# ACOS 6.0.7
# Scaleout Configuration Guide

**April, 2025**

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

# Introduction

This chapter introduces Scaleout and its benefits. It also describes how each ACOS device can be configured with a priority value that is used during the master election process.

The following topics are covered:

# About Scaleout

Scaleout is a solution where multiple ACOS devices form a cluster to provide the same set of services as a single infrastructure. The Scaleout technology enables CGN, Gi/SGi Firewall, and standalone Firewall services to be provided across multiple devices for load distribution and better scalability. The services provided by one or more NAT pools can be spanned across multiple devices that form a Scaleout cluster.

Scaleout supports dynamically adding or removing devices to the cluster. When the number of devices in the cluster changes, the traffic gets rebalanced automatically.

# Scaleout Benefits

The following are some of the use cases where Scaleout can be successfully implemented.

- **Scale Throughput Capacity As You Grow** - An organization may begin with 2 ACOS devices to meet its initial needs. However, over time, as the demand grows, the application capacity requirements grow with it.
  Scaleout provides support for multiple devices acting as a single cluster to increase the aggregate throughput capacity. Additional devices can be added or removed seamlessly as needed, without changing the network topology design when the current devices are already functioning as Scaleout cluster. Traffic can be distributed between the devices without any intervention, modification, or change to the network design.
- **Dynamically Provisioned Network** - Not all network resources have to be available one hundred percent of the time. Usually, network utilization has peak times throughout the day. For example, fixed-interval software updates (Microsoft's Patch Tuesday) result in predictable network peaks. The service providers often do not need all their resources available outside of working hours. In such cases, fewer ACOS devices can be enabled, and other parts of the network can be switched off during the off-peak hours. This decision can be made based on both the time of the day, the amount of user traffic, or resource utilization. This results in a much more efficient network, power usage saving, and reduced cooling requirements.

Scaleout allows new devices to be provisioned in or out of the cluster seamlessly to accommodate high or low utilization scenarios.

- **Redundancy** - The dimensions of the Scaleout cluster (that is the number of ACOS devices) can be adjusted to achieve n + m redundancy (such as n + 1) while utilizing all ACOS devices actively. This helps organizations to utilize all available resources inside the Scaleout cluster in normal operation and also have a certain level of redundancy at the same time during device failures.
- **Surge Relief** - In situations where an ACOS device is suddenly presented with a flow of traffic that causes it to work near full capacity, the load of other traffic can be temporarily distributed across other ACOS devices in the Scaleout cluster. When the load reduces, the service can easily be rolled back to the devices originally handling that service.
  Coupled with a controller, surge relief becomes available when additional processing nodes can be added, when required.

**NOTE:** This capability is achieved only by a third-party controller.

# System Requirements

Scaleout is currently supported on all shipping A10 Thunder systems. To configure Scaleout, the following are the requirements:

Table 1 : System Requirements

| Element | Description |
|---|---|
| ACOS Device | • Minimum - 1 ACOS device<br>• Maximum - 8 ACOS devices with the same hardware and software models |
| ACOS Version (Beginning from) | • ACOS 4.1.2-P1 for CGN<br>• ACOS 4.1.4-P3 for Gi/SGi Firewall and Standalone Firewall |

Throughout this document, devices in the cluster are referred to as “ACOS devices”. However, Scaleout is only supported on physical A10 Thunder Series, vThunder devices, Bare Metal, Thunder Container and services such as Carrier Grade Network

(CGN), Gi/SGi Firewall, and Standalone Firewall. Scaleout is not supported on AX Series devices.

Consider the following points:

- For the physical A10 Thunder Series, the devices in a cluster must be identical. For vThunder devices, the number of cores and memory must be the same across all instances.
- In a Scaleout cluster, every device must be running the same software version. For example, if a device within a cluster is running 6.0.3-P1, all the devices in that cluster must run 6.0.3-P1.
- When upgrading the software version of any device within the Scaleout cluster, ensure that all devices within the cluster are upgraded to the same software version.

For information on vThunder devices, see the vThunder installation guide for your hypervisor. All installation instructions are available for download on the Support Portal.

The installation Guides are located at: https://documentation.a10networks.com/Install/Software/A10_ACOS_Install/vThunder.html.

# Limitations

Scaleout has the following limitations:

- Scaleout supports only a single cluster.
- A Scaleout cluster supports a minimum of 1 ACOS device. The maximum number of ACOS devices supported is 16 for vThunder, Thunder Containers, and Bare Metal, and 8 for Hardware platforms.
- For the L2 redirection to work properly, the sender and the receiver Scaleout device must be in the same Layer 2 network with unique IP and MAC addresses.
- The time difference between multiple failures, within the limit of failures, must be at least 5 minutes. For example, in a Scaleout cluster with 5 nodes, up to 2 nodes can fail as long as the time difference between the failures is more than 5 minutes. This is required to enable a seamless flow of traffic.

- SNMP is only supported at the device level.
- When Scaleout is enabled on a device, VRRP-A cannot be enabled on the same device. However, the configuration items within VRRP-A such as device-ID and set-ID are used for ACOS Virtual Chassis System (aVCS).
- When Scaleout is configured in L3V mode and when follow-shared is configured to redirect the packets to the L3V partition using the Shared Partition redirection table, the “L3 Redirect Follow Shared Table Error” is displayed.
- For the Scaleout cluster to operate, you must configure at least 1 IPv4 address on the inside or outside and other Scaleout interfaces. In a pure IPv6 deployment, you must additionally configure at least 1 IPv4 address on the inside or outside and other Scaleout interfaces.

# Scaleout Overview

The following topics are covered:

# Scaleout Topology

Figure 1 illustrates a sample basic Scaleout topology.

Figure 1 : Sample Scaleout Topology

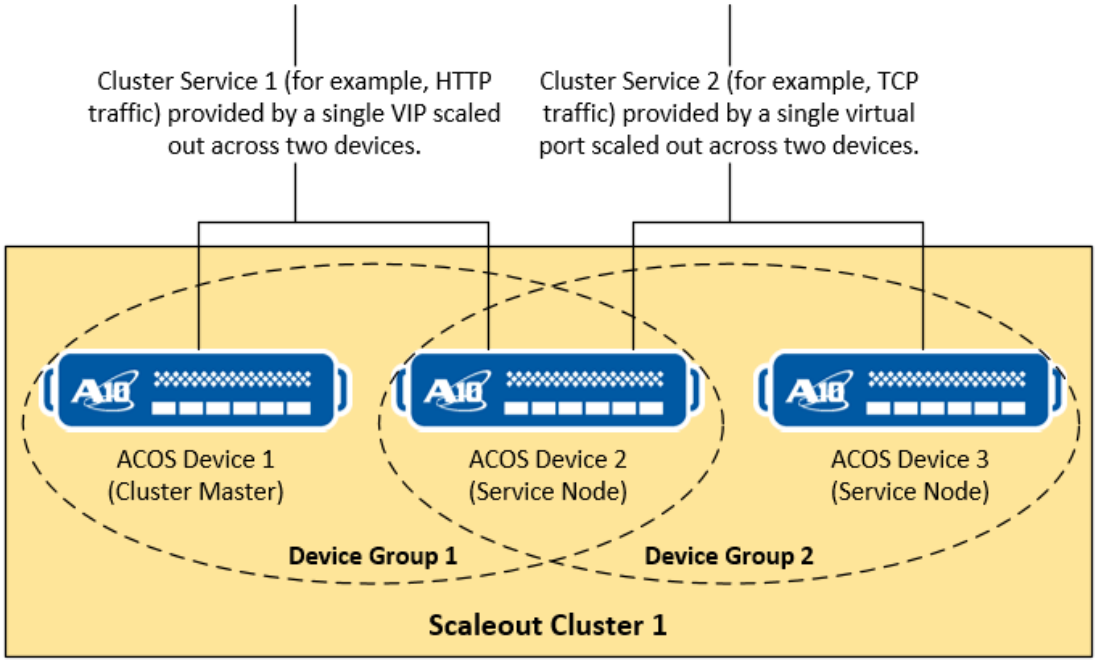


The elements in Figure 1 are described below in Table 2.

Table 2 : Explanation of the Sample Scaleout Topology

| Element | Description |
|---|---|
| Scaleout Cluster | The group of ACOS devices configured to provide Scaleout functionalities. |
| Cluster Service | The cluster service relates to a specific type of service provided by a VIP or virtual port being scaled out across multiple devices (cluster nodes).<br><br>Each cluster service can be provided by different sets of devices in the cluster. In the example above, HTTP traffic to a single VIP is scaled out across Device 1 and Device 2, and TCP traffic to a single virtual port is scaled out across Device 2 and Device 3. |

Table 2 : Explanation of the Sample Scaleout Topology

| Element | Description |
|---|---|
| | Use the **`template scaleout`** command to identify the services that you want to scaleout. For more information, see the *Command Line Interface Reference*. |
| Traffic Map | Each ACOS device in the cluster may have multiple traffic maps keeping track of the active and the standby ACOS device for each traffic user group. Sessions get synced between the Active and the Standby devices for each traffic user group. And traffic gets internally redirected using the traffic map for the traffic user group.<br>• Traffic Map<br>◦ This traffic map is updated and synced across the ACOS devices whenever there is any change to the cluster (device gets added/removed).<br>◦ The incoming traffic streams are classified and distributed to the devices within a cluster for processing. Incoming traffic is distributed based on a hash of the source IP. For more information, see Understand Traffic Distribution. The downstream router generally uses ECMP to distribute traffic to different devices. On ACOS, incoming traffic is hashed based on the source IP, then being L2-redirected to peer devices based on the traffic map.<br>◦ A traffic map shows how the incoming traffic is distributed across devices. Devices inside the cluster distribute packets based on the traffic map. It shows which device will process incoming client-side traffic based on source IP hashing (it is called as User-Group). If incoming traffic arrives at a device other than the active node in the traffic map, it is L2-redirected to the correct device based on the traffic map. You can use the **`show scaleout traffic-map`** command to view information about the traffic maps on each device. For more information, see `show scaleout` command in the *Command Line Reference Guide*.<br>Example: |

Table 2 : Explanation of the Sample Scaleout Topology

| Element | Description |
|---|---|
| | By default, **`show scaleout traffic-map`** is per partition level.<br>`DEV4#show scaleout traffic-map virtual-server?`<br>`NAME<length:1-63> Specify the virtual service name`<br>`DEV4#show scaleout traffic-map virtual-server vs`<br>`Number of Traffic Map Tables: 5`<br>`service type: virtual server - virtual server: vs`<br>`User-Group Active Device Standby Device New Act Device New Stby Device`<br>`0 4 5 - -`<br>`DEV4#show scaleout traffic-map virtual-server vs ?`<br>`src-ip Specify the source ip address`<br>`virtual-port Specify the virtual port number`<br>`\| Output modifiers` |
| Traffic Slice | A traffic slice is a combination of the implicit (default user group) or explicit service-config template and the application configuration (CGN/GiFW). |
| User Group | • The entire set of traffic destined for a cluster service is classified into multiple subsets or user groups.<br>• Each traffic user group is assigned to a single ACOS device.<br>• Each ACOS device can have multiple user groups assigned to it.<br>• Different user groups from different cluster services may be assigned to a single ACOS device. This is useful for properly distributing traffic destined for a particular cluster service across service nodes. |
| User-group-count | User group count is the number of traffic user groups.<br>It can be configured in the L3V and shared partitions and applied to the default IPv4 and IPv6 traffic slices and the associated traffic maps. |
| Cluster Master | A single device in the cluster that keeps track of all other devices in that cluster. |

Table 2 : Explanation of the Sample Scaleout Topology

| Element | Description |
|---|---|
| | The cluster master is responsible for programming and distributing the traffic-maps to the various devices in the cluster. It also keeps track of service nodes leaving and joining the cluster and takes appropriate actions. For more information, see Election of the Cluster Master. |
| Service Nodes | Each device in the Scaleout cluster is a service node. All service nodes process incoming traffic and each node act as a traffic classification and distribution engine. The upstream router or switch may forward traffic to one of the devices in the cluster based on routing or ARP responses. The devices inspect the packet and may redirect it to other nodes within the cluster if the session for handling that packet is on another node based on the traffic map.<br><br>The devices also serve (provide) the corresponding virtual service for the packet's destination. The mapping of traffic segments and their owners is created by the cluster master and pushed to all service nodes.<br><br>The Cluster Master and the Service Nodes are also more generically referred to as Cluster nodes. |

# Scaleout Concepts

The following are the Scaleout Concepts:

| Element | Description |
|---|---|
| Cluster Mode | The cluster-mode enabled on each Scaleout cluster node at the partition level for traffic redirection and session synchronization.<br><br>The cluster-mode can be one of the following:<br><br>• Layer-2 (default)<br>• Layer-3 |

| Element | Description |
|---|---|
| | The Layer-2 and Layer-3 are mutually exclusive. |
| Start Delay | Start Delay is the time Scaleout must wait to join the cluster and to be operational after the device boot up. |

# Understand Traffic Distribution

Figure 2 shows a topology illustrating traffic distribution in a Scaleout topology. This type of traffic distribution is called Symmetric Distribution of Traffic.

Figure 2 : Scaleout Traffic Distribution - Symmetric Distribution

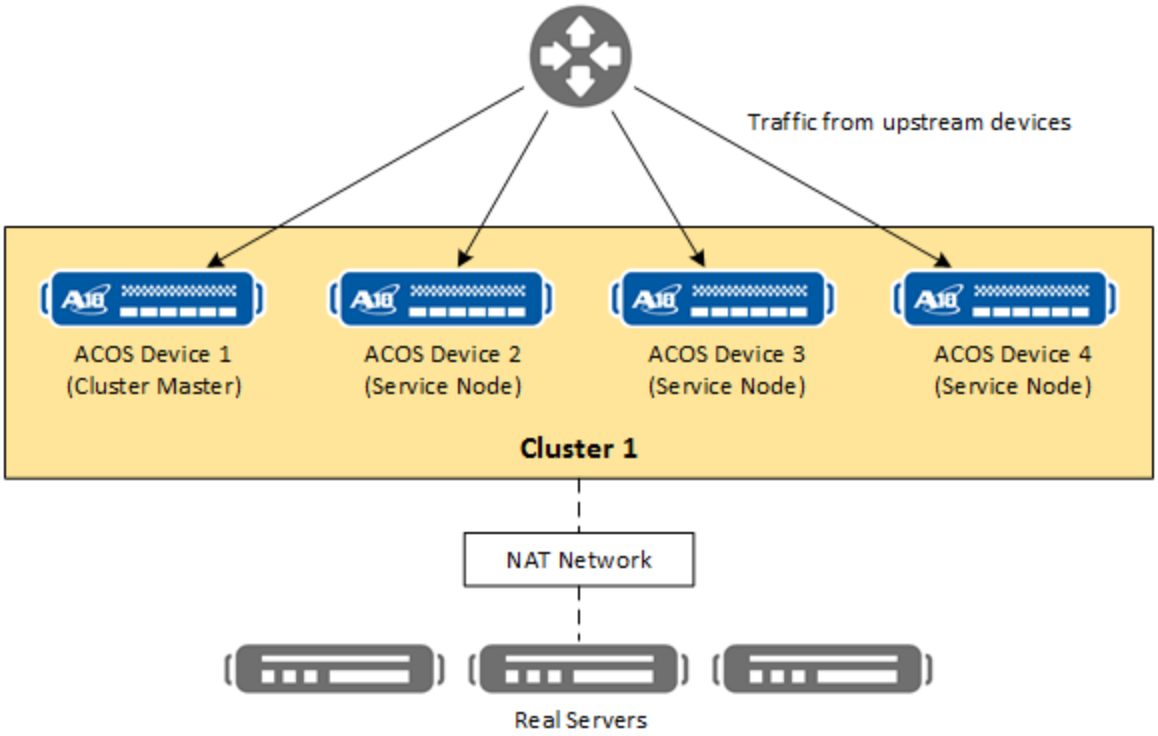


Traffic from upstream devices is distributed across the devices in the cluster based on a hash of the client IP address.

The entire set of traffic destined for a cluster service is classified into multiple subsets or user groups:

- Each traffic user group is assigned to a single ACOS device.
- Each ACOS device can have multiple user groups assigned to it.
- Different user groups from different cluster services may be assigned to a single ACOS device. This is useful for properly distributing traffic destined for a particular cluster service across service nodes.

The cluster devices may be connected to the rest of the network through an upstream router or a switch. Based on how that device forwards the traffic (for example, a router might have a default route for a VIP in a different subnet or ARP entry for a VIP in the same subnet), it will reach one of the cluster nodes. The cluster node may need to redirect the incoming packet based on a hash of the client IP if that node is not the owner of the set of traffic for the client IP.

All cluster nodes must be Layer-2 connected with high throughput links so that any redirection of packets can be done efficiently.

