit 1 and Unit 2 both have equal weight and priority, yet a failover is triggered based on the device ID. For VRID1, Unit 1 is lower than Unit 2, it is now in an Active state instead of the Standby state:

```

ACOS(config)# show vrrp-a
vrid 0
Unit          State          Weight          Priority
2 (Local)     Standby        65534           150             *
              became      Standby at:    May 20 12:10:05 2013
              for 0 Day,22 Hour,54 min
1 (Peer)     Active         65534           150
vrid 1
Unit          State          Weight          Priority
2 (Local)     Standby        65534           150             *
              became      Standby at:    May 20 12:20:05 2013
              for 0 Day,22 Hour,54 min
1 (Peer)     Active         65534           150

```

Events Tracked for Weight via the Templates

Several events can be tracked for weight or priority. However, the failover policy templates only allow you to specify the weight for each event, not the priority. If configured using the GUI or the CLI, the following events may cause a change in the VRID and result in an ACOS device failover:

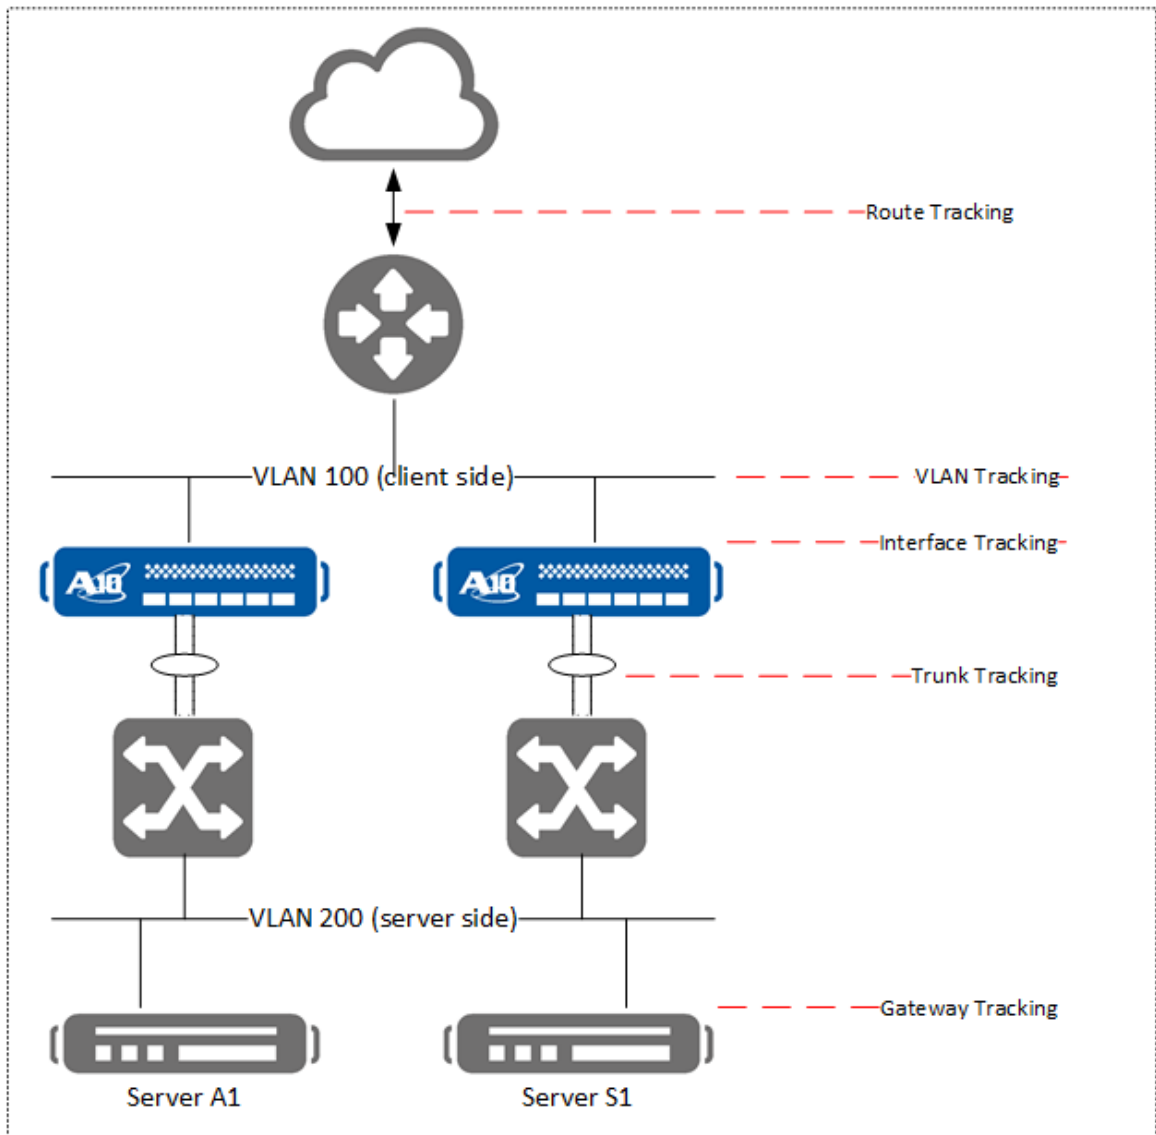
Table 2 : Configurable Event Tracking

Event	Description
Lost link a default gateway	VRRP-A periodically tests connectivity to the IPv4 and IPv6 default gateways connection to a real server by pinging them. If a gateway stops responding, VRRP-A reduces the weight or priority for the VRID. The weight value to subtract can be specified individually for each gateway's IP address that you configure in the tracking events list.
VLAN inactivity	If VRRP-A stops detecting traffic on a VLAN, VRRP-A reduces the device's priority for the VRID. The priority value to subtract can be specified individually for each VLAN. VLAN tracking is only supported in the shared partition.
Lost link on a trunk	If the trunk or individual ports in the trunk go down, VRRP-A reduces the device's priority for the VRID. The priority value to subtract can be specified for the trunk and individual ports within the trunk. If you track a trunk that does not exist, VRRP-A will not reduce the device priority. To track a trunk port within an L3V partition, you must verify that the tracked port is being used within that partition (for example, verify that the port is used in a VLAN of the partition): <ul style="list-style-type: none"> • If the tracked trunk is down, VRRP-A reduces the device priority based on the trunk value and ignores the status of the ports within the trunk. • If the tracked trunk is up, VRRP-A checks the ports