## Scaleout Redirection

The packets that arrive on an unrelated node are redirected to the destined active device based on the traffic map. The nodes in the Scaleout cluster can be L2 connected or L3 connected. By default, the Layer-2 cluster mode is enabled. If the nodes are L3 connected, then the Layer-3 cluster mode must be enabled manually. When the Layer-3 cluster mode is enabled, the Layer-2 cluster mode gets automatically disabled.

L2 redirection is supported in the L3V partition.

ACOS previously required IPv6 interfaces to have IPv4 addresses configured. This limitation is removed. An interface used in a Scaleout configuration can be configured with an IPv6 address only as required. Redirection and session synchronization over IPv6 interfaces are fully supported. It also adds another layer of resiliency to the network. Both IPv4 and IPv6 traffic may use IPv4 or IPv6 interfaces to redirect traffic and synchronize sessions within the cluster.

**NOTE:**

- For IPv4 traffic, the redirection takes place over the IPv4 interfaces. If the IPv4 interfaces are unavailable and the IPv6 interfaces are available, then the traffic will be redirected using the IPv6 interfaces.
- For IPv6 traffic, the redirection takes place over the IPv6 interfaces. If the IPv6 interfaces are unavailable and the IPv4 interfaces are available, then the traffic will be redirected using the IPv4 interfaces.

The following is the configuration example. The IPv4 interfaces are VLANs 200 and 300, while the IPv6 interfaces are VLANs 400 and 500.

```
ACOS(config)# interface ve 200
ACOS(config-if:ve200)# name IPv4_Trust
ACOS(config-if:ve200)# ip address 30.30.30.1 255.255.255.0
ACOS(config-if:ve200)# ip client
ACOS(config-if:ve200)# ip nat inside
ACOS(config-if:ve200)# exit
```

```
ACOS(config)# interface ve 300
ACOS(config-if:ve300)# name IPv4_UnTrust
ACOS(config-if:ve300)# ip address 30.30.31.1 255.255.255.1
ACOS(config-if:ve300)# ip server
ACOS(config-if:ve300)# ip nat outside
ACOS(config-if:ve300)# exit
```

```
ACOS(config)# interface ve 400
ACOS(config-if:ve400)# name IPv6_Trust
ACOS(config-if:ve400)# ip client
ACOS(config-if:ve400)# ipv6 address fd36:a94c:26e9:6::2/64
ACOS(config-if:ve400)# ipv6 nat inside
ACOS(config-if:ve400)# exit
```

```
ACOS(config)# interface ve 500
ACOS(config-if:ve500)# name IPv6_UnTrust
ACOS(config-if:ve500)# ip server
ACOS(config-if:ve500)# ipv6 address fd36:a94c:26e9:7::2/64
```

```
ACOS(config-if:ve500)# exit
```

```
ACOS(config)# rule-set firewall
ACOS(config-rule set:firewall)# rule nat64
ACOS(config-rule set:firewall-rule:nat64)# ip-version v6
ACOS(config-rule set:firewall-rule:nat64)# action permit cgnv6 lsn-lid 1
log
ACOS(config-rule set:firewall-rule:nat64)# source ipv6-address any
ACOS(config-rule set:firewall-rule:nat64)# source zone GI-Trust
ACOS(config-rule set:firewall-rule:nat64)# dest object NAT64-Prefix
ACOS(config-rule set:firewall-rule:nat64)# dest zone GI-Untrust
ACOS(config-rule set:firewall-rule:nat64)# service any
ACOS(config-rule set:firewall-rule:nat64)# exit
```

```
ACOS(config-rule set:firewall)# rule Native-IPv6
ACOS(config-rule set:firewall-rule:Native...)# ip-version v6
ACOS(config-rule set:firewall-rule:Native...)# action permit forward
ACOS(config-rule set:firewall-rule:Native...)# source ipv6-address any
ACOS(config-rule set:firewall-rule:Native...)# source zone GI-Trust
ACOS(config-rule set:firewall-rule:Native...)# dest ipv6-address any
ACOS(config-rule set:firewall-rule:Native...)# dest zone GI-Untrust
ACOS(config-rule set:firewall-rule:Native...)# service any
ACOS(config-rule set:firewall-rule:Native...)# exit
```

The following topics are covered:



# Scaleout L2 Redirection

When an ACOS device receives packets on an incoming interface and VLAN, it checks the traffic map to find the destination device. If the current device is not the destination device, the packets are redirected using the Layer-2 network to the destination device through the same incoming interface and VLAN. See Figure 3.

Figure 3 : Sample CGN Scaleout L2 Redirection

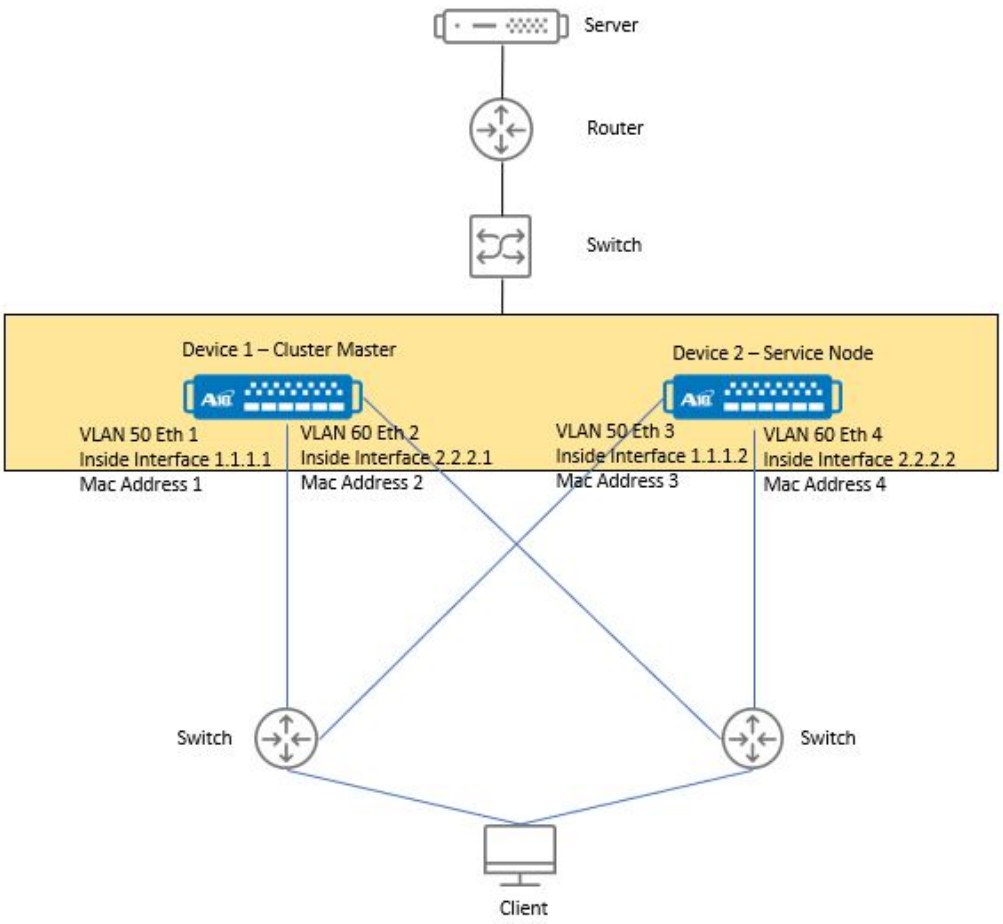


In Figure 3, the packets that arrive on Device 1, VLAN 50, Eth 1 interface is inspected to find out the device to which the packets must be sent. If the packets belong to Device 1, they are retained in the same device. If the packets are to be sent to Device 2, they are redirected over Device 2, VLAN 50, Eth 3 interface.

If the destination device's redirect interface link is down or unreachable due to various reasons, the scaleout architecture uses an alternate path between the devices in the cluster to protect against packet loss. In the event of the original incoming interface being down or unreachable, the traffic is forwarded to the destination device through the redundant link. See Figure 4.

Figure 4 : Sample CGN Scaleout L2 Redirect Link Failure

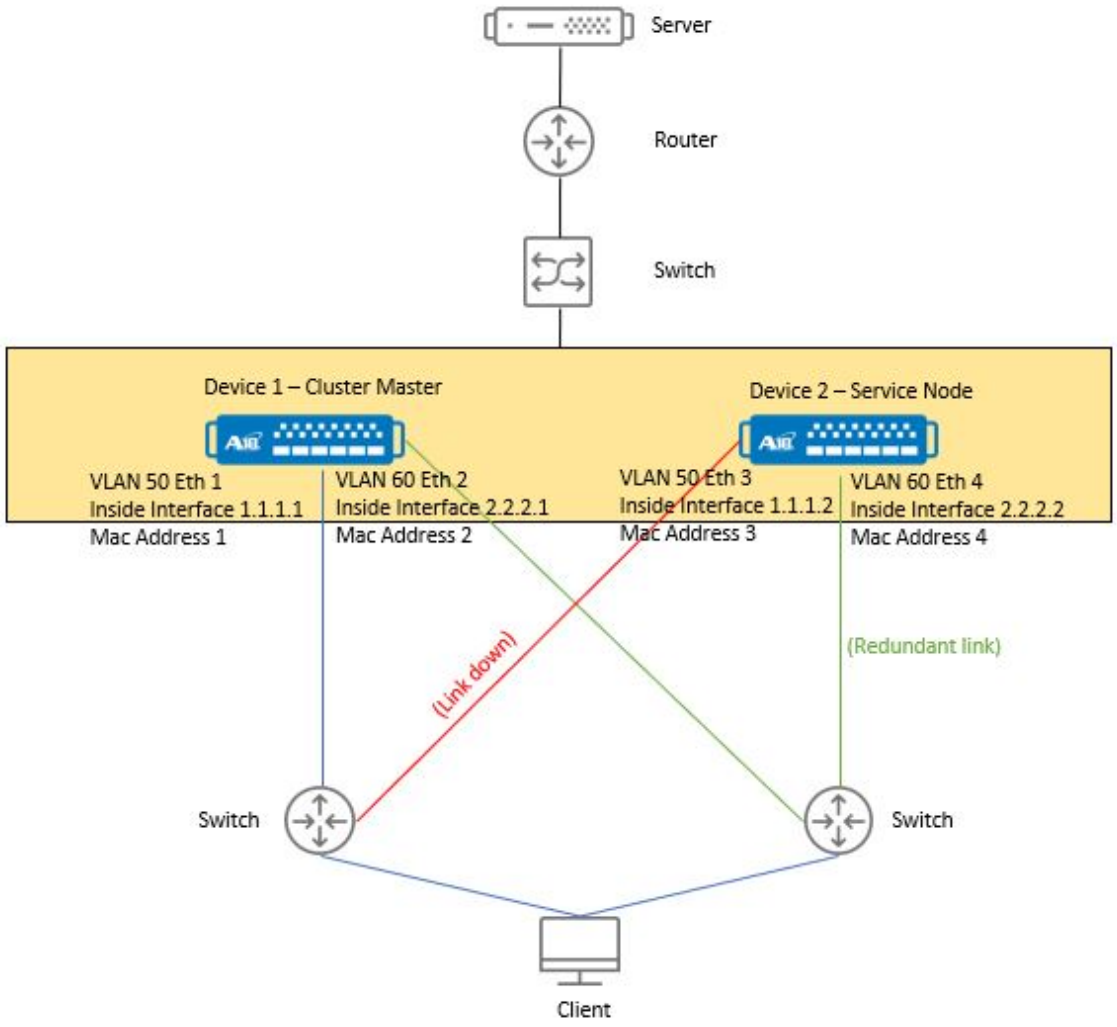


In the cluster shown in Figure 4, the packets arriving on Device 1, VLAN 50, Eth 1 belong to Device 2. If the Device 2, VLAN 50, Eth 3 interface is down, then Device 1 forwards the packets to Device 2 through VLAN 60, Eth 4 interface which is a redundant link.

When the original interface link fails, Scaleout chooses a redundant link based on the following criteria:

- The link connecting both devices must be up. For example, VLAN 60, Eth 4 is up.
- The interface type must be the same as the original link. For example, the Inside interface.

- The redundant links on Devices 1 and 2 must have the same IP subnet, and they must be L2 reachable to each other.

## L2 Redirection Using Dedicated Interface

The ACOS device uses the incoming traffic interface and VLAN for L2 redirection. This may cause network congestion on the existing network infrastructure as the amount of L2 redirection traffic increases. You can configure the device to use a dedicated interface to send L2 redirection traffic packets.

Figure 5 : Sample CGN Scaleout L2 Redirection over Dedicated Interface

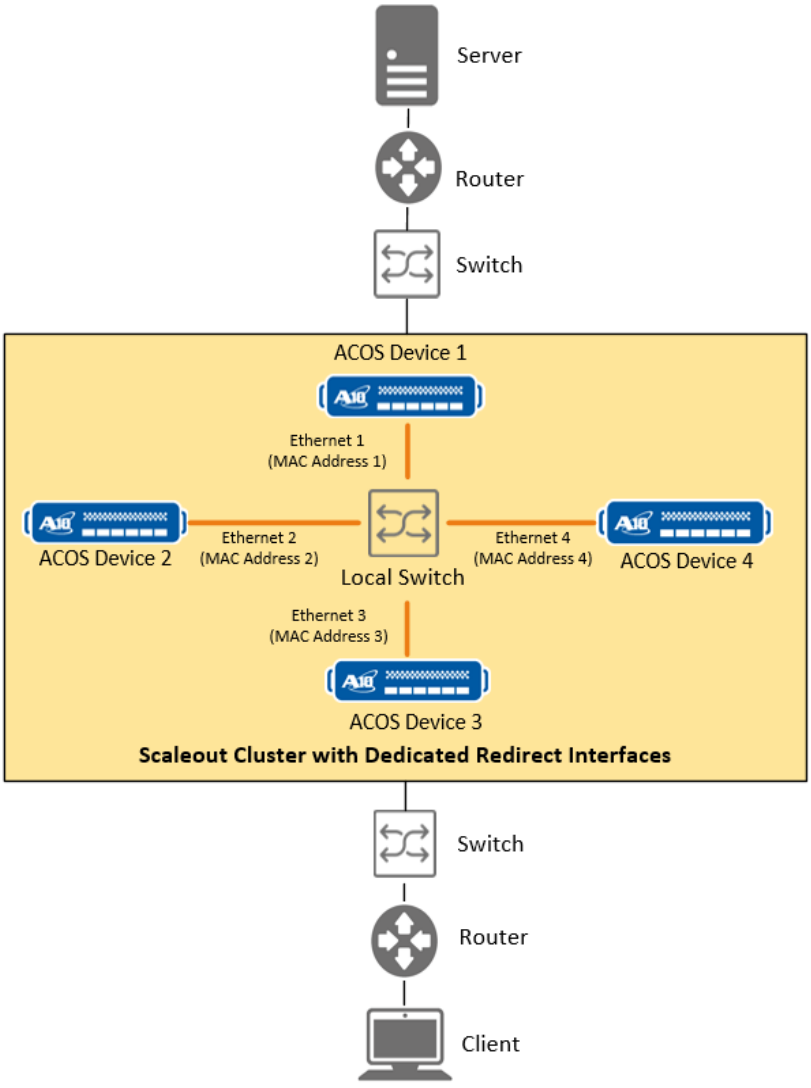

Ethernet 1 (using MAC Address 1), Ethernet 2 (using MAC Address 2), Ethernet 3 (using MAC Address 3), Ethernet 4 (using MAC Address 4) are configured as the redirect interface.

Upon receiving packets from the incoming interface and VLAN, ACOS can (Layer 2) redirect the traffic packets to the active Scaleout ACOS device using the dedicated redirect interface.

Under the Scaleout local device configuration level, use the **`l2-redirect`** command to configure a dedicated interface for L2 redirection.

All L2 redirection packets are sent out from the configured redirect interface.

When the dedicated interface for L2 redirection goes down, the traffic is dropped.

Consider the following points:

- All scaleout nodes are connected over the same L2 network through a switching network.
- Only 1 interface can be configured for L2 redirection.
- The interface configured for L2 redirection can be either an L2 trunk or a physical port. You can use the default VLAN or a specified VLAN while configuring the L2 trunk or the physical port for the dedicated redirect interface.
- The configuration of a dedicated redirect interface is only supported for CGN Scaleout.
- L2 dedicated redirection interface may only be configured in the shared partition. All shared and L3V redirected packets will use this configured redirection interface.

For information about L2 redirection configuration, see L2 Redirection Configuration.

## Limitations

- Respond to user mac is not supported with CGN Scaleout and L2 Redirection.
- If the dedicated L2 interface is configured, then the dedicated L2 ethernet interface must be in the L2 mode. The interface must not have any Virtual Ethernet or the IP addresses configured under it. The configured dedicated L2 interface can be used only for redirection.

For information about L2 configuration, see Scaleout L2 Redirection.

# Scaleout L3 Redirection

When the Scaleout devices are L3 connected, you must enable Layer-3 cluster-mode on all the devices in the cluster for forwarding the packets between the nodes.

Use the following command to enable Layer-3 cluster-mode:

```
ACOS(config-cluster:2-local-device)# cluster-mode layer-3
```

Figure 6 shows a sample view of a Scaleout cluster in which the ACOS devices are connected to each other over the Layer-3 network.

Figure 6 : Scaleout L3 Redirection

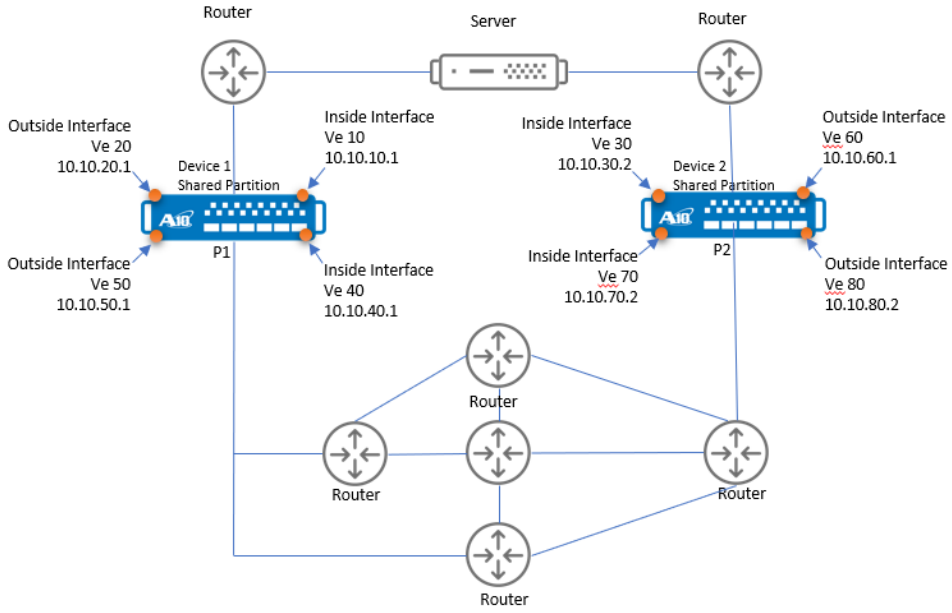


When an ACOS device receives packets on an incoming interface, it checks the traffic map to find the destination device. Once the destination device is identified, it finds the reachable redirection IPs and saves the next hop IP addresses for the destination IP. The packets are then forwarded over a VxLAN tunnel to the destination device's redirect interface.

**NOTE:** The redirected packets are dropped when all the peer redirect interfaces are down.

In the Layer-3 connected scaleout cluster, the IP of the redirect interface is the IP on which the device receives the redirected traffic from the other cluster nodes. You can configure more than one redirect interface for a particular cluster node. The maximum number of interfaces that can be added is 8 at any given time per partition. The redirect interface can be one of the following:

- ethernet
- trunk
- ve

Scaleout uses VXLAN (Virtual eXtensible Local Area Network) to encapsulate and send the tunneled redirected packets to the destination device. The redirected packets are decapsulated at the tunnel destination. In order to avoid fragmentation, it is recommended that the MTU of the redirected interface to be at least 50 bytes larger than the anticipated incoming data traffic to account for IPv4 VXLAN encapsulation and 70 bytes for IPv6 VXLAN encapsulation.

**NOTE:**

- The redirection of IPv4 traffic is through IPv4 VXLAN (Virtual eXtensible Local Area Network) tunnel.
- The redirection of IPv6 traffic is through IPv6 VXLAN tunnel.
- During an upgrade from the older releases to 5.2.1-P7 and the later releases, there may be a redirection failure. To resolve such failures, use the `use-v4-vxlan` CLI, which is configured under `encap vxlan` under the `traffic-redirection` CLI command to force all redirections to occur over IPv4 VxLAN tunnels.

The L3 connected Scaleout clusters can be configured in the L3V partition. You can configure the L3V partition to redirect and perform session sync based on the shared partition configuration using the **`follow-shared`** CLI command. This means that the packets are redirected to the peer cluster device using the shared partition redirection table.

You can also configure to skip the use of default routes to forward the redirected or session synchronized packets using the **`skip-default-route`** CLI command.

For information about L3 redirection configuration, see L3 Redirection Configuration.

### Limitation

The **`fw respond-to-user-mac`** command is not supported.

## Session Synchronization

Session synchronization interfaces can be configured to allow session synchronization requests to be handled specifically through such interfaces.

On the local device, configure a set of specific interfaces that handle session synchronization with other cluster nodes in the Scaleout cluster.

From the configured session synchronization interfaces, ACOS elects a "primary" interface, of which the following attributes are satisfied:

- Contains a valid L2/L3 configuration
- Is operational
- Shares at least 1 common IPv4 subnet with other cluster nodes

**NOTE:** If a session synchronization interface has not been configured, the first interface within the scaleout database with an IPv4 subnet in common with other cluster devices is used for session synchronization.

On the other cluster nodes, this set of configured interfaces do not control the interface/IP subnet to be used for session synchronization.

For more information, see Configure Scaleout Session Synchronization.

## aVCS Synchronization

ACOS supports both configuration synchronization and database synchronization. In configuration synchronization, any configuration made on aVCS vMaster will automatically sync with all the active devices in the cluster. Similarly, in database synchronization, any IP address configured adds to the database and syncs with all the devices in the cluster.

With the VCS database synchronization feature, you can enable and disable configuration synchronization and database synchronization separately. This can be done on any node within the cluster.

The vcs enable option enables configuration synchronization, while the vcs-database-distribution enable option enables database synchronization on the nodes. Similarly, the vcs disable option disables configuration synchronization, and the vcs database-distribution disable option disables database synchronization.

When the vcs disable option is enabled on a node, the configuration synchronization will not take place on that node, while database synchronization still takes place. Similarly, when the vcs database-distribution disable option is enabled on a node, the database synchronization will not take place on that node, while configuration synchronization still takes place.

Consider the following points:

- If you disable configuration and database synchronization on a node within the aVCS cluster where both services are enabled, the node will go down.
- In the cluster, on vBlade in which the vcs enable option is enabled and vcs database-distribution option is disabled, if the vcs database-distribution option is enabled again, you must run the `vcs reload device` command to restart the node and to join the cluster.
- If you enable the vcs enable or vcs database-distribution option on a disabled node, the changed configuration of aVCS will take effect only after running the vcs reload process.
- If you disable configuration and/or database synchronization on vMaster, a confirmation message will be prompted, asking whether you want to disable the service. If you choose Yes, aVCS on the node will shut down completely. Additionally, aVCS vMaster must have more services than vBlade.
- Before disabling either configuration synchronization or database synchronization on vMaster, you must switch over vMaster to another node that has more services than the existing vMaster.
- The vBlade cannot takeover vMaster role if it has less services than the current vMaster.

- The new node joins the cluster in the following scenarios:
  - If vMaster has both configuration and database synchronization enabled and running, the new vBlade will join the cluster successfully, even if vBlade has only configuration synchronization or database synchronization enabled.
  - If vMaster has configuration synchronization or database synchronization enabled and running, the new vBlade will not join the cluster if vBlade does not have the same service enabled or if vBlade has both services enabled. This is because vMaster cannot provide the missing service to vBlade.

You can use the `show vcs summary` to verify the status.

For more information, see Configure aVCS Synchronization.

# aVCS Reload

In the aVCS cluster, after completing configuration changes such as configuration synchronization, database synchronization, and device-specific parameter changes on the devices, you must reload the aVCS process for the changes to take effect.

In the vcs reload process, you can use the following options:

- **cluster-discovery** - This option applies the configuration changes on the aVCS chassis, while the device <device-id> option helps to reload a specific device when aVCS is enabled.
- **disable-merge** - When you are replacing a virtual chassis vBlade by removing an ACOS device and replacing it with another ACOS device of the same model, the disable-merge option allows the replacement device to be configured by vMaster. After the initial configuration migration, configuration synchronization operates normally. Without using this option, when you add an ACOS device to a virtual chassis that is already running, the device's configuration information is migrated to vMaster.
- **db-safe** - To keep database safe and avoid interruptions in the Scaleout cluster, the vcs reload process must be run in the db-safe mode. This mode ensures that Scaleout is stable and the traffic is uninterrupted and ignores the node leaving or joining. You can start the db-safe mode and specify the timeout value to automatically end the db-safe mode. The db-safe reload start completes within the

specified vcs dead-interval time. Also, you can complete the db-safe vcs reload within the specified timeout or after the timeout using the force option.

The force option can be used in the following scenarios:

- If the vcs reload db-safe is not completed automatically.
- If the vcs reload db-safe specified time is not expired.
- If the vcs reload is not completed for some reason and it remains in the db-safe mode even after the specified timeout period.

## Service-Config Template

To distribute traffic efficiently, you can configure the service-config template, specify the user group, and bind it to the IPv4 application (CGN/GiFW) configuration for the distribution of traffic.

The combination of the implicit (default user group) or explicit service-config template and the application configuration (CGN/GiFW) is a traffic slice.

Each service-config template is associated with a traffic map that distributes traffic and application objects, such as NAT IPs across the Scaleout cluster. The default IPv4 and IPv6 traffic may also be considered as an implicit traffic slice and have its associated default traffic maps.

## User Group

User group count can be configured in the L3V and shared partitions and applied to the default IPv4 and IPv6 traffic slices and the associated traffic maps. The user group can only have values of 64, 128, and 256. 256 is the default value. It allows the creation of a partition for smaller NAT pools without configuring a service-config template (Traffic Slice).

The user group count per partition or within the service-config template level supports smaller NAT pool prefixes such as /28, /30, and /32.

When the user group count is configured at the shared partition level, it serves as the default value for that partition. Likewise, at the L3V level, it only applies to the L3V partition.

## Limitations

- On a multi-PU platform, the minimum user-group-count that can be configured is 2.
- When the LSN NAT Pools for Fixed NAT are configured in the shared partition or any L3V has an explicit ‘default-user-group-count’ configured, the user-group-count in the shared partition cannot be modified.
- When the LSN NAT Pools for Fixed NAT are configured in the L3V partition, the user-group-count in the L3V partition cannot be modified.

# Scaleout Configuration

The following topics are covered:

# Configure Scaleout

This section shows how to configure Scaleout in a cluster of 3 devices using ACOS Virtual Chassis System (aVCS).

**NOTE:** ACOS Virtual Chassis System (aVCS) enables you to manage a cluster of ACOS devices like a single, virtual chassis.

Figure 7 illustrates a sample basic Scaleout topology.

Figure 7 : Sample Scaleout Topology

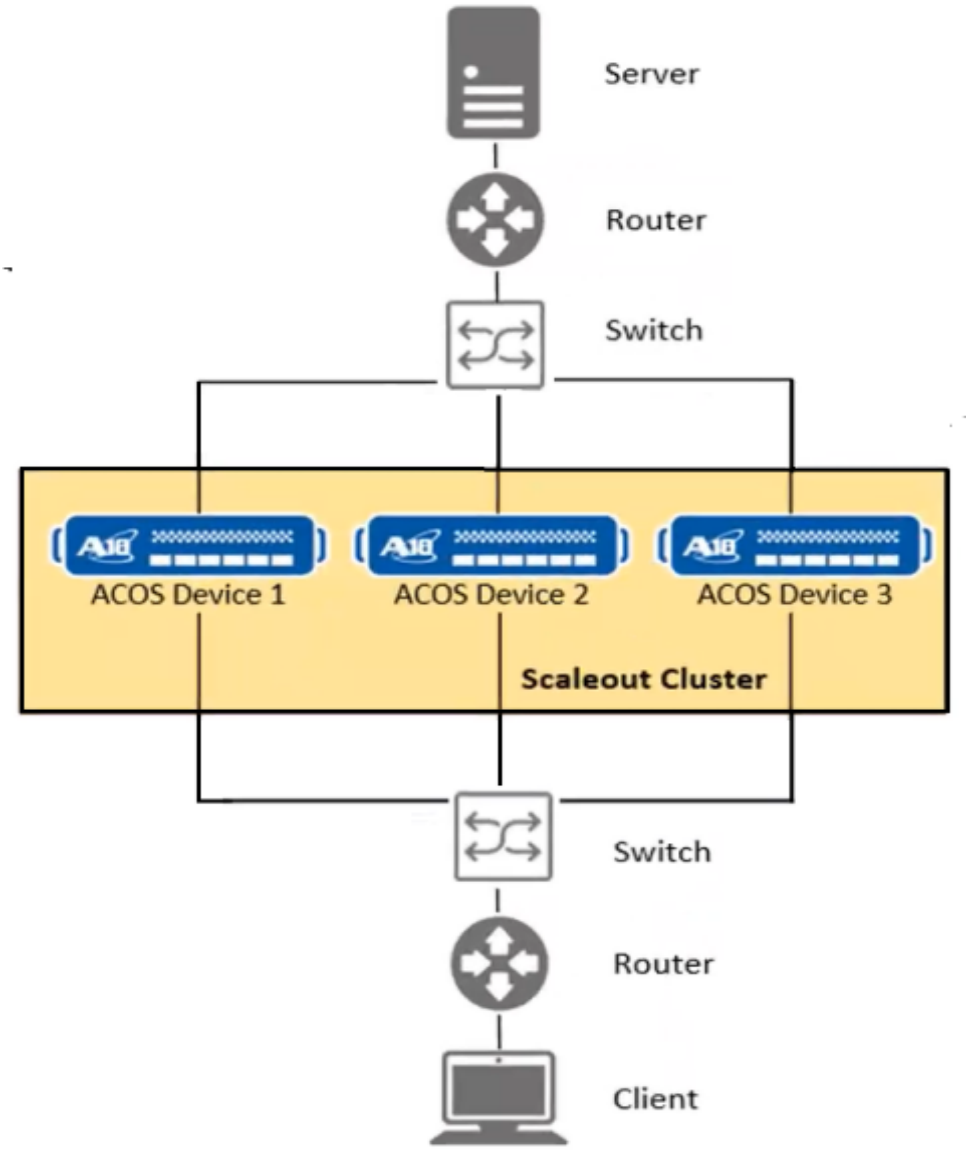


You can configure Scaleout in two modes:

- aVCS Unicast Mode - traffic flows from one node to another node.
- aVCS Multicast Mode - traffic flows from one node to multiple nodes.

The steps to configure Scaleout using aVCS Unicast or aVCS Multicast Mode are as follows:

1. Configure aVCS on each device.

   In a general deployment scenario, ACOS uses aVCS for Scaleout-related configuration synchronization. For Scaleout to function efficiently, it is essential to enable aVCS so that the Scaleout-related configurations are applied and synchronized across all devices within the cluster. When the aVCS is enabled, the aVCS configuration synchronization is also enabled. If you choose not to use aVCS, you must disable the aVCS configuration synchronization on all devices. The `vcs database-distribution enable` command enables the aVCS database synchronization for the Scaleout cluster. You must configure this command on each node in the cluster.

   a. For aVCS Unicast Mode, see Configure aVCS on Each Device for Unicast Mode.

   b. For aVCS Multicast Mode, see Configure aVCS on Each Device for Multicast Mode.

2. Set up the Scaleout configuration on vMaster. (See Configure Scaleout Cluster.)

   With aVCS configured and enabled, configuration changes on vMaster are automatically synchronized to the Service Nodes.

3. Enable Scaleout. (See Enable Scaleout.)

4. Set up the following configurations:

   a. For Carrier Grade Networking (CGN) - See Scaleout for Carrier Grade Networking (CGN).

   b. For Gi/SGi Firewall and Standalone Firewall - See Scaleout for Gi/SGi Firewall and Standalone Firewall.