Table 2 : Configurable Event Tracking

Event	Description
	<p>within the trunk, and if any of them are down, VRRP-A reduces the device priority (1x configured port priority). If two ports are down, VRRP-A reduces the device priority by twice the priority assigned to the ports within the trunk (2x configured port priority).</p>
Lost data route	<p>If an IPv4 or IPv6 route matching the specified options is not in the data route table, VRRP-A reduces the device's priority for the VRID.</p>
Lost link on an Ethernet port	<p>If the link goes down on an Ethernet data port, VRRP-A reduces the device's priority for the VRID. The priority value to subtract can be specified individually for each Ethernet data port.</p> <p>If a tracked interface is a member of a trunk, only the lead port in the trunk is shown in the tracking configuration and statistics. For example, if a trunk contains ports 1-3 and you configure tracking of port 3, the configuration will show that tracking is enabled on port 1. Likewise, tracking statistics will show port 1, not port 3. Similarly, if port 1 goes down but port 3 is still up, statistics still will show that port 1 is up since it is the lead port for the trunk.</p>

The following graphic visually represents where tracking can be enabled for an ACOS device:



Configuring VRRP-A Policy-Based Failovers

The following topics are covered:

- [Using the GUI to Configure VRRP-A Policy-Based Failovers](#) 63
- [Using the CLI to Create VRRP-A Policy-Based Failovers](#) 64

NOTE: For details on additional template rules, refer to [Template Rules and Partitions](#).

Using the GUI to Configure VRRP-A Policy-Based Failovers

To configure the policy-based failover feature using the GUI, follow these steps:

1. Go to **System > VRRP-A > Failover Policy Template** to view the Failover Policy Template page.
2. Click **Create** to view the Create Failover Policy Template page.
3. Specify a template name in the Template Name field.
4. Configure the remainder of the page as needed.

For a description of the different events, see [Events Tracked for Weight via the Templates](#).

When an event occurs, the specified weight will be subtracted from the weight of the VRID, possibly causing a failover.

Create Failover Policy Template

Template Name *

Gateway Tracking

IP Address Type	Gateway	Weight	
IP v4	10.10.10.1	50	<input checked="" type="checkbox"/> ✕

Route Tracking

IP Address Type	IP Address	Mask	Gateway	Protocol	Distance	Weight	
IP v4	10.10.10.0	/24	10.10.10.1	static	4	55	<input checked="" type="checkbox"/> ✕

Trunk Tracking

Trunk ID	Weight	Per Port Weight	

Interface Tracking

Interface Ethernet	Weight	
1	40	<input checked="" type="checkbox"/> ✕

Vlan Tracking

VLAN ID	Weight	Timeout	
10	35	30	<input checked="" type="checkbox"/> ✕

For more information about this screen, refer to the online help.

5. Click **Create** when you are finished.

Your failover template will be listed on the Failover Policy Template page.

Using the CLI to Create VRRP-A Policy-Based Failovers

Failover Policy Templates can be configured using the CLI in either the shared or the private partition.

Configuring Templates in the Shared Partition

To configure the policy-based failover feature in the shared partition using the CLI, follow these steps:

1. Create your failover policy template. For example., create a template with the name `template1`:

```
ACOS(config)# vrrp-a fail-over-policy-template template1
ACOS(config-fail-over-policy-template:tem...)#
```

2. Configure the events you wish to track:

- a. For gateway tracking, indicate the gateway IP address and assign a weight for the gateway in the event of a failure. For example, if gateway 10.10.10.1 fails, 50 will be subtracted from the weight of the VRID:

```
ACOS(config-fail-over-policy-template:tem...)# gateway 10.10.10.1
weight 50
```

- b. For interface tracking, indicate the interface type, interface number, and assign a weight for that interface in the event of a failure. For example, if Ethernet interface 1 fails, 35 will be subtracted from the weight of the VRID:

```
ACOS(config-fail-over-policy-template:tem...)# interface ethernet 1
weight 35
```

- c. For route tracking, indicate the route number and assign a weight for that route in the event of failure. For example, if route 20.20.20.0 /24 fails, 70 will be subtracted from the weight of the VRID:

```
ACOS(config-fail-over-policy-template:tem...)# route 20.20.20.0 /24
weight 70
```

- d. For trunk tracking, indicate the trunk identification number and assign a weight for that trunk in the event of failure. For example, if trunk 4 fails, 25 will be subtracted from the weight of the VRID:

```
ACOS(config-fail-over-policy-template:tem...)# trunk 4 weight 25
```

- e. For VLAN tracking, indicate the VLAN identification number, the timeout value, and assign a weight for that trunk in the event of failure. For example, if VLAN 6 fails, 40 will be subtracted from the weight of the VRID:

```
ACOS(config-fail-over-policy-template:tem...)# vlan 6 timeout 30
weight 40
```

3. Assign a failover template to a VRID from the Global configuration level. For example, to assign template1 to VRID 0:

```
ACOS(config)# vrrp-a vrid 0
ACOS(config-vrid:0)# blade-parameters
ACOS(config-vrid:0-blade-parameters)# fail-over-policy-template
template1
```

NOTE: Only one template can be assigned to a VRID in any partition.

Configuring Templates in L3V Partition

To configure the policy-based failover feature in an L3V partition, follow these steps:

1. Switch to the active partition (for example, companyA):

```
ACOS(config)# active-partition companyA id 3
```

2. Follow the instructions in [Configuring Templates in the Shared Partition](#).

Verifying the Fail-Over Policy Template Configuration

To view the reason for failover, the `show vrrp-a` command on the Active ACOS device will display the template event that has caused the failover:

```
ACOS(config)# show vrrp-a
vrid 0
Unit                State                Weight                Priority
```