## Configure aVCS on Each Device for Unicast Mode

The following topics are covered:

## Configure aVCS on Device 1 for Unicast Mode

To configure aVCS on device 1, use the following commands:

1. Specify a VRRP-A device ID and set ID:

```
ACOS(config)# hostname ACOS1
ACOS1(config)# vrrp-a common
ACOS1(config-common)# device-id 1
ACOS1(config-common)# set-id 1
ACOS1(config-common)# exit
```

2. Enable aVCS:

```
ACOS1(config)# vcs enable
```

**NOTE:** If you want to configure Scaleout without using aVCS, do not enable aVCS.

3. Configure a floating IP address for the virtual chassis:

```
ACOS1(config)# vcs floating-ip 192.168.220.24 / 25
The changed configuration of aVCS will take effect only after 'vcs
reload'
```

4. Configure an aVCS profile for the device. The `vcs discovery-mode unicast` command will remove aVCS interface commands.

```
ACOS1(config:1)# vcs device 1
ACOS1(config:1-device:1)# priority 200
ACOS1(config:1-device:1)# enable
ACOS1(config:1-device:1)# exit
ACOS1(config:1)# vcs unicast-election
ACOS1(config:1-unicast-election)# members
ACOS1(config:1-unicast-election-members)# ip-address 2.2.2.115
ACOS1(config:1-unicast-election-members)# ip-address 2.2.2.116
ACOS1(config:1-unicast-election-members)# ip-address 2.2.2.117
ACOS1(config:1-unicast-election-members)# exit
ACOS1(config:1)# vcs discovery-mode unicast
ACOS1(config:1)# vcs reload
```

## Configure aVCS on Device 2 for Unicast Mode

To configure aVCS on device 2, use the following commands:

1. Specify a VRRP-A device ID and set ID:

```
ACOS(config)# hostname ACOS2
ACOS2(config)# vrrp-a common
ACOS2(config-common)# device-id 2
ACOS2(config-common)# set-id 2
ACOS2(config-common)# exit
```

2. Enable aVCS:

```
ACOS2(config)# vcs enable
```

**NOTE:** If you want to configure Scaleout without using aVCS, do not enable aVCS.

3. Configure a floating IP address for the virtual chassis:

```
ACOS2(config:2)# vcs floating-ip 192.168.220.24 / 25
The changed configuration of aVCS will take effect only after 'vcs
reload'
```

4. Configure an aVCS profile for the device:

```
ACOS2(config:2)# vcs device 2
ACOS2(config:2-device:2)# priority 190
ACOS2(config:2-device:2)# enable
ACOS2(config:2-device:2)# exit
ACOS2(config:2)# vcs unicast-election
ACOS2(config:2-unicast-election)# members
ACOS2(config:2-unicast-election-members)# ip-address 2.2.2.115
ACOS2(config:2-unicast-election-members)# ip-address 2.2.2.116
ACOS2(config:2-unicast-election-members)# ip-address 2.2.2.117
ACOS2(config:2-unicast-election-members)# exit
ACOS2(config:2)# vcs discovery-mode unicast
ACOS2(config:2)# vcs reload
```

## Configure aVCS on Device 3 for Unicast Mode

To configure aVCS on device 3, use the following commands:

1. Specify a VRRP-A device ID and set ID:

```
ACOS(config)# hostname ACOS3
ACOS3(config)# vrrp-a common
ACOS3(config-common)# device-id 3
ACOS3(config-common) set-id 1
ACOS3(config-common)# exit
```

2. Enable aVCS:

```
ACOS3(config)# vcs enable
```

**NOTE:** If you want to configure Scaleout without using aVCS, do not enable aVCS.

3. Configure a floating IP address for the virtual chassis:

```
ACOS3(config:3)# vcs floating-ip 192.168.220.24 / 25
The changed configuration of aVCS will take effect only after 'vcs
reload'
```

4. Configure an aVCS profile for the device:

```
ACOS3(config:3)# vcs device 3
ACOS3(config:3-device:3)# priority 202
ACOS3(config:3-device:3)# enable
ACOS3(config:3-device:3)# exit
ACOS3(config:3)# vcs unicast-election
ACOS3(config:3-unicast-election)# members
ACOS3(config:3-unicast-election-members)# ip-address 2.2.2.115
ACOS3(config:3-unicast-election-members)# ip-address 2.2.2.116
ACOS3(config:3-unicast-election-members)# ip-address 2.2.2.117
ACOS3(config:3-unicast-election-members)# exit
ACOS3(config:3)# vcs discovery-mode unicast
ACOS3(config:3)# vcs reload
```

5. Use the following show command to view the aVCS information:

```
ACOS-vMaster[1/1]# show vcs summary
aVCS Chassis:
VCS Configuration-Sync Enabled:            Yes
VCS Database-distribution Enabled:         No
```

```
Chassis ID:                                 1
Floating IP:                                192.168.220.24 / 25
Unicast Election port:                      41473
Multicast IP:                               224.0.0.210
Multicast Port:                             41473
Version:                                    6.0.0-d.b284
Current Discover mode:                      Unicast
Members(* means local device):
ID  State       Priority IP:Port                    Location
------------------------------------------------------------------------
--------
1   vMaster(*)  200      2.2.2.115:41216            Local
2   vBlade      190      2.2.2.116:41216            Remote
3   vBlade      180      2.2.2.117:41216            Remote

Total: 3
```

For the detailed output field descriptions for the `show vcs summary` command, see the *Command Line Interface Reference* guide.

# Configure aVCS on Each Device for Multicast Mode

The following topics are covered:



## Configure aVCS on Device 1 for Multicast Mode

To configure aVCS on device 1, use the following commands:

1. Specify a VRRP-A device ID and set ID:

```
ACOS(config)# hostname ACOS1
ACOS1(config)# vrrp-a common
ACOS1(config-common)# device-id 1
ACOS1(config-common)# set-id 1
ACOS1(config-common)# exit
```

2. Enable aVCS:

```
ACOS1(config)# vcs enable
```

**NOTE:** If you want to configure Scaleout without using aVCS, do not enable aVCS.

3. Configure a floating IP address for the virtual chassis:

```
ACOS1(config)# vcs floating-ip 192.168.209.23 / 24
The changed configuration of aVCS will take effect only after 'vcs
reload'
```

4. Configure an aVCS profile for the device:

```
ACOS1(config:1)# vcs device 1
ACOS1(config:1-device:1)# interfaces management
ACOS1(config:1-device:1)# priority 200
ACOS1(config:1-device:1)# enable
ACOS1(config:1-device:1)# exit
ACOS1(config:1)# vcs reload

System configuration has been modified. Save? [yes/no]:yes
Building configuration...
Write configuration to default primary startup-config
[OK]
Running configuration is saved
```

## Configure aVCS on Device 2 for Multicast Mode

To configure aVCS on device 2, use the following commands:

1. Specify a VRRP-A device ID and set ID:

```
ACOS(config)# hostname ACOS2
ACOS2(config)# vrrp-a common
ACOS2(config-common)# device-id 2
ACOS2(config-common)# set-id 1
ACOS2(config-common)# exit
```

2. Enable aVCS:

```
ACOS2(config)# vcs enable
```

**NOTE:** If you want to configure Scaleout without using aVCS, do not enable aVCS.

3. Configure a floating IP address for the virtual chassis:

```
ACOS2(config:2)# vcs floating-ip 192.168.210.24 / 25
The changed configuration of aVCS will take effect only after 'vcs
reload'
```

4. Configure an aVCS profile for the device:

```
ACOS2(config:2)# vcs device 2
ACOS2(config:2-device:2)# interfaces management
ACOS2(config:2-device:2)# priority 190
ACOS2(config:2-device:2)# enable
ACOS2(config:2-device:2)# exit
ACOS2(config:2)# vcs reload

System configuration has been modified. Save? [yes/no]:yes
Building configuration...
Write configuration to default primary startup-config
[OK]
Running configuration is saved
```

## Configure aVCS on Device 3 for Multicast Mode

To configure aVCS on device 3, use the following commands:

1. Specify a VRRP-A device ID and set ID:

```
ACOS(config)# hostname ACOS3
ACOS3(config)# vrrp-a common
ACOS3(config-common)# device-id 3
ACOS3(config-common)# set-id 1
ACOS3(config-common)# exit
```

2. Enable aVCS:

```
ACOS3(config)# vcs enable
```

**NOTE:** If you want to configure Scaleout without using aVCS, do not enable aVCS.

3. Configure a floating IP address for the virtual chassis:

```
ACOS3(config:3)# vcs floating-ip 192.168.211.24 / 25
The changed configuration of aVCS will take effect only after 'vcs
reload'
```

4. Configure an aVCS profile for the device:

```
ACOS3(config:3)# vcs device 3
ACOS3(config:3-device:3)# priority 180
ACOS3(config:3-device:3)# interfaces management
ACOS3(config:3-device:3)# enable
ACOS3(config:3-device:3)# exit
ACOS3(config:3)# vcs reload

System configuration has been modified. Save? [yes/no]:yes
Building configuration...
Write configuration to default primary startup-config
[OK]
Running configuration is saved
```

5. Use the following show command to view the aVCS information:

```
ACOS-vMaster[1/1]# show vcs summary

aVCS Chassis:
VCS Configuration-Sync Enabled:            Yes
VCS Database-distribution Enabled:         No
Chassis ID:                                1
Floating IP:                               192.168.211.24 / 25
Unicast Election port:                     41473
Multicast IP:                              224.0.0.210
Multicast Port:                            41473
Version:                                   6.0.0-d.b280
Current Discover mode:                     Multicast

Members(* means local device):
ID  State       Priority IP:Port                    Location
```

```
------------------------------------------------------------------------
--------
1    vMaster(*)   200      192.168.105.121:41216        Local
2    vBlade       190      192.168.105.123:41216        Remote
3    vBlade       180      192.168.105.124:41216        Remote

Total: 3
```

# Configure Scaleout Cluster

1. To configure Scaleout on a vMaster device:

```
ACOS-vMaster[1/1](config:1)# scaleout 1
ACOS-vMaster[1/1](config:1-cluster:1)# local-device
This operation applied to device 1
ACOS-vMaster[1/1](config:1-cluster:1-local-device)# priority 210
This operation applied to device 1
ACOS-vMaster[1/1](config:1-cluster:1-local-device)# device-context 2
All the following configuration will go to device 2
ACOS-vMaster[1/1](config:2-cluster:1-local-device)# priority 220
This operation applied to device 2
ACOS-vMaster[1/1](config:2-cluster:1-local-device)# device-context 3
All the following configuration will go to device 3
ACOS-vMaster[1/1](config:3-cluster:1-local-device)# priority 200
This operation applied to device 3
ACOS-vMaster[1/1](config:3-cluster:1-local-device)# exit
ACOS-vMaster[1/1](config:3-cluster:1)# exit
ACOS-vMaster[1/1](config:3)# scaleout apps enable
```

2. To enable the aVCS cluster for Scaleout on all nodes:

```
ACOS-vMaster[1/1](config:1)# vcs database-distribution enable
This operation applied to device 1
ACOS-vMaster[1/1](config:1)# device-context 2
All the following configuration will go to device 2
ACOS-vMaster[1/1](config:2)#vcs database-distribution enable
This operation applied to device 2
ACOS-vMaster[1/1](config:2)# device-context 3
All the following configuration will go to device 3
ACOS-vMaster[1/1](config:3)# vcs database-distribution enable
```

```
This operation applied to device 3
ACOS-vMaster[1/1](config:3)# exit
ACOS-vMaster[1/1]# vcs reload
```

Run the `vcs reload` command on all nodes. Scaleout is up and running on all nodes.

3. Use the following show command to view Scaleout information:

```
ACOS-vMaster[1/1]# show scaleout

Device Role: Cluster Master

Cluster Mode: Layer-3

Device 1 - Active (Local) (Master)
Device 2 - Active
Device 3 - Active
```

4. Use the following show command to view the aVCS information.

```
ACOS-vMaster[1/1]# show vcs summary

aVCS Chassis:
VCS Configuration-Sync Enabled:            No
VCS Database-distribution Enabled:         Yes
Chassis ID:                                1
Floating IP:                               192.168.211.24 / 25
Unicast Election port:                     41473
Multicast IP:                              224.0.0.210
Multicast Port:                            41473
Version:                                   6.0.0-d.b280
Current Discover mode:                     Multicast

Members(* means local device):
ID  State       Priority IP:Port                    Location
--------------------------------------------------------------------------
--------
1   vMaster(*)  200      192.168.105.121:41216      Local
2   vBlade      190      192.168.105.123:41216      Remote
3   vBlade      180      192.168.105.124:41216      Remote
```

```
Total: 3
```

The *VCS Configuration-Sync Enabled* field displays as No if aVCS is not enabled.

5. Use the following show command to view the aVCS information.

```
ACOS-vMaster[1/1]# show vcs summary

aVCS Chassis:
VCS Configuration-Sync Enabled:            Yes
VCS Database-distribution Enabled:         Yes
Chassis ID:                                1
Floating IP:                               192.168.211.24 / 25
Unicast Election port:                     41473
Multicast IP:                              224.0.0.210
Multicast Port:                            41473
Version:                                   6.0.0-d.b280
Current Discover mode:                     Multicast

Members(* means local device):
ID  State       Priority IP:Port                     Location
--------------------------------------------------------------------------------
1   vMaster(*)  200      192.168.105.121:41216       Local
2   vBlade      190      192.168.105.123:41216       Remote
3   vBlade      180      192.168.105.124:41216       Remote

Total: 3
```

The *VCS Configuration-Sync Enabled* field displays as Yes if aVCS is enabled.

The following is the show running configuration for the aVCS Multicast mode:

```
!
vrrp-a common
  device-id 1
  set-id 1
!
device-context 1
  vcs database-distribution enable
  vcs enable
!
```

```
device-context 2
  vcs database-distribution enable
  vcs enable
!
device-context 3
  vcs database-distribution enable
  vcs enable
!
vcs floating-ip 192.168.211.24 / 25 255.255.255.240
!
vcs device 1
  priority 200
  interface management
  enable
!
vcs device 2
  priority 190
  interface management
  enable
!
vcs device 3
  priority 180
  interface management
  enable
!
!
scaleout 1
 device-context 1
  local-device
   priority 210
device-context 2
 local-device
  priority 220
device-context 3
 local-device
  priority 200
!
scaleout apps enable
!
```

This following is the show running configuration for the aVCS Unicast mode:

```
!
vrrp-a common
  device-id 1
  set-id 1
!
device-context 1
  vcs database-distribution enable
  vcs enable
!
device-context 2
  vcs database-distribution enable
  vcs enable
!
device-context 3
  vcs database-distribution enable
  vcs enable
!
vcs floating-ip 192.168.211.24 / 25 255.255.255.240
!
vcs device 1
  priority 200
  enable
!
vcs device 2
  priority 190
  enable
!
vcs device 3
  priority 180
  enable
!
vcs unicast-election
  members
    ip-address 2.2.2.115
    ip-address 2.2.2.116
    ip-address 2.2.2.117
!
vcs discovery-mode Unicast
!
```

```
!
scaleout 1
 device-context 1
  local-device
   priority 210
device-context 2
 local-device
  priority 220
device-context 3
 local-device
  priority 200
!
scaleout apps enable
!
```

# Enable Scaleout

Scaleout can be enabled for CGN, Gi/SGi Firewall, and standalone Firewall.

## Prerequisites

- A scaleout cluster must be configured.
- No application configuration must be present prior to enabling scaleout.

Use the following command to enable Scaleout:

```
ACOS(config)# scaleout apps enable
```

**NOTE:** If you disable Scaleout using the **`no scaleout apps enable`** command, reload the device for the change to take effect.

# Election of the Cluster Master

Each ACOS device can be configured with a priority value that is used during the master election process. Devices with a higher priority have precedence over devices with lower priorities.

Within the Scaleout cluster, the following summarizes the states of the roles supported:

- When Scaleout is not configured or is inactive, a message “Scaleout is not Active” is displayed. Use the command “`show scaleout`” to view the device status.
- During Scaleout initialization and election, the node has the “Unknown Node” role.
- Once the cluster master is elected, the Master node has the “Cluster Master” role, while other nodes in the cluster have the “Service Node” role.

However, since the election process is non-preemptive, if an existing master is present, it will retain its master status. Among devices that have an equal priority value, any one of them can potentially be selected as the master.

## Add a Device to the Cluster

By default, all device additions to a cluster are graceful; either when a new device joins an existing cluster, or a previously disabled device is re-enabled.

In both cases:

1. The device notifies the cluster master that it wants to join the cluster.

   If any of the existing devices in the cluster have open sessions for which this new device will be the new active device, the sessions are synced to this device.