VRRP-A Policy-Based Failover Template

```

1 (Local)      Active      65534      150
               became    Active at:  May 20 12:04:05 2013
               for 0 Day,22 Hour,54 min
2 (Peer)      Standby    65534      150      *
vrid 1
Unit          State      Weight     Priority
1 (Local)    Standby    65334      150      *
               became    Standby at:  May 20 12:04:05 2013
               for 0 Day,22 Hour,54 min
2 (Peer)      Active    65534      150
Detected local policy based fail over events:
  interface ethernet 1 weight 200

```

If you specify the name of the template, the `show vrrp-a fail-over-policy-template` command will display the contents of that template:

```

ACOS(config)# show vrrp-a fail-over-policy-template templatel
vrrp-a fail-over-policy-template templatel
  interface ethernet 1 weight 35
  gateway 10.10.10.1 weight 50
  vlan 6 timeout 30 weight 40
  trunk 4 weight 25
  route 20.20.20.0 255.255.255.0 weight 70

```

Use the `show vrrp-a` command with the `detail` option to display comprehensive information on your current VRRP-A configuration:

```

ACOS(config)# show vrrp-a detail
vrid 0
Unit          State      Weight     Priority
2 (Local)    Standby    65534      150      *
               became    Standby at:  May 20 12:04:05 2013
               for 0 Day,22 Hour,54 min
1 (Peer)      Active    65534      150
...

VRRP-A stats
Peer: 1, vrid 0
Port 1: received 25685 missed 3
...

```

```
Total packets sent for vrid 0: 58524
Sent from port 1: 29262
Sent from port 7: 29262

...

Peer IP[1]: 7.7.7.1
```

To view information on all the fail-over policy templates, you can use the **show vrrp-a fail-over-policy-template** command;

```
ACOS(config)# show vrrp-a fail-over-policy-template
vrrp-a fail-over-policy-template interface
  interface ethernet 1 weight 200
vrrp-a fail-over-policy-template trunk
  trunk 1 weight 1 per-port-weight 20
vrrp-a fail-over-policy-template gateway
  gateway 41.41.41.20 weight 200
  gateway 20.20.20.20 weight 50
vrrp-a fail-over-policy-template route
  route 4.4.4.0 255.255.255.0 weight 200
vrrp-a fail-over-policy-template vlan
  vlan 10 timeout 2 weight 200
vrrp-a fail-over-policy-template template1
  interface ethernet 1 weight 35
  gateway 10.10.10.1 weight 50
  vlan 6 timeout 30 weight 40
  trunk 4 weight 25
  route 20.20.20.0 255.255.255.0 weight 70
```

VRRP-A Service Migration

Service Migration enables the migration of all services from a particular virtual server in a source VRRP-A pod to another virtual server located in a destination VRRP-A pod. This migration is independent of partitions and VRIDs. Service Migration occurs in a stateful process where all the sessions from the virtual server of the source VRRP-A pod are synced to a virtual server on the destination VRRP-A pod.

A VRRP-A pod is defined by the VRRP-A set id. All ACOS devices within a VRRP-A pod have the same set-id.

Service Migration is currently supported for DSR virtual servers only. This feature also requires VRRP-A pods configured with OSPF route advertisements.

The following topics are covered:

Guidelines for Service Migration	69
Deploying Service Migration	69

Guidelines for Service Migration

The following are the guidelines for Service Migration:

- Service Migration can be initiated for only one virtual server at one time.
- Service Migration requires the virtual server to have DSR configured.
- Initiate a second migration only after the first migration is complete.
- The virtual server to be migrated requires a dedicated route map so that the migration does not affect other virtual servers within the VRRP-A pod.
- The floating IP address for default gateways used by the downstream devices of the source VRRP-A pod must be reachable from the destination pod.
- VRRP-A must be enabled for both the source pod and destination pod.
- The session syncs are from an active VRRP-A virtual server to another active VRRP-A virtual server. The active device then syncs the session to their corresponding HA standby devices.