2. The new device takes over the sessions after a short delay and starts processing the traffic actively.

To add a device to the Scaleout cluster in the aVCS unicast mode, perform the following:

```
ACOS(config)# vcs database-distribution enable
ACOS(config)# vcs unicast-election
ACOS(config-unicast-election)# port 41217
ACOS(config-unicast-election)# members
ACOS(config-unicast-election-members)# ip-address 2.2.2.115
ACOS(config-unicast-election-members)# ip-address 2.2.2.116
ACOS(config-unicast-election-members)# ip-address 2.2.2.117
ACOS(config-unicast-election-members)# ip-address 2.2.2.118
ACOS(config-unicast-election-members)# exit
```

```
ACOS(config)# vcs discovery-mode unicast
ACOS(config)# vcs device 4
ACOS(config-device:4)# priority 125
The changed configuration of aVCS will take effect only after 'vcs
reload'.
ACOS(config-device:4)# enable
ACOS(config)# vcs reload
```

After configuring the new device, the same must be added to the existing devices of aVCS on vMaster.

```
ACOS-vMaster[1/1]#config
ACOS-vMaster[1/1](config:4)#device-context 1
All the following configuration will go to device 1
ACOS-vMaster[1/1](config:1)#vcs unicast-election
ACOS-vMaster[1/1](config:1-unicast-election)# members
ACOS-vMaster[1/1](config:1-unicast-election-members)#ip-address 2.2.2.118

The changed configuration of aVCS will take effect only after 'vcs
reload'.
Run 'vcs reload cluster-discovery' if only aVCS unicast member IP address
was changed.

ACOS(config-device 1)# vcs reload cluster-discovery

System configuration has been modified. Save? [yes/no]: yes
Building configuration…
Write configuration to default primary startup-config
[OK]
Running configuration is saved
```

To add a device to the Scaleout cluster in the aVCS multicast mode, perform the following:

```
ACOS(config)# vcs database-distribution enable
ACOS(config)# vcs device 1
ACOS(config-device:1)# priority 126
ACOS(config-device:1)# interface management
ACOS(config-device:1)# enable
The changed configuration of aVCS will take effect only after 'vcs
reload'.
ACOS(config)# vcs reload
```

```
System configuration has been modified. Save? [yes/no]:yes
Building configuration...
Write configuration to default primary startup-config
[OK]
Running configuration is saved
```

# Remove a Device from the Cluster

In a shutdown scenario:

1. The device notifies the cluster master that it will be shut down. This triggers the cluster master to assign a new owner for each bucket.
2. The old owner (device going to be shut down) performs session sync to the new owner for all Layer 4 connections.
3. The cluster master updates the bucket mappings and pushes the configuration to all active devices in the cluster.
4. Since existing Layer-7 sessions are processed by the old owner, the admin should wait for all Layer-7 sessions on the device to close, at which point it is safe to shut down the device.

To remove a device from the Scaleout cluster in the aVCS unicast mode, perform the following:

```
ACOS(config)# vcs database-distribution disable
ACOS(config)# vcs unicast-election
ACOS(config-unicast-election)# port 41217
ACOS(config-unicast-election)# members
ACOS(config-unicast-election-members)# ip-address 2.2.2.115
ACOS(config-unicast-election-members)# ip-address 2.2.2.116
ACOS(config-unicast-election-members)# ip-address 2.2.2.117
ACOS(config-unicast-election-members)# no ip-address 2.2.2.118
ACOS(config-unicast-election-members)# exit

ACOS-vMaster[1/1]#config
ACOS-vMaster[1/1](config:4)#device-context 1
All the following configuration will go to device 1
ACOS-vMaster[1/1](config:1)#vcs unicast-election
```

```
ACOS-vMaster[1/1](config:1-unicast-election)# members
ACOS-vMaster[1/1](config:1-unicast-election-members)# no ip-address
2.2.2.118

The changed configuration of aVCS will take effect only after 'vcs
reload'.
Run 'vcs reload cluster-discovery' if only aVCS unicast member IP address
was changed.

ACOS(config-unicast-election-members)# vcs reload cluster-discovery

System configuration has been modified. Save? [yes/no]: yes
Building configuration…
Write configuration to default primary startup-config
[OK]
Running configuration is saved
```

To remove a device from the Scaleout cluster in the aVCS multicast mode, perform the following:

```
ACOS(config)# vcs database-distribution disable
ACOS(config)# no vcs device 1
```

# Configuration Notes

Consider the following point:

In a Scaleout cluster, all devices must run the same ACOS version.

# L2 Redirection Configuration

The following commands configure the local device in a cluster, assign a priority, and configure an interface for L2 redirection:

```
ACOS(config)# scaleout 2
ACOS(config-cluster:2)# local-device
ACOS(config-cluster:2-local-device)# priority 100
ACOS(config-cluster:2-local-device)# l2-redirect interface ethernet 5 vlan
60
```

**NOTE:** At the local device configuration level, use the `no l2-redirect` command to remove the L2-redirect configuration.

# L3 Redirection Configuration

The following commands configure the interface for L3 redirection:

```
ACOS(config)# scaleout 10
ACOS(config-cluster:2)# local-device
ACOS(config-cluster:2-local-device)# priority 100
ACOS(config-cluster:2-local-device)# cluster-mode layer-3
ACOS(config-cluster:2-local-device)# traffic-redirection
ACOS(config-cluster:2-local-device-traff...)# interfaces
ACOS(config-cluster:2-local-device-traffic-redirection-int...)# ethernet
<1-12>
OR
ACOS(config-cluster:2-local-device-traffic-redirection-int...)# trunk <1-
4096>
OR
ACOS(config-cluster:2-local-device-traffic-redirection-int...)# ve <2-
4094>
```

The following command configures the loopback interface:

```
ACOS(config-cluster:2-local-device-traffic-redirection-int...)# loopback
<0-10>
```

The following commands configure the cluster nodes to send the packets to the L3V partition using the Shared Partition redirect table:

```
ACOS(config-cluster:2-local-device)# traffic-redirection
ACOS(config-cluster:2-local-device-traff...)# session-sync
ACOS(config-cluster:2-local-device-traff...)# follow-shared
```

The following configuration skips the use of default routes to forward the redirected or session synchronized packets:

```
ACOS(config-cluster:2-local-device)# traffic-redirection
ACOS(config-cluster:2-local-device-traff...)# reachability-options
ACOS(config-cluster:2-local-device-traff...)# skip-default-route
```

# Configure Scaleout Session Synchronization

This chapter describes the overview of Scaleout session synchronization and configuration example.

The following topics are covered:



## Overview

Consider the following:

- While configuring the session synchronization interface, there is no option to assign preference or priority for the interface.
- For each interface type, a maximum of 6 entries can be configured, with the following order of default preference among interface types:
  - Ethernet
  - Trunk
  - VE
- Once the primary interface has been elected, adding a new interface with higher preference will not trigger re-election of the primary interface.
- If no primary interface is elected when multiple interfaces are added, the aforementioned interface preference order is used to elect the interface for session synchronization.
- If the IPv4 subnet on the primary interface used for session synchronization is removed, ACOS will match another IPv4 subnet based on the preference order to ensure that the new IPv4 subnet will match the cluster device's subnets.
- When an IP address on a neighbor scaleout node is added to the cluster, if a primary interface is not present, then ACOS will perform primary interface election. When an IP address on a neighbor scaleout node is deleted from the cluster, then ACOS will perform primary interface election.

- When adding new Scaleout cluster nodes, configure at least 1 interface with a subnet that can be shared among other cluster nodes for session synchronization purposes. If not, scaleout session synchronization will fail among all cluster nodes.
- To configure the session synchronization interface, use the **`session-sync-interface`** command under the local-device configuration.
- When a node goes down, the traffic user groups that were active on that node switch over to another node. When the failed node comes back up, some traffic user groups will be active on it again. For scaleout to become operational, configure the time scaleout must wait to join the cluster after the device boot up. Use the `start-delay` command in the scaleout cluster local device configuration to configure the delay time.

```
ACOS (config-cluster:1-local-device)# start-delay <10-300>
```

- When you enable session synchronization, a native and seamless failover mechanism with session L3-L4 continuity is provided. Active sessions from a node are synchronized to the backup nodes (based on traffic-group mapping) and can be used immediately after failover. Syncing sessions between nodes consumes session table space; each node maintains backup sessions for other nodes, which increases the number of sessions stored on each node. As a result, the session capacity is reduced by half.

## Configuration Example

The following commands enter the scaleout local device configuration mode to configure the local device in a cluster.

```
ACOS(config)# scaleout 1
ACOS(config-cluster:1)# local-device
ACOS(config-cluster:1-local-device)# priority 200
```

The following commands configure the session synchronization interface for the local device:

```
ACOS(config-cluster:1-local-device)# session-sync
ACOS(config-cluster:1-local-device-sess...)# interface
ACOS(config-cluster:1-local-device-sess...)# ethernet 1
ACOS(config-cluster:1-local-device-sess...)# trunk 1
ACOS(config-cluster:1-local-device-sess...)# ve 50
ACOS(config-cluster:1-local-device-sess...)# exit
```

```
ACOS(config-cluster:1-local-device)# exit
```

# Configure aVCS Synchronization

With the VCS database synchronization feature, you can enable and disable configuration synchronization and database synchronization separately. This can be done on any node within the Scaleout cluster.

**Scenario 1**: If you disable configuration and database synchronization on vBlade. The aVCS becomes inactive on vBlade.

```
ACOS-vBlade#vcs disable
ACOS(config)#vcs database-distribution disable
```

To check the status, use the following command:

```
ACOS(config)#show vcs summary
VCS is not active.

ACOS(config)#show scaleout
Device Role: Scaleout is not active.
```

**Scenario 2**: If you disable configuration synchronization on vBlade while keeping the database synchronization active:

```
ACOS-vBlade#vcs disable
```

To check the status, use the following command:

```
ACOS(config)#show vcs summary
aVCS Chassis:
VCS Configuration-Sync Enabled:            No
VCS Database-distribution Enabled:         Yes
Chassis ID:                                15
Unicast Election port:                     41473
Multicast IP:                              224.0.1.210
Multicast Port:                            41473
Version:                                   6.0.1.b4
Current Discover mode:                     Multicast
Members (* means local device) (C:cfg-sync | D:db-sync | B:both):
ID  State        PriorityIP:Port
Location
```

```
-------------------------------------------------------------------------
-----
1   vBlade(D)(*) 0       10.64.19.32:41216
Local
2   vBlade(B)    0       10.64.19.33:41216
Remote
3   vMaster(B)   0       10.64.19.34:41216
Remote
Total: 3
```

Here,

- C represents that configuration synchronization is active on the node.
- D represents that database synchronization is active on the node.
- B represents that both configuration and database synchronization is active on the node.

**Scenario 3**: If you disable database synchronization on vBlade while keeping the configuration synchronization active:

```
ACOS-vMaster[1/1](config:3)#device-context 1
All the following configuration will go to device 1
ACOS-vMaster[1/1](config:1)#vcs database-distribution disable
This operation applied to device 1
```

To check the status, use the following command:

```
ACOS(config)#show vcs summary
aVCS Chassis:
VCS Configuration-Sync Enabled:            Yes
VCS Database-distribution Enabled:         No
Chassis ID:                                15
Unicast Election port:                     41473
Multicast IP:                              224.0.1.210
Multicast Port:                            41473
Version:                                   6.0.1.b4
Current Discover mode:                     Multicast
Members (* means local device) (C:cfg-sync | D:db-sync | B:both):
ID  State        PriorityIP:Port
Location
-------------------------------------------------------------------------
-----
```

```
1    vBlade(C)(*) 0        10.64.19.32:41216
Local
2    vBlade(B)    0        10.64.19.33:41216
Remote
3    vMaster(B)   0        10.64.19.34:41216
Remote
Total: 3
```

Scenario 4: If you run disable command on vMaster, a message is displayed as shown below. If Yes, it will shut down aVCS on vMaster, even if only one of the synchronization (configuration or database) is disabled.

```
ACOS-vMaster[1/3](config:3)#vcs database-distribution disable
You are running this command on vMaster, and it will shutdown both
database-distribution and configuration synchronization. Confirm?
[yes/no]:yes
Confirmed
This operation applied to device 3

OR

ACOS-vMaster[1/3]#vcs disable
You are running this command on vMaster, and it will shutdown both
database-distribution and configuraiton synchronization.Confirm?
[yes/no]:yes
Confirmed
```

To check the status, use the following command:

```
ACOS(config)#show vcs summary
VCS is not active.
```

Scenario 5: If you disable configuration or database synchronization on node 1 while the other node has both configuration and database synchronization enabled and try to take-over on the node 1.

The vBlade cannot take over as vMaster if it has fewer services than vMaster.

```
ACOS(config)#show vcs summary
aVCS Chassis:
VCS Configuration-Sync Enabled:            Yes
VCS Database-distribution Enabled:         No
Chassis ID:                                15
```

```
Unicast Election port:                          41473
Multicast IP:                                   224.0.1.210
Multicast Port:                                 41473
Version:                                        6.0.0-P1.b4
Current Discover mode:                          Multicast
Members (* means local device) (C:cfg-sync | D:db-sync | B:both):
ID  State         PriorityIP:Port
Location
-------------------------------------------------------------------------
-----
1   vBlade(c)(*) 0       10.64.19.32:41216
Local
2   vBlade(B)    0       10.64.19.33:41216
Remote
3   vMaster(B)   0       10.64.19.34:41216
Remote
Total: 3
```

Try to take-over on the node 1:

```
ACOS#vcs vmaster-take-over 1
Cannot switch-over because I am running less service than other(s)
```

# Configure aVCS Reload

During the vcs reload process, you can use the following commands:

- To apply the configuration changes on the aVCS chassis:

  ```
  ACOS(config)# vcs reload cluster-discovery
  ```

- To prevent configuration information from being migrated from vBlade to vMaster:

  ```
  ACOS(config)# vcs reload disable-merge
  ```

- To reload a specific device when aVCS is enabled:

  ```
  ACOS(config)# vcs reload device <device-id>
  ```

  The device ID can be between 1-16.

- To keep the Scaleout cluster uninterrupted:

```
ACOS(config)# vcs reload db-safe start
```

- To end the db-safe mode automatically based on the specified timeout value:

```
ACOS(config)# vcs reload db-safe start timeout <seconds>
```

The timeout value is 5-300 seconds (about 5 minutes).

- To complete the db-safe vcs reload within the specified timeout or after the timeout:

```
ACOS (config)# vcs reload db-safe complete force
```

For example, the vcs reload is run in the db-safe mode with a timeout of 100 seconds.

```
ACOS-vMaster[1/1](config:1)#vcs reload db-safe start timeout 100
```

To check the status, use the following command:

```
ACOS-vMaster[1/1]#show vcs summary

aVCS Chassis:
VCS Configuration-Sync Enabled:            Yes
VCS Database-distribution Enabled:         Yes (DB-SAFE-RELOAD-MODE)
[expire: 90 seconds]
Chassis ID:                                1
Floating IP:                               10.240.163.206
Mask:                                      255.255.255.240
Unicast Election port:                     12121
Multicast IP:                              224.0.1.210
Multicast Port:                            41473
Version:                                   6.0.1-d.b119
Current Discover mode:                     Unicast

Members (* means local device) (C:cfg-sync | D:db-sync | B:both):
ID  State        PriorityIP:Port
Location
-------------------------------------------------------------------------
-------
1   vMaster(B)(*)125     172.16.10.100:41216
Local
```

```
3    vBlade(B)     123      172.16.10.102:41216
Remote
4    vBlade(B)     122      172.16.10.103:41216
Remote
Total: 3
```

When db-safe vcs is reloading, the DB-SAFE-RELOAD-MODE message displays in the VCS Databse-distribution Enabled field, along with the running timer or timeout.

# Configure Service-Config Template

The following commands configure the service-config template. This is currently only supported with CGN.

The user group count must be a power of 2 value starting from 1, 2, 4, … 256.

```
ACOS(config)#scaleout 1
ACOS(config-cluster:1)#service-config
ACOS(config-cluster:1-service-config)#template <name>
ACOS(config-cluster:1-service-config-name)#user-group-count <num>
```

**NOTE:**

- To change the user group count, you must unbind the service-config template from the application (like NAT pool), make the change, and rebind it. OR make the change without unbinding, but you must clear all the existing sessions for the application traffic slice.
- The number of service-config templates is limited system-wide on a platform basis.

# Configure User Group

The following commands configure shared and per-partition user group count. It can be applied independently both in the shared and L3V partitions:

```
ACOS(config)#scaleout1
ACOS(config-cluster:1)#service-config
ACOS(config-cluster:1-service-config)#user-group-count <num>
```

# Configure Policy-Based Failover

This chapter describes how to use policy-based failover templates to trigger failover actions in Scaleout and how to use routing protocols to withdraw or redistribute routes when a Scaleout node is down or exits the cluster.

The following topics are covered:

# Overview

ACOS supports event tracking and policy-based failover support via a template. The template provides a flexible way of defining events that will trigger failover once a tracked event takes place.

Using a policy-based failover template, you can allocate a weight of 1-255 for each event. When the event occurs, the cost of the policy-based template increases. For example, the Ethernet interface 2 connected to a node in the Scaleout cluster is assigned a weight of 100. When the Ethernet interface is up and running, the policy cost is 0. When the Ethernet interface 2 goes down, the cost increases to 100, possibly causing the failover.

Policy-based failover templates are available in both shared and L3V partitions. When a default tracking template is configured in a partition, it affects all the IPv4 and IPv6 traffic. If the tracking template is not configured in the L3V partition, it follows the tracking templates configured in the shared partition.

You can define the policy-based failover template on a per-partition basis for IPv4 traffic. This template affects all IPv4 traffic slices and IPv4 default traffic. Similarly, you can define a policy-based failover template on a per-partition basis for IPv6 traffic (like IPv6 Firewall traffic).

For example, one tracking template can work for an IPv4 data flow's traffic slice and another can work for an IPv6 data flow (like IPv6 Firewall).