Deploying Service Migration

Service Migration requires the following steps:

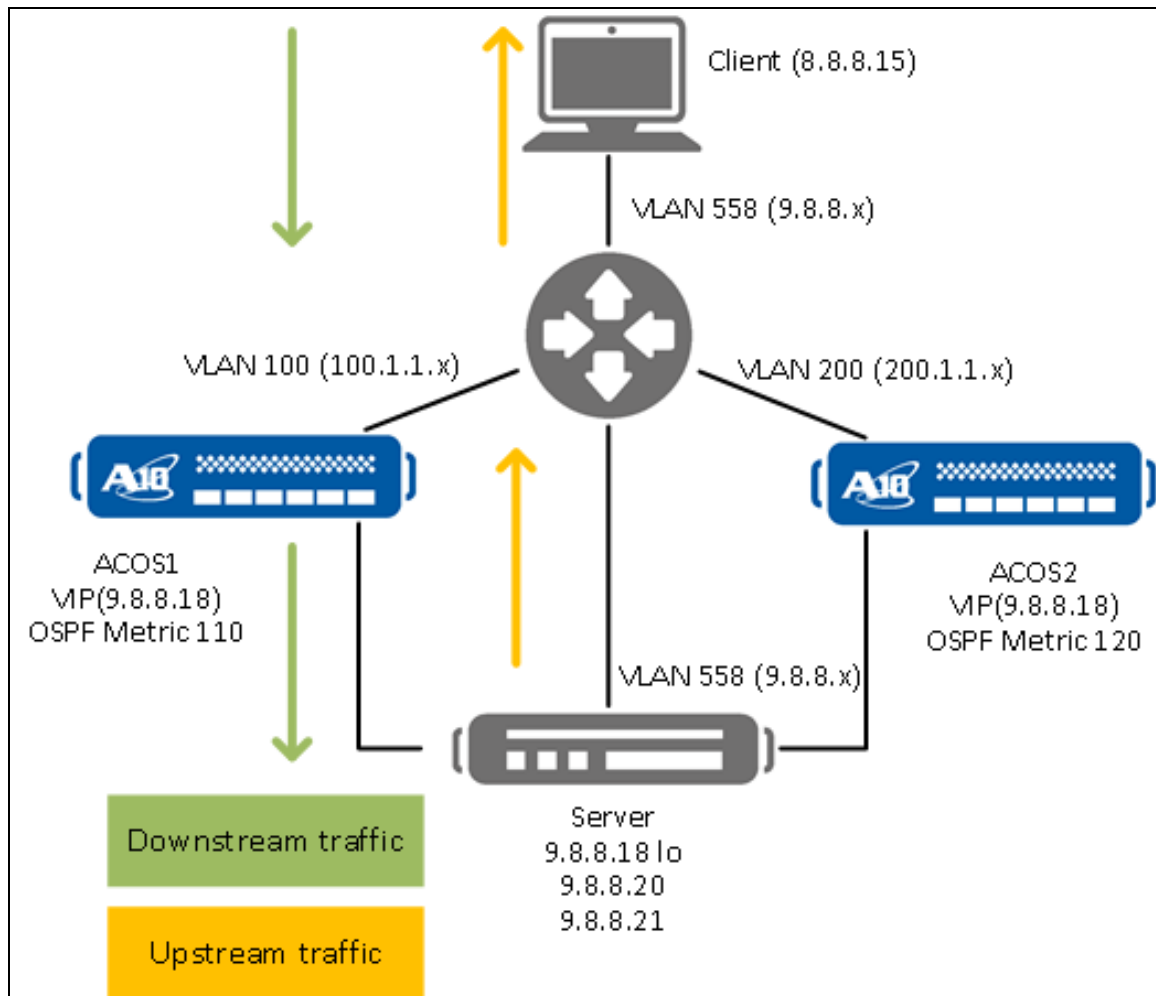
- [Step 1: Configuring the Destination VRRP-A Pod](#)
- [Step 2: Running the Migration Command](#)
- [Step 3: Lowering the Destination Pod OSPF Cost Metric](#)
- [Step 4: Running the Finish Migration Command](#)

Step 1: Configuring the Destination VRRP-A Pod

The following network topology displays two ACOS devices, each belonging to a different VRRP-A pod. In this example, the client is connected to both ACOS1 and ACOS2 through a gateway router. ACOS1 belongs to the source VRRP-A pod and ACOS2 belongs to the destination VRRP-A pod. ACOS1 and ACOS2 are then connected to the server.