When there is no tracking template in a partition or in a shared partition, the traffic map updates are not subject to any policy-based failover template. However, the traffic map updates will be affected by node-level changes such as a crash, power-off, or reboot.

Multiple policy-based failover templates can be created and associated to a Scaleout cluster.

## Events Tracked via Policy-based Failover Templates

The events that can be tracked from the policy-based failover templates are as follows:

Table 3 : Event Tracking

| Event | Description |
|---|---|
| BGP neighbor unreachable | If a BGP neighbor is unreachable, the cost of the policy-based template increases causing the failover. |
| Lost link in a default gateway | The failover template periodically checks connectivity to the IPv4 and IPv6 default gateways connection to a real server by pinging them. If a gateway stops responding, the cost of the template increases resulting in failover.<br><br>The weight can be specified individually for each gateway's IP address that you configure in the tracking events list.<br><br>**Note**: You must configure the ACOS device to track the status of the gateway. |
| Lost link on an Ethernet port | If the link on an Ethernet data port goes down, the cost of the policy-based template increases.<br><br>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 to 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 will show that port 1 is up since it is the lead port for the trunk. |
| Lost data route | If an IPv4 or IPv6 route matching the specified options is not in the data route table, the cost of the policy-based template increases. |
| Lost link on a trunk | If the trunk or individual ports in the trunk go down, the cost of the policy-based template increases.<br><br>To track a trunk port within an L3V partition, you must verify that the tracked port is being used within that partition:<br><br>If the tracked trunk is down, the cost is increased based on the trunk value and ignores the status of the ports within the trunk. |

Table 3 : Event Tracking

| Event | Description |
|---|---|
| | If the tracked trunk is up, it checks the ports within the trunk, and if any of them are down, the failover template increases the cost. If 2 ports are down, the template increases the cost by twice. |
| VLAN inactivity | If ACOS stops detecting traffic on a VLAN, the cost of the policy-based template increases. |

The cost of a failover template is calculated based on the total cost of the events configured in the template.

## Configure Policy-Based Failover Templates

Use the `resource-track` command to configure the policy-based failover templates.

To configure a policy-based failover template:

1. At the global configuration level, create a policy-based failover template.

```
ACOS(config)#resource-track template_1
ACOS(config-resource-track:template_1)#
```

**NOTE:** Resource tracking updates come into effect on time upon the changes in the status of the tracked events such as link status, BGP peer status, Gateway, and so on. The events are not dampened and are directly delivered to the resource tracking module.

2. At the config-resource-track configuration level, configure any or all of the following events:
   a. **`bgp`**—Specify the BGP IP address and assign a weight for the BGP in the event of a failure. For example, if BGP 12.12.10.1 fails, the cost is increased to 100:

```
ACOS(config-resource-track:template_1)#bgp 12.12.10.1 weight 100
```

   b. **`gateway`**—Before configuring the gateway event, the ACOS device must have a gateway tracking option configured. To configure the tracking option, use the following commands:

```
ACOS(config)#cgnv6 server gateway 10.10.10.1
ACOS(config-real server)#health-check gateway
ACOS(config-real server)#exit
```

Specify 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, the cost is increased to 100:

```
ACOS(config-resource-track:template_1)#gateway 10.10.10.1 weight
100
```

c. **interface**— Specify 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, the cost increases to 40:

```
ACOS(config-resource-track:template_1)#interface ethernet 1 weight
40
```

d. **route**—Specify 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, the cost increases to 100:

```
ACOS(config-resource-track:template_1)#route 20.20.20.1 /24 weight
100
```

e. **trunk**—Specify the trunk identification number and assign a weight for that trunk in the event of failure. For example, if trunk 1 fails, the cost increases to 20:

```
ACOS(config-resource-track:template_1)#trunk 1 weight 20
```

f. **VLAN**—Specify the VLAN identification number, the timeout value, and assign a weight for that trunk in the event of failure. For example, if VLAN 2 fails, the cost increases to 30:

```
ACOS(config-resource-track:template_1)#vlan 2 timeout 20 weight 30
```

3. Assign the failover template to a Scaleout cluster. For more information on associating the template, see Associate Failover Templates to Scaleout Cluster.

## Displaying the Configured Failover Template

Use the **show resource-tracked** command to view all the configured templates.

The command is as follows:

```
ACOS (config)#show resource-tracked
Resource Tracking Name: BGP
 bgp 20.20.20.1 weight 50


User-Idx 1 | User name MULTI_scaleout_BGP_multi_PART-2 | Cost 50
User-Idx 2 | User name MULTI_scaleout_BGP_multi_PART-3 | Cost 50
User-Idx 3 | User name MULTI_scaleout_BGP_abc_PART-2 | Cost 50
User-Idx 4 | User name MULTI_scaleout_BGP_abc_PART-3 | Cost 50


Totally 4 event(s) tracked
```

Use the **`show resource-tracked [template name]`** command to view a specific template.

The command is as follows:

```
ACOS (config)#show resource-tracked BGP
Resource Tracking Name: BGP
 bgp 20.20.20.1 weight 50


User-Idx 1 | User name MULTI_scaleout_BGP_multi_PART-2 | Cost 50
User-Idx 2 | User name MULTI_scaleout_BGP_multi_PART-3 | Cost 50
User-Idx 3 | User name MULTI_scaleout_BGP_abc_PART-2 | Cost 50
User-Idx 4 | User name MULTI_scaleout_BGP_abc_PART-3 | Cost 50


Totally 4 event(s) tracked
```

Use the **`show resource-tracked-by-user`** command to view a template based on user information.

```
ACOS (config)#show resource-tracked-by-user


User-Idx 1| User name MULTI_scaleout_BGP_multi_PART-2 | Cost 50
 Resource Tracking Name: BGP
   bgp 20.20.20.1 weight 50
```

```
User-Idx 2| User name MULTI_scaleout_BGP_multi_PART-3 | Cost 50
 Resource Tracking Name: BGP
      bgp 20.20.20.1 weight 50
```

```
User-Idx 3| User name MULTI_scaleout_BGP_abc_PART-2 | Cost 50
Resource Tracking Name: BGP
      bgp 20.20.20.1 weight 50
```

```
User-Idx 4| User name MULTI_scaleout_BGP_abc_PART-3 | Cost 50
Resource Tracking Name: BGP
      bgp 20.20.20.1 weight 50

Totally 4 event(s) tracked
```

## Associate Failover Templates to Scaleout Cluster

One or more failover template can be associated to a Scaleout cluster to track events such as the operational state of BGPs, gateways, interfaces, trunks, and VLANs. When the operational state of one or any of the configured events goes down, the cost of the template increases, which results in executing the configured action on the Scaleout node.

For example, when an Ethernet port connected to a Scaleout node is down, all traffic to this port is dropped. To avoid such packet drops, a tracking template is configured at the config-cluster configuration level, and the policy-based failover template is associated. Scaleout inquires the tracking templates bound to all the Scaleout clusters every second to check if there is any change in the events configured in the policy-based failover template. When there is a change in the event, based on the cost of the template and the threshold value set, the Scaleout node is disabled or exited from the cluster.

When you associate a template to a Scaleout cluster, a threshold value is specified. When the cost of the template is below the threshold value, Scaleout continues to be in the active state. When the cost exceeds the threshold value, then the configured failover action is executed.

You can configure 2 failover actions in the templates. They are down and exit-cluster.

- down—Specifies that the Scaleout node remains in the cluster.
- exit-cluster—Specifies that the Scaleout node leaves the cluster.

## Failover Template Configuration Example

The following example creates a failover policy-based template named policy_ template. This template tracks the health of Ethernet interface 2 and gateway IP address 15.15.1.0. The template is then bound to Scaleout 1 for tracking the events.

1. At the global configuration level, create a policy-based template named **policy_ template**.

```
ACOS(config)#resource-track policy_template
```

2. At the config-resource-track configuration level, configure the interface Ethernet and gateway events to track their operational state.

```
ACOS(config-resource-track:policy_template)#interface ethernet 2 weight
40
ACOS(config-resource-track:policy_template)#gateway 15.15.1.0 weight
100
```

3. Bind the newly created **policy_template** to **Scaleout 1**.

```
ACOS(config)#scaleout 1
ACOS(config-cluster:1)#local-device
ACOS(config-cluster:1-local-device)#tracking-template
ACOS(config-cluster:1-local-device-tracking-template)#template policy_
template
ACOS(config-cluster:1-local-device-tracking-template-t...)#threshold
100 down
ACOS(config-cluster:1-local-device-tracking-template-t...)#threshold
250 exit-cluster
```

4. Verify the cost of the template and scaleout status before failover.

**Display the failover template configuration**

```
ACOS(config)#show resource-tracked policy_template
Resource Tracking Name: policy_template
```

```
interface ethernet 2 weight 40
gateway 15.15.1.0 weight 100

User-Idx 1 | User name policy_template | Cost 0

Totally 2 event(s) tracked
```

**Display the Scaleout Status**

```
ACOS(config-cluster:1-local-device-tracking-template-t...)#show
scaleout
Role - Service Node

Device 1 - Active (Local)
Device 2 - Active
Device 3 - Active
Device 4 - Active
```

When the Ethernet interface 2 and gateway IP address tracked by the failover template are up and running, the cost is 0. The Scaleout status is also Active.

5. View the cost of the template and the scaleout status after the failover event.

   In this example, the gateway 15.15.1.0 stops responding and the cost of the event is increased to 100.

   **Display the failover template configuration**

```
ACOS(config)#show resource-tracked policy_template
Resource Tracking Name: policy_template
interface ethernet 2 weight 40
gateway 15.15.1.0 weight 100

User-Idx 1 | User name policy_template | Cost 100

Totally 2 event(s) tracked
```

   **Display the Scaleout Status**

```
ACOS(config-cluster:1-local-device-tracking-template-t...)#show
scaleout
Role - Service Node
```

```
Device 1 - Disabled (Local)
Device 2 - Active
Device 3 - Active
Device 4 - Active
```

As the threshold is set to 100 and the action configured is down, device 1 in the Scaleout cluster is moved to the disabled state.

## Associate Failover Multi-Template Example

The following example shows how to associate a multi-template to a Scaleout cluster to resource track L3V partition objects such as BGP or routes or Gateway across the partition.

**NOTE:** A maximum of 8 multi-templates are supported.

```
ACOS(config)#scaleout 1
ACOS(config-cluster:1)#local-device
ACOS(config-cluster:1-local-device)#priority 1
ACOS(config-cluster:1-local-device)#tracking-template
ACOS(config-cluster:1-local-device-tracki...)#multi-template multi
ACOS(config-cluster:1-local-device-tracki...)#template BGP partition PART-
2
ACOS(config-cluster:1-local-device-tracki...)#template BGP partition PART-
3
ACOS(config-cluster:1-local-device-tracki...)#threshold 100 down
```

The following is the show running configuration in the shared partition.

```
scaleout 64
  local-device
   priority 1
   tracking-template
    multi-template multi
     template BGP partition PART-2
     template BGP partition PART-3
     threshold 100 down
```

The following is the show running configuration in the respective partition.

```
ACOS(PART-2)#show run
```

```
active-partition PART-2

resource-track BGP
bgp 20.20.20.1 weight 50

ACOS(PART-3)#show run
active-partition PART-3

resource-track BGP
bgp 20.20.20.1 weight 50
```

## Limitations

- Tracking templates must be configured before binding them to a Scaleout cluster.
- Up to 2 templates can be bound to a single Scaleout cluster.
- Per template, up to two thresholds and actions can be configured.
- Operationally disabled Scaleout node takes preference over configured resource threshold action. For example, if a Scaleout node is already disabled from CLI, a change in the event will not trigger any action.
- The Exit-Cluster will only be supported for the default template in the shared partition.
- The IPv4/IPv6 option cannot be changed dynamically. The tracking template must be removed and readded with the required option.

# Use Route Map to Withdraw or Redistribute Routes

Using route maps, Scaleout enables routing protocols to advertise or withdraw route advertisements when a Scaleout node is disabled or exits the cluster.

When the status of a Scaleout node changes from active to disabled, Scaleout triggers an event to the route map to update the Scaleout state. A rule configured in the route map checks the Scaleout node status specified in the rule with the current state of the Scaleout node in the cluster. If the state does not match, the route map will withdraw or redistribute routes from the neighboring router (if configured to advertise).

A rule can be configured in the route map to specify Scaleout whether to permit or deny the route advertisements.

Either an ipv4 or ipv6 option can be configured in the match command in the route map. It informs the routing protocol when the ipv4 or ipv6 tracking template for a node, or the default tracking template goes up or down.

For example,

```
route-map test1 permit 10
  match scaleout 1 [ipv4/ipv6] up/down
```

When the ipv4 option is specified in the match command, it matches when the state of the ipv4 tracking template, or the default tracking template goes up or down.

When the ipv6 option is specified in the match command, it matches when the state of the ipv6 tracking template, or the default tracking template goes up or down.

When neither ipv4 nor ipv6 option is specified, it implies a default value that matches the state of the default tracking template.

The route map can be bound to multiple bind points in BGP, OSPF, ISIS, and so on.

## Configure Route Map to Withdraw or Redistribute Routes

Use the **`match scaleout`** command at the route-map configuration level to specify the Scaleout node status in the route map rule. This rule is used to permit or deny routes based on the Scaleout status.

```
ACOS(config)# route-map 1 permit 10
ACOS(config-route-map)# match scaleout 1 up
```

## Display the Configured Route Map

To display the configured route maps, use the following command:

```
ACOS(config-route-map:10)#show route-map 1
route-map 1, permit, sequence 10
  Match clauses:
    scaleout 1 up
  Set clauses:
```

## Limitation

If the route-map is configured to use a Scaleout cluster id that is not configured on the system, all the routes on the protocols using route-map are either withdrawn or not distributed.

# Scaleout for Carrier Grade Networking (CGN)

This chapter describes the Scaleout feature in a CGN deployment.

The following topics are covered:

# Overview

In CGN Scaleout, inside users and all the available NAT IP addresses are separated into user groups. When a cluster node is added or removed, the associated user groups can be moved from one cluster node to another. The NAT IP addresses associated with that user group are also moved to the new cluster node.

For Fixed NAT, when mapping inside client IP addresses to public NAT IP addresses, inside users belonging to a user group will be mapped to NAT IP belonging to the same user group.

## Traffic Distribution

Traffic from upstream devices is distributed across the nodes based on the user group associated with the packet. Each node is responsible for processing packets belonging to a subset of user groups. In a network, all the traffic coming from the same subscriber may have different IPv6 addresses. However, ACOS expects the traffic coming from a subscriber to have the same prefix (of length IPv6 prefix length) in the source IPv6 address and accordingly identifies the subscriber associated with the packet based on the configured IPv6 prefix length. All the packets with the same prefix (of length IPv6 prefix length) in the source IPv6 address will map to the same user group and hence be processed by the same node in the Scaleout cluster. This will allow these packets to be assigned the same NAT IPv4 address (using the user-quota-prefix-length configuration). For more information about user-quota-prefix-length, see *IPv4-to-IPv6 Transition Solutions Guide*).

The `ipv6-prefix-length` command is a system-level attribute that indicates the length of the source and destination IPv6 prefix. This prefix is used to determine the user group for processing a packet.

The IPv6 prefix length can be configured only in the shared partition but is applicable to all L3V partitions, LIDs, and Class-lists. The default value of this attribute is set to 128 on all devices.

The IPv6 prefix length must be less than or equal to all configured values for the user-quota-prefix-length. If the IPv6 prefix length is greater than the user-quota-prefix-length, all packets with the same user-quota prefix will not receive the same

NAT IPv4 address. For more information about user-quota-prefix-length, see *IPv4-to-IPv6 Transition Solutions Guide*.

Use the following command to configure the IPv6 prefix length:

```
ACOS(config)#system ipv6-prefix-length length
```

You can select a value between 16 and 128 for length. By default, the IPv6 prefix length is 128.

**NOTE:** The IPv6 prefix length must be modified only when the Scaleout cluster is not in use. Else, unpredictable disruptions may occur to the traffic.