Before starting Service Migration, traffic flows from the downstream router to ACOS1 to the upstream server. The response from the upstream server goes to the downstream router directly.

Figure 11 : Flow of Traffic Before Service Migration



The following is the sample configuration of ACOS1. ACOS1 connects to the router on tagged ethernet 3 and VLAN 100. ACOS1 connects to the server on tagged ethernet 6 and VLAN 558. The VIP of the virtual server is 9.8.8.18 and is associated with the OSPF route map ospf_red. The OSPF cost metric of the route is 110. The VIP has the service group sg-http associated.

```
!
vlan 100
```

```
tagged ethernet 3
router-interface ve 100
name toRTR
!
vlan 558
tagged ethernet 6
router-interface ve 558
name toServer
!
slb virtual-server vip1 9.8.8.18
redistribution-flagged
redistribute route-map ospf_red
port 80 tcp
ha-conn-mirror
service-group sg-http
no-dest-nat
!
router ospf
network 100.1.1.0 0.0.0.255 area 0
redistribute vip only-flagged
!
route-map ospf_red permit 1
set metric 110
!
```

To qualify ACOS2 for service migration as the destination pod, you must configure ACOS2 with the same VIP, service groups, and route map. Before initiating migration, ensure that the OSPF cost metric for the OSPF map in ACOS2 is higher than that of ACOS1. In this example, since the OSPF metric for the source pod is 110, configure the OSPF metric of the destination pod as 120.

The following is the sample configuration for ACOS2. ACOS2 connects to the router on tagged ethernet 3 and VLAN 200. ACOS2 connects to the server on tagged ethernet 2 and VLAN 558. The VIP of the virtual server is the same as that of the source POD and is configured as 9.8.8.18. Similar to the source pod, the VIP is associated with the OSPF route map `ospf_red`. The OSPF cost metric of the route is set higher than that of the source pod and configured as 120. Similar to the source pod, the VIP has the service group `sg-http` associated.

```
!
vlan 200
```

```
tagged ethernet 3
router-interface ve 200
name toRTR
!
vlan 558
tagged ethernet 2
router-interface ve 558
name toServer
!
slb virtual-server vip1 9.8.8.18
vrid 1
redistribute route-map ospf_red
port 80 tcp
ha-conn-mirror
service-group sg-http
no-dest-nat
!
router ospf
network 200.1.1.0 0.0.0.255 area 0
redistribute vip only-flagged
redistribute vip only-not-flagged
!
route-map ospf_red permit 1
set metric 120
!
```

Step 2: Running the Migration Command

1. From the source pod, run the migrate-vip command to start the migration.

```
ACOS1-Active(config-slbg vserver)#migrate-vip 5 target-floating-ipv4
9.8.8.50
```

When running the command:

- Specify the number of data CPUs on the destination pod. Here, the value is 5.
- Specify the floating IP or IPv6 address on the destination POD. The floating address must belong to the destination L3V partition id and vrid that the destination virtual server will be on. Here, the address is 9.8.8.50.

2. Run the following command to display the status of migration:

```
ACOS1-Active(config-slb vserver)# show slb virtual-server vip1
```

A result similar to the following is displayed:

```
Virtual server name:                vip1
Virtual server IP address:          9.8.8.18
Virtual server MAC:                 021f:a004:0007
Virtual server template:            default
TCP Stack TFO Cookie Time Limit:   60 seconds
Current connection:                 0
Current request:                    0
Total connection:                   0
Total request:                      0
Total request success:              0
Total forward bytes:                0
Total forward packets:              0
Total reverse bytes:                 0
Total reverse packets:              0
Peak connections:                   0
Current connection rate:            0 per second
Migration status:                   Sync started
```