## Scaleout Mapping

In Scaleout mode, inside user addresses are mapped to user groups based on hashing calculation. The following illustrates how respective insider user addresses are mapped to a specific user group:

```
Address          User group
1.1.1.1               1
1.1.1.2               2
1.1.1.3               3
…

Address          User group
2.2.2.1               1
2.2.2.2               2
2.2.2.3               3
…
```

In this example, use the following command to configure the Fixed NAT mapping:

```
ACOS(config)# cgnv6 fixed-nat inside 1.1.1.0 1.1.2.255 netmask /16 nat
2.2.2.0 2.2.2.255 netmask /24
```

To view the port-mapping for the NAT IP address 2.2.2.1 and 2.2.2.2, use the following commands:

```
ACOS4(config:4)# show cgnv6 fixed-nat nat-address 2.2.2.1 port-mapping
NAT IP Address: 2.2.2.1
Inside User: 1.1.1.1
```

```
TCP:  1024 to 33279
UDP:  1024 to 33279
ICMP: 1024 to 33279
Inside User: 1.1.2.1
TCP:  33280 to 65535
UDP:  33280 to 65535
ICMP: 33280 to 65535

ACOS4(config:4)# show cgnv6 fixed-nat nat-address 2.2.2.2 port-mapping
NAT IP Address: 2.2.2.2
Inside User: 1.1.1.2
TCP:  1024 to 33279
UDP:  1024 to 33279
ICMP: 1024 to 33279
Inside User: 1.1.2.2
TCP:  33280 to 65535
UDP:  33280 to 65535
ICMP: 33280 to 65535
```

In a non-Scaleout mode, inside user addresses are mapped to NAT IPs using the default method as follows:

```
ACOS(config)# show cgnv6 fixed-nat nat-address 2.2.2.1 port-mapping
NAT IP Address: 2.2.2.1
Inside User: 1.1.1.1
TCP:  1024 to 22527
UDP:  1024 to 22527
ICMP: 1024 to 22527
Inside User: 1.1.1.2
TCP:  22528 to 44031
UDP:  22528 to 44031
ICMP: 22528 to 44031
Inside User: 1.1.1.3
TCP:  44032 to 65535
UDP:  44032 to 65535
ICMP: 44032 to 65535
```

```
ACOS(config)# show cgnv6 fixed-nat nat-address 2.2.2.2 port-mapping
NAT IP Address: 2.2.2.2
Inside User: 1.1.1.4
```

```
TCP:  1024 to 22527
UDP:  1024 to 22527
ICMP: 1024 to 22527
Inside User: 1.1.1.5
TCP:  22528 to 44031
UDP:  22528 to 44031
ICMP: 22528 to 44031
Inside User: 1.1.1.6
TCP:  44032 to 65535
UDP:  44032 to 65535
ICMP: 44032 to 65535
```

## NAT IP Route Aggregation and Redistribution

In CGN Scaleout mode, every node advertises the routes of the NAT IPs allocated to that node. For example, if there are 256 NAT IPs, each NAT IP is associated with a user group as follows:

| User Group | NAT IPs | Route Prefix |
|---|---|---|
| 0 | 15.15.1.0 | 15.15.1.0/32 |
| 1 | 15.15.1.1 | 15.15.1.1/32 |
| 2 | 15.15.1.2 | 15.15.1.2/32 |
| .<br>.<br>.<br>255 | .<br>.<br>.<br>15.15.1.255 | .<br>.<br>.<br>15.15.1.255/32 |

When the number of NAT IPs in a NAT pool is more than 256 IP addresses, ACOS aggregates or consolidates specific routes into a single route advertisement. The route aggregation helps to minimize the number of route advertisements and reduces the number of entries in the routing tables on the external server.

For route aggregation, the total number of NAT IPs in a NAT pool is divided into groups before allocating them to a user group. For example, consider the following configuration:

```
ACOS(config)#cgnv6 nat pool P1 15.15.1.0 15.15.4.255 netmask /24
```

In this example, there are 1024 NAT IPs. To distribute 1024 NAT IPs into 256 user groups, you can combine a group of 4 NAT IPs to form a single route and then associate this aggregated route to a single user group.

The following table summarizes the combination of NAT IPs in each user group and their route prefix.

| User Group | NAT IPs | Route Prefix |
|---|---|---|
| 0 | 15.15.1.0 to 15.15.1.3 | 15.15.1.0 /30 |
| 1 | 15.15.1.4 to 15.15.1.7 | 15.15.1.4 /30 |
| 2 | 15.15.1.8 to 15.15.1.11 | 15.15.1.8 /30 |
| .<br>.<br>.<br>255 | .<br>.<br>.<br>15.15.2.252 to 15.15.2.255 | .<br>.<br>.<br>15.15.2.252 /30 |

When the total number of IP addresses is equal to 3 * /24 or 7 * /24 or a value that is not equal to two to the power of n, then the NAT IPs are divided into sub-groups where each sub-group prefix is equal to two to the power of n.

Consider the following example.

```
ACOS(config)#cgnv6 nat pool P1 15.15.1.0 15.15.7.255 netmask /24
```

In this example, the total number of NAT IPs available is 1792. When 1792 NAT IPs are divided equally into 256 user groups, each user group contains 7 NAT IPs. ACOS cannot aggregate 7 NAT IPs to advertise a route. Therefore, the routes are divided into sub-groups.

The 7 * /24 routes are divided into sub-groups as follows—4 * /24 is aggregated as one route, 2 * /24 is aggregated as another route, and 1 * /24 is aggregated as yet another route.

Therefore, the maximum routes per pool are equal to (No. of user groups) * (No. of sub-groups in a pool).

The following table summarizes the combination of NAT IPs in each user group and their route prefix.

| User Group | NAT IPs | Route Prefix |
|---|---|---|
| 0 | 15.15.1.0 to 15.15.1.3<br>15.15.5.0 to 15.15.5.1<br>15.15.7.0 | 15.15.1.0 /30<br>15.15.5.0 /31<br>15.15.7.0 /32 |
| 1 | 15.15.1.4 to 15.15.1.7<br>15.15.5.2 to 15.15.5.3<br>15.15.7.1 | 15.15.1.4 /30<br>15.15.5.2 /31<br>15.15.7.1 /32 |
| 2 | 15.15.1.8 to 15.15.1.11<br>15.15.5.4 to 15.15.5.5<br>15.15.7.2 | 15.15.1.8 /30<br>15.15.5.4 /31<br>15.15.7.2 /32 |
| .<br>.<br>.<br>255 | .<br>.<br>.<br>15.15.2.252 to 15.15.2.255<br>15.15.5.254 to 15.15.5.255<br>15.15.7.255 | ..<br>.<br>.<br>15.15.2.252 /30<br>15.15.5.254 /31<br>15.15.7.255 /32 |

Depending on the state of scaleout, the following scenarios might occur:

- If the CGN Scaleout traffic map is stabilized, when NAT pool or Fixed NAT configuration is performed, routes are advertised immediately.
- If CGN Scaleout is not active when NAT pool or Fixed NAT configuration is performed, routes are advertised when Scaleout traffic map is stabilized and traffic map is changed.

When the addition of a node or failure of a node causes the traffic map to be updated, each node must advertise the newly owned NAT IPs and withdraw the formerly owned NAT IPs.

- As the number of NAT IPs increases, convergence time during a traffic map update may increase.
- During convergence time, packets may be sent to the wrong node. The misplaced packets will be redirected to the correct node via L2. Henceforth, L2 redirection is possible.

## Configuring NAT IP Route Aggregation

By default, the NAT IP routes are aggregated and then associated to user groups. Use the following CLI command to display the configuration:

```
ACOS(config)#cgnv6 scaleout nat-ip-hashing-scheme ?
  route-aggregation  Chunk contiguous NAT IPs for route aggregation
(default)
  mod-user-groups    Hash NAT IPs by taking mod of user-groups
```

## Displaying NAT IP to user group mapping

Use the following command to display the NAT IP to user group mapping in scaleout mode:

```
ACOS(config)#show cgnv6 scaleout address-mapping nat-address 1.1.1.1
User-Group                Active-Node     Standby-Node
-----------------------------------------------------
  1                          1                2
```

## Support for Hashing NAT IPs

This section helps to revert from the NAT IP route aggregation to the pre-4.1.4-P2 approach of hashing NAT IPs.

Prior to software version 4.1.4-P2, the NAT IPs were hashed similar to inside users and assigned to the user groups to which the inside users with the same hash value belonged. With this approach, the consecutive NAT IP addresses were mapped to different user groups making it impossible to aggregate NAT IPs for route advertisement. Starting in software version 4.1.4-P2, the NAT IP route aggregation is introduced to overcome these limitations.

The default hashing scheme for NAT IP is route aggregation. To revert to the old approach of hashing NAT IPs, use the `cgnv6 scaleout nat-ip-hashing-scheme` command before CGN Scaleout is enabled as it requires reboot. To preserve the hashing scheme while upgrading from versions prior to 4.1.4-P2, add the following command in the start-up config file:

```
ACOS(config)#cgnv6 scaleout nat-ip-hashing-scheme ?
  route-aggregation  Chunk contiguous NAT IPs for route aggregation
(default)
  mod-user-groups    Hash NAT IPs by taking mod of user-groups
```

```
ACOS(config)#cgnv6 scaleout nat-ip-hashing-scheme mod-user-groups
```

### Recommendations

The following are recommended:

- As a minimum, the number of NAT IPs required for CGN Scaleout must be 256 consecutive addresses.
- When using hairpinning in CGN Scaleout:
  - Endpoint-Independent Mapping (EIM) and Endpoint-Independent Filtering (EIF) must be enabled.
  - If the destination NAT IP doesn't belong to the same user group as the source IP, the packet is forwarded to a router, which then is routed to the correct node.
- The range of addresses in the CGN class-list and Fixed-NAT inside users should match only CGN client addresses.

## RADIUS Message Distribution in a Cluster

When a RADIUS message is received on a node, ACOS identifies the traffic map associated with the IP address (IPv4 or IPv6) and checks the active status of the node using a class-list lookup. If the node is not the active node, the message is redirected to the active node for processing. On the active node, a Session Management Protocol (SMP) entry is created based on the content of the RADIUS message, which includes attributes such as IMSI, IMEI, and IPv4 or IPv6 addresses. The same SMP entry is then synchronized to the standby node.

If a RADIUS message contains IPv4 and IPv6 addresses, ACOS identifies two traffic maps associated with the IP addresses. As per the traffic maps, the active node is identified, and packets are redirected based on the following:

- If both traffic maps point to the same node, the packet is redirected to the active node.
- If both traffic maps point to different active nodes, the packet is redirected to both nodes. An SMP entry with both IPv4 and IPv6 is created on the active node and synchronized to the standby node. This ensures consistent session information across active and standby nodes, facilitating seamless failover transitions.

- If both traffic maps point to different active nodes but has the same standby nodes, the HA sync message is sent from both the nodes and only a single SMP entry is created.

**NOTE:** RADIUS message distribution is supported with traffic-slice configuration or any non-default Scaleout user group configuration.

To distribute the RADIUS traffic equally among all the nodes in the cluster, you must configure the loopback interface with the same IP address on all the nodes. The RADIUS messages must be sent to the loopback IP. You must configure the routing protocols such as OSPF or BGP between the router and the ACOS device so that the client-side router allows the RADIUS message to reach any node in the cluster.

The RADIUS message distribution in a cluster is applicable for both CGN and Firewall devices. Both L2 and L3 redirection are supported on the RADIUS message distribution.

When a node fails or a new node is added to the cluster, the following is observed:

- **Node Addition**: ACOS periodically checks whether the node is active or standby and synchronizes the entries to the corresponding nodes. When there is a new node added, it takes approximately 30 seconds for the traffic map to reflect the changes. If the new node is the active node for the subscriber IP, then the RADIUS entry is created in the new active node and is synchronized to the standby node. The entry on the old node is deleted.
- **Node Deletion or Failure**: Similar to the behavior described under node addition, when an active node fails or is deleted, the standby node becomes the active node. The RADIUS entries are then synchronized to the standby node for that subscriber.

Scaleout supports the following RADIUS accounting message types:

- Accounting Start
- Accounting Stop
- Accounting Interim Update
- Accounting On

Each message type has a default action configured, which can be overwritten in the radius server configuration.

| Message Type | Description | Default Actions |
|---|---|---|
| `start` | Initiates accounting for a new session. | `append-entry` |
| `stop` | Ends accounting for a session. | `delete-entry` |
| `interim-update` | Provides periodic updates on session activity. | `ignore` |
| `on` | Indicates that accounting is active or enabled. | `ignore` |

The following are the supported configurable actions for the RADIUS message. When a node receives a RADIUS message with the default or configured action, the specified action is performed and synced to the standby node.

| Configurable Actions | Description |
|---|---|
| `delete-entry` | Deletes an entry associated with the RADIUS message. |
| `replace-entry` | Replaces an existing entry with a new one. |
| `append-entry` | Updates attributes to the existing entry. |
| `delete-entry-using-attribute` | Allows the deletion of entries associated with a particular attribute. |

For detailed information about RADIUS accounting, see *IPv4-to-IPv6 Transition Solutions Guide*.

## Example Configuration for Distributing RADIUS Traffic

The following example configuration demonstrates how to configure the routing protocols, loopback interfaces, and RADIUS server configuration to ensure the RADIUS traffic is redistributed across all nodes:

1. Configure the routing protocols such as OSPF or BGP within the client network such that the client-side router allows the RADIUS message to reach any node in

the cluster.

```
router bgp 2
bgp router-id 50.50.50.1
maximum-paths 10
default-information originate
neighbor 50.50.50.2 remote-as 20
neighbor 50.50.50.3 remote-as 20
neighbor 50.50.50.4 remote-as 20
redistribute connected
redistribute static
```

2. Configure the same RADIUS server IP address as a loopback interface on all the nodes in the cluster.

```
interface loopback 1
ip address 70.70.70.70 255.255.255.255

router bgp 2
neighbor 50.50.50.1 remote-as 10
redistribute connected
redistribute static
```

3. Configure the RADIUS server parameters with RADIUS Client IP as 20.20.200.1 ip-list RADIUS.

```
ip-list RADIUS
20.20.200.1
secret 154698
system RADIUS server
  remote ip-list RADIUS

  secret a10
  attribute msisdn number 31

  attribute imei vendor 10415 number 20

  attribute imsi vendor 10415 number 1

  attribute custom2 Location number 32
```

```
attribute inside-ip number 8

attribute inside-ipv6-prefix prefix-length 96 number 97

accounting start replace-entry
accounting stop delete-entry-and-sessions

accounting interim-update replace-entry

accounting on delete-entries-using-attribute msisdn
```

It is important to exclude the loopback IP address from the IP routing table in order to ensure the radius message distribution works properly.

In the following example, as loopback 1 is used as the server IP address, you can exclude this IP from being used for routing by using the **exclude-interfaces** CLI command as shown below:

```
ACOS(config)# scaleout 64
ACOS(config-cluster:64)# device-context 1
ACOS(config-cluster:64)# local-device
ACOS(config-cluster:64-local-device)# priority 1
ACOS(config-cluster:64-local-device)# exclude-interfaces
ACOS(config-cluster:64-local-device-exclud...)# loopback 1
ACOS(config-cluster:64-local-device-exclud...)# exit
ACOS(config-cluster:64-local-device)# start-delay 30
ACOS(config-cluster:64-local-device)# disable
```

## Handling RADIUS Message in Multi-PU

In multi-PU, the RADIUS packet will land on an appropriate PU based on the client IP status (odd to PU1 and even to PU2). The packet will be redistributed to the appropriate PU based on the RADIUS message's IP addresses. When both IPv4 and IPv6 addresses are available in the RADIUS message, the redistribution will be applied with IPv4 and IPv6 addresses respectively.

## Limitations

Consider the following limitations:

- When RADIUS is configured in Scaleout mode, the packet must be sent to the loopback address only. If the packet is sent to a non-loopback address, it will be lost.
- During the node addition process, the RADIUS entry is synchronized from the old active node to the newly added node. An entry that is created after the synchronization is complete but before the new traffic map takes effect is not synchronized to the newly added node and is lost once the traffic map is updated.
- When a RADIUS entry from an old active node is sent to the new active node for that subscriber, the entry from the old active node is deleted. There is no backup taken for this deleted entry. In case of the synchronization packet being lost, the subscriber must log out and log back in to create a new RADIUS entry.
- When a RADIUS message is received on a non-active node for a subscriber, the message will be redirected to the correct active node. When the cluster is in L2 mode and a dedicated interface is not configured, the interface chosen on the active node (for redirecting the message) must be the one that has the same tag as the interface on which the message is received on the current node. If there is no tag, ACOS may not know where to redirect and drop the packet. Hence, you must tag the interfaces as follows:
  - For CGN, tag the interfaces with `ip nat inside` and `ip nat outside`.
  - For firewall, tag the interfaces with `ip client` and `ip server`.
- If a RADIUS message has an IPv6 address prefix, the prefix-length must be consistent across all RADIUS messages. In addition, it should be equal to or greater than the “system ipv6-prefix-length” configuration to ensure the message is redirected to the correct node.
- If a RADIUS message contains both IPv4 and IPv6 prefixes, the message is redirected based on the IPv6 prefix. The user quota must be configured based on the same IPv6 prefix. If the user quota is not configured with the same IPv6 prefix, the prefix configured in the system will be used to redirect the packets.