Step 3: Lowering the Destination Pod OSPF Cost Metric

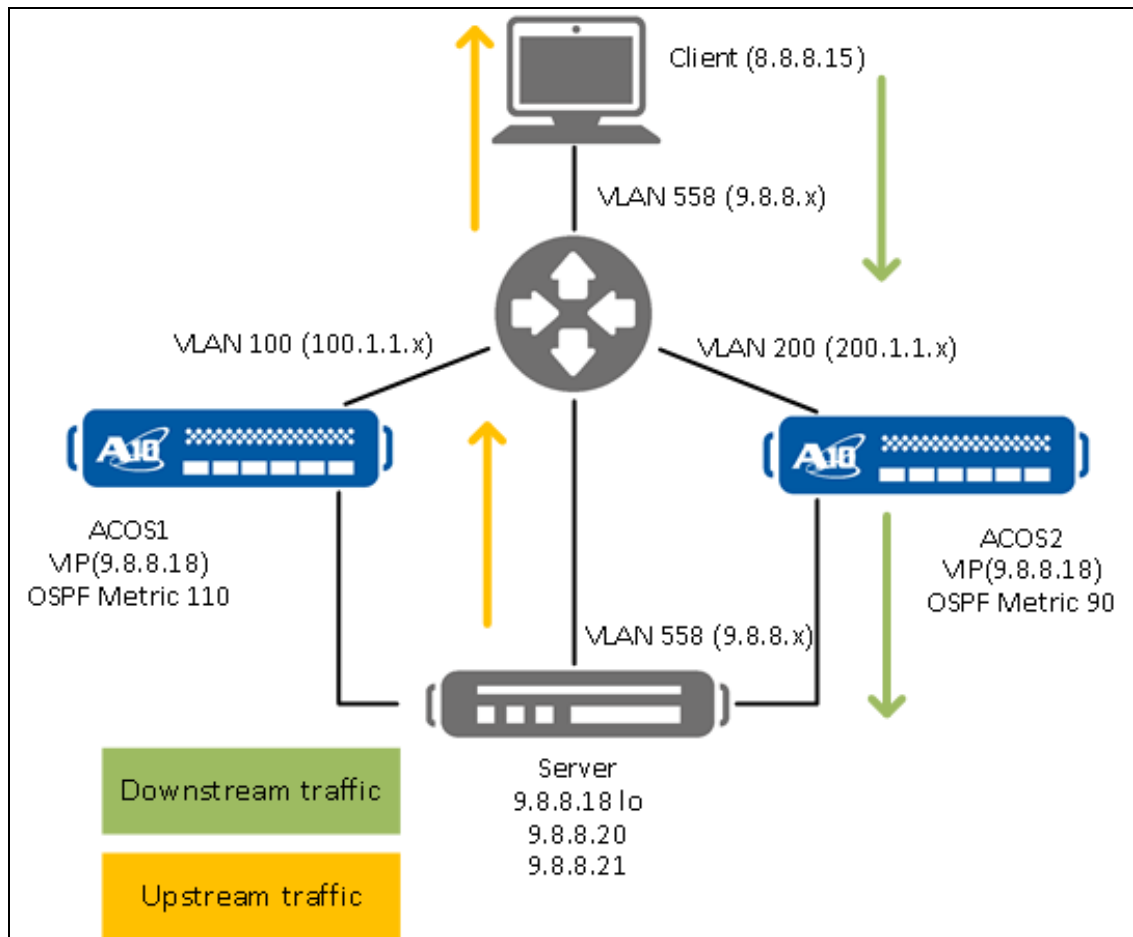
After the status of the migration changes to **Sync Complete**, change the OSPF cost metric on the destination pod to become lower than the source pod. This enables the traffic to flow from the gateway to the destination pod. The following command changes the OSPF cost metric on the destination pod to 90.

```
ACOS2-Active(config)#route-map ospf_red permit 1
ACOS2-Active(config-route-map:1)#set metric 90
```

Wait for a couple of minutes for the traffic to stabilize and all flows to move from the source pod to the destination pod.

After Service Migration completes, traffic flows from the downstream router to ACOS2 to the upstream server. The response from the upstream server goes to the downstream router directly.

Figure 12 : Flow of Traffic After Service Migration



Step 4: Running the Finish Migration Command

To finish the migration, run the `migrate-vip finish-migration` command on the source pod. You can now either let the sessions age or delete the VIP on the source pod.