## Displaying and Clearing the RADIUS Sever Statistics

You can enter the following commands to display the RADIUS server statistic:

```
ACOS# show system radius server statistics
RADIUS packets dropped due to redirect failure (SO)                10
```

To clear the RADIUS server statistics, use the following command:

```
ACOS# clear system radius server statistics
```

## CGN Scaleout Limitations

Scaleout has the following limitations:

- CGN Scaleout supports LSN (NAT 44, NAT64), Fixed NAT (NAT44 and NAT64), and DS-Lite. The following CGN IPv6 technologies are not supported:
  - Static-NAT or range list
  - One-to-One NAT
  - Stateless technologies such as MAP-E, MAP-T, and Lightweight 4over6
- Port Reservation with Prefix Quota for DS-Lite is not supported.
- An explicit route for NAT Pool subnet with the outbound router as next hop is required for hairpinning.

**NOTE:** It is not recommended to configure unsupported technologies in Scaleout mode.

- CGN Scaleout supports only the default device group which contains all the devices configured within the cluster.
- Scaleout does not support the NAT pool configuration using GUI.
- Scaleout does not support SCTP and PCP.
- While integrating with AAA for obtaining subscriber identity via accounting messages, the messages are not replicated between nodes.
- The default number of User-Groups is 256. This value cannot be modified.
- There is no correlation between IPv4 and IPv6 addresses belonging to the same subscriber.

# Configure Scaleout for CGN

This section describes the steps to configure scaleout for CGN.

Figure 8 illustrates a sample basic CGN Scaleout topology.

Figure 8 : Sample CGN Scaleout Topology

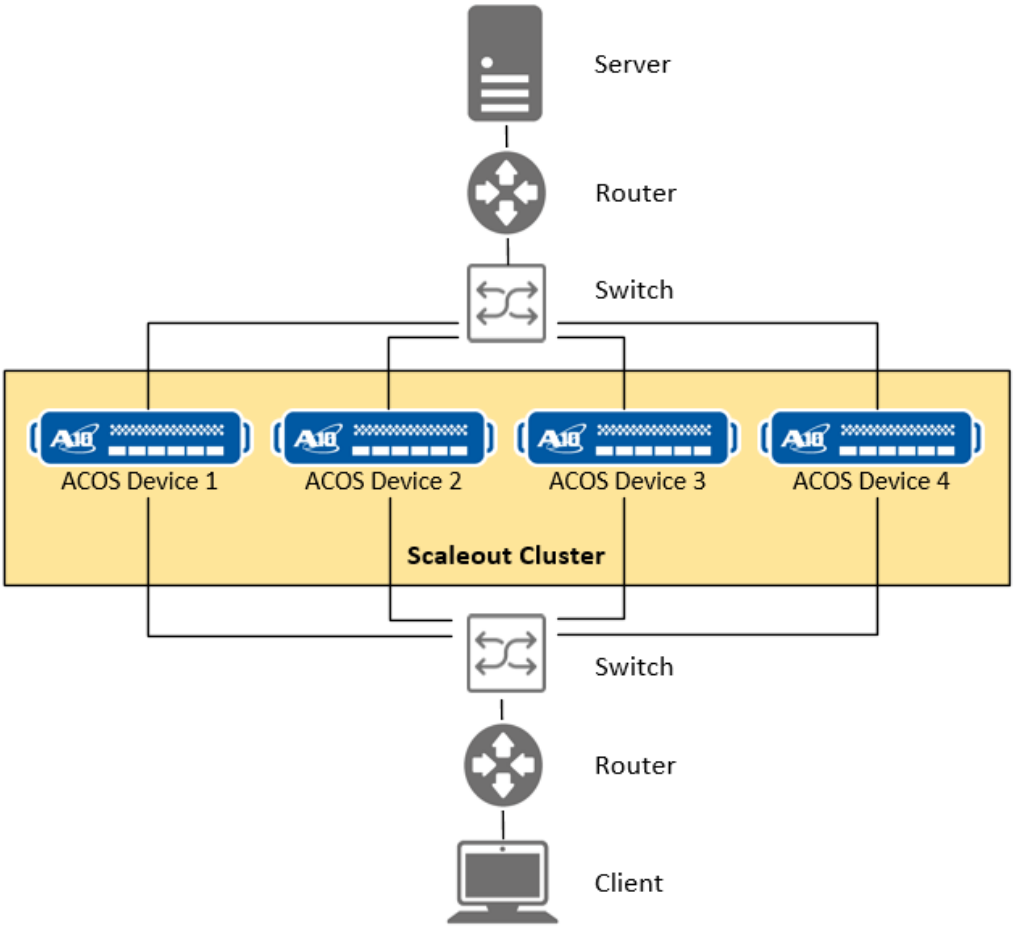


The steps to configure CGN:

1. Configure aVCS on each device. (See Configure aVCS on Each Device.)

   aVCS is not mandatory for Scaleout-related configuration. For Scaleout to function efficiently, the Scaleout-related configuration on all devices in the cluster must be applied and synchronized. To accomplish this, you may use aVCS to automatically synchronize the configurations on all devices. Alternatively, if you choose not to use aVCS, then you must manually replicate the configuration on all devices.

**NOTE:** aVCS is not mandatory for Scaleout-related configuration. When aVCS is used, configuration synchronization and cluster management will be automatically done in a cluster.

2. Set up the Scaleout configuration on the vMaster. (See Configure Clusters.)

   With aVCS configured and enabled, configuration changes on the vMaster are automatically synchronized to the Service Nodes.

3. Enable Scaleout. (See Enable Scaleout.)
4. Set up CGN configurations. (See Configuring CGN.)

## Configure aVCS on Each Device

This section describes how to configure aVCS on Cluster Node 1.

The following commands configure IPv4 access on the management interface on the new device:

```
ACOS(config)# hostname ACOS1
ACOS1-Active(config:1-device:1)# interface management
ACOS(config-if:management)# ip address 10.0.18.101 /24
ACOS(config-if:management)# ip default-gateway 10.0.18.1
```

The following commands enable the Ethernet interface, configure IPv4 access, and enable inside source NAT on the interface connected to the internal hosts:

```
ACOS(config)# interface ethernet 1
ACOS(config-if:ethernet:1)#  enable
ACOS(config-if:ethernet:1)#  ip address 20.20.20.101 /24
ACOS(config-if:ethernet:1)#  ip nat inside
```

The following commands enable the Ethernet interface, configure IPv4 access, and enables source NAT on the interface connected to the external hosts:

```
ACOS(config)# interface ethernet 2
ACOS(config-if:ethernet:2)#  enable
ACOS(config-if:ethernet:2)#  ip address 30.30.30.101 /24
ACOS(config-if:ethernet:2)#  ip nat outside
```

Specify a 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 1
ACOS(config-common)# set-id 8
ACOS-Active(config-common)# exit
```

The following command enables aVCS.

```
ACOS-1-Active(config)# vcs enable
```

The following commands configure the aVCS profile for the device:

```
ACOS-1-Active(config:1)# vcs device 1
ACOS-1-Active(config:1-device:1)# priority 200
ACOS-1-Active(config:1-device:1)# interfaces management
ACOS-1-Active(config:1-device:1)# enable
```

Repeat the same procedures to configure other Cluster Nodes.

## Configure CGN with LSN

This section describes how to configure CGN with LSN on vMaster for all Cluster Nodes.

In the following commands, the nat pools, cp1 and cp3 are associated with the service-config template traffic-slice-1 while cp2 is associated with the service-config template traffic-slice-2.

NAT pools configured for example, cp100, without service-config template configuration use the default number of user groups for that partition.

```
ACOS(config)# cgnv6 nat pool cp1 30.30.30.0 30.30.30.3 netmask /30
service-config template traffic-slice-1

ACOS(config)#cgnv6 nat pool cp3 30.30.30.7 30.30.30.10 netmask /30
service-config-template traffic-slice-1

ACOS(config)#cgnv6 nat pool cp2 200.0.0.5 200.0.0.5 netmask /32 service-
config-template traffic-slice-2

ACOS(config)#cgnv6 nat pool cp100 100.1.1.0 100.1.1.255 netmask /24
```

The following command binds the IPv4 NAT pool to the LID.

```
ACOS(config)# cgnv6 lsn-lid 1
ACOS(config-lsn lid)# source-nat-pool cp1

ACOS(config)# cgnv6 lsn-lid 2
ACOS(config-lsn lid)# source-nat-pool cp2

ACOS(config)# cgnv6 lsn-lid 3
ACOS(config-lsn lid)# source-nat-pool cp3

ACOS(config)# cgnv6 lsn-lid 100
ACOS(config-lsn lid)# source-nat-pool cp100

ACOS(config)# class-list lsn
ACOS(config-class list)# 10.1.1.0/24 lsn-lid 1
ACOS(config-class list)# 20.1.1.0/24 lsn-lid 2
ACOS(config-class list)# 10.2.2.0/24 lsn-lid 3
ACOS(config-class list)# 30.1.1.0/24 lsn-lid 100

ACOS(config)# cgnv6 source inside class-list lsn
```

## Configuration with lsn-rule-list

The lsn-rule-list can be used to apply different NAT pools for the same client based on the destination of traffic.

When traffic-slice is used with lsn-rule-list, the nat pool under lsn-lid and the pool under lsn-rule-list MUST be associated with the same service-config template.

In this example, pool cp1 and cp3 MUST be associated with the same service-config template.

```
ACOS(config)#cgnv6 lsn-rule-list r1
ACOS(config-lsn-rule-list)#ip 70.1.1.0/24
ACOS(config-lsn-rule-list-ip)#default action snat pool cp3
ACOS(config)#cgnv6 lsn-lid 1
ACOS(config-lsn-lid)#source-nat-pool cp1
ACOS(config-lsn-lid)#lsn-rule-list destination r1
```

# Deploy LSN with FW Using lsn-lid in Rule Action

In some cases, when LSN is deployed with Firewall, LSN class-list is not required if the Firewall rule action has lsn-lid configured.

But in case traffic-slice is used by NAT pools, LSN class-list MUST ALWAYS be configured, even if the Firewall rule action has lsn-lid configured.

```
ACOS(config)#rule-set slices
ACOS(config-rule-set slices)#rule slice-1
ACOS(config-rule-set slices-rule-slice-1)#source ipv4-address 10.1.1.0/24
ACOS(config-rule-set slices-rule-slice-1)#action-group
ACOS(config-rule-set slices-rule-slice-1)#permit cgnv6 lsn-lid 1
ACOS(config-rule-set slices)#rule slice-2
ACOS(config-rule-set slices-rule-slice-2)#source ipv4-address 20.1.1.0/24
ACOS(config-rule-set slices-rule-slice-2)#action-group
ACOS(config-rule-set slices-rule-slice-2)#permit cgnv6 lsn-lid 2

ACOS(config)#class-list lsn
ACOS(config-class list)#10.1.1.0/24 lsn-lid 1
ACOS(config-class list)#20.1.1.0/24 lsn-lid 2
```

If NAT pools use a traffic-slice for the same clients, then there must not be multiple Firewall rules that select different NAT pools based on the destination. If such rules exist, NAT pools must be associated with the same traffic-slice service-config template.

The following configuration displays what cannot be configured:

```
ACOS(config)#rule-set slices
ACOS(config-rule-set slices)#rule slice-1a
ACOS(config-rule-set slices-rule-slice-1a)#source ipv4-address 10.1.1.0/24
ACOS(config-rule-set slices-rule-slice-1a)#dest zone outside
ACOS(config-rule-set slices-rule-slice-1a)#action-group
ACOS(config-rule-set slices-rule-slice-1a)#permit cgnv6 lsn-lid 1
ACOS(config-rule-set slices)#rule slice-1b
ACOS(config-rule-set slices-rule-slice-1b)#source ipv4-address 10.1.1.0/24
ACOS(config-rule-set slices-rule-slice-1b)#dest zone dmz
ACOS(config-rule-set slices-rule-slice-1b)#action-group
ACOS(config-rule-set slices-rule-slice-1b)#permit cgnv6 lsn-lid 100
```

```
ACOS(config)#cgnv6 nat pool cp1 100.0.0.0 100.0.0.3 netmask /30 service-
config-template traffic-slice-1
ACOS(config)#cgnv6 nat pool cp100 100.1.1.0 100.1.1.255 netmask /24

ACOS(config)#cgnv6 lsn-lid 1
ACOS(config lsn-lid)#source-nat-pool cp1

ACOS(config)#cgnv6 lsn-lid 100
ACOS(config lsn-lid)#source-nat-pool cp100
```

If NAT pools use a traffic-slice for the same clients, NAT rule and transparent rule must not be configured.

The following configuration displays what cannot be configured:

```
ACOS(config)#rule-set slices
ACOS(config-rule-set slices)#rule slice-1a
ACOS(config-rule-set slices-rule-slice-1a)#source ipv4-address 10.1.1.0/24
ACOS(config-rule-set slices-rule-slice-1a)#dest zone outside
ACOS(config-rule-set slices-rule-slice-1a)#action-group
ACOS(config-rule-set slices-rule-slice-1a)#permit cgnv6 lsn-lid 1
ACOS(config-rule-set slices)#rule slice-1b
ACOS(config-rule-set slices-rule-slice-1b)#source ipv4-address 10.1.1.0/24
ACOS(config-rule-set slices-rule-slice-1b)#dest zone dmz
ACOS(config-rule-set slices-rule-slice-1b)#action-group
ACOS(config-rule-set slices-rule-slice-1b)#permit forward

ACOS(config)#cgnv6 nat pool cp1 100.0.0.0 100.0.0.3 netmask /30 service-
config-template traffic-slice-1

ACOS(config)#cgnv6 lsn-lid 1
ACOS(config lsn-lid)#source-nat-pool cp1
```

## Configure CGN with Fixed-NAT

This section describes how to add the service-config template to the Fixed-NAT configuration.

```
ACOS(config)#cgnv6 fixed-nat inside 10.1.1.0 10.1.1.255 netmask /24 nat
100.0.0.0 100.0.0.3 netmask /30 service-config-template traffic-slice-1
```

```
ACOS(config)#cgnv6 fixed-nat inside 20.1.1.0 20.1.1.255 netmask /24 nat
100.0.0.7 100.0.0.10 netmask /30 service-config-template traffic-slice-1

ACOS(config)#cgnv6 fixed-nat inside 10.2.2.0 10.2.2.255 netmask /24 nat
200.0.0.5 200.0.0.5 netmask /30 service-config-template traffic-slice-2
```

## Fixed-NAT Using Rule-Set Configuration

The following is the Fixed-NAT using the rule-set configuration.

```
ACOS(config)#rule-set slices
ACOS(config-rule-set slices)#rule slice-1
ACOS(config-rule-set slices-rule-slice-1)#source ipv4-address 10.1.1.0/24
ACOS(config-rule-set slices-rule-slice-1)#action-group
ACOS(config-rule-set slices-rule-slice-1)#permit cgnv6 fixed-nat
ACOS(config-rule-set slices)#rule slice-2
ACOS(config-rule-set slices-rule-slice-2)#source ipv4-address 20.1.1.0/24
ACOS(config-rule-set slices-rule-slice-2)#action-group
ACOS(config-rule-set slices-rule-slice-2)#permit cgnv6 fixed-nat

ACOS(config)#cgnv6 fixed-nat inside 10.1.1.0 10.1.1.255 netmask /24 nat
100.0.0.0 100.0.0.3 netmask /30 service-config-template traffic-slice-1

ACOS(config)#cgnv6 fixed-nat inside 20.1.1.0 20.1.1.255 netmask /24 nat
100.0.0.7 100.0.0.10 netmask /30 service-config-template traffic-slice-2
```

If Fixed-NAT uses the traffic-slice for the same clients, NAT rule and transparent rule must not be configured.

The following configuration displays what cannot be configured:

```
ACOS(config)#rule-set slices
ACOS(config-rule-set slices)#rule slice-1a
ACOS(config-rule-set slices-rule-slice-1a)#source ipv4-address 10.1.1.0/24
ACOS(config-rule-set slices-rule-slice-1a)#dest zone outside
ACOS(config-rule-set slices-rule-slice-1a)#action-group
ACOS(config-rule-set slices-rule-slice-1a)#permit cgnv6 fixed-nat
ACOS(config-rule-set slices)#rule slice-1b
ACOS(config-rule-set slices-rule-slice-1b)#source ipv4-address 10.1.1.0/24
ACOS(config-rule-set slices-rule-slice-1a)#dest zone dmz
ACOS(config-rule-set slices-rule-slice-1b)#action-group
```

```
ACOS(config-rule-set slices-rule-slice-1b)#permit cgnv6 forward

ACOS(config)#cgnv6 fixed-nat inside 10.1.1.0 10.1.1.255 netmask /24 nat
100.0.0.0 100.0.0.3 netmask /30 service-config-template traffic-slice-1
```

## Configuring IPv6 Prefix Length

The following command configures the IPv6 prefix length in the source IPv6 address:

```
ACOS(config)#system ipv6-prefix-length length
```

**NOTE:** The IPv6 prefix length can be configured only in the shared partition.

## Configure Clusters

This section describes how to configure clusters on vMaster for all Cluster