At any time during the migration, you can cancel the migration by using the `migrate-vip cancel-migration` command on the source pod as follows:

```
ACOS1-Active(config-slb vserver)#migrate-vip cancel-migration
```

NOTE: If you cancel the migration, ensure that the OSPF cost metric on the source pod is made lower than that of the destination pod.

Glossary

B

bypass

A go-around functionality, where a specific action is performed via an external or alternative route instead of the intended route. In network security, a bypass is defined as a security system flaw that allows attackers access to network by circumventing the security mechanism.

D

DNS

Domain Name System. A hierarchical model and decentralized naming system that identifies computers, resources and network-based services over a private network or the Internet. It specifies information on web domain names associated with respective entities.

DSR

Direct Server Return. A load balancing mode where packets are routed to the backend server by modifying only the destination MAC address.

F

failover

A backup operational mode that allows the functions of a system component such as a network, server or database to be assumed by secondary components in instances when the primary components are unavailable due to failure or downtime.

G

gateway

A hardware device such as a firewall, router, or server, that acts as a gate between two networks and allows the inward and outward flow of traffic among the networks. It secures the nodes within a network and also serves as a node itself.

H

HTML

Hypertext Markup Language. The standard markup language developed for displaying documents in a web browser.

HTTP

HyperText Transfer Protocol. An underlying web protocol that defines the way messages can be formatted and sent, and the actions to be taken by web servers and browsers for responding to multiple commands.

I

IPv4

The fourth version of the Internet Protocol used as a core protocol in standardized internetworking methods over the Internet and packet-switched networks.

L

L3

A Network Layer, the third layer in the seven-layered OSI reference model used for routing traffic and forwarding packets across intermediate routers.

L4

A Transport Layer, the fourth layer of the seven-layered OSI reference model used for establishing host-to-host communications for applications.

load balancing

The process of distributing a set of tasks over a set of resources for making their overall processing more efficient and improving the performance.

P

packet flow

A sequence of packets sent from a source to a destination over packet-switching networks, across a host, a broadcast domain, or a multicast group.

S

service-group

A group of one or more services linked together for making object configurations simple.

SIP

Session Initiation Protocol. A signaling protocol that initiates, maintains, and ends real-time voice, video and messaging sessions of applications.

SMPP

Short Message Peer-to-Peer. A standardized protocol that provides scalable data communication interfaces for short message data transmissions between ESME, message centres, and routing entities.

source NAT pool

A pool of source NAT protocols used when an internal host begins a session with an external host and a dual NAT without using the switch IP.

SPDY

Speedy. A deprecated and open-spec networking protocol, developed by Google, which transports web content and manages HTTP traffic at a high speed by lowering webpage load latency and boosting web security.

subnet

An IP network subdivision.

T

TCP

Transmission Control Protocol. Key part of the main IP suite protocols used during initial network implementation.

TLS

Transport Layer Security. A cryptographic protocol that provides data transport and communications security over a network.

U

UDP

User Datagram Protocol. An alternative to TCP and used for setting up connections with low-latency and loss-tolerance between internet applications.

URL

Uniform Resource Locator. A web address that works as a reference to specify the location of a web resource on a computer network and also runs a mechanism for its retrieval.

V

VIP

Virtual Internet Protocol. An IP address which does not correspond to any real physical network interface but is used for mobility, network address translation, and fault-tolerance.

virtual port

An emulation or virtualization of a hardware port.

VLAN

Virtual Local Area Network. A LAN broadcast domain which is separated and isolated at the data link layer in a network.



©2025 A10 Networks, Inc. All rights reserved. A10 Networks, the A10 Networks logo, ACOS, A10 Thunder, Thunder TPS, A10 Harmony, SSLi and SSL Insight are trademarks or registered trademarks of A10 Networks, Inc. in the United States and other countries. All other trademarks are property of their respective owners. A10 Networks assumes no responsibility for any inaccuracies in this document. A10 Networks reserves the right to change, modify, transfer, or otherwise revise this publication without notice. For the full list of trademarks, visit: www.a10networks.com/company/legal/trademarks/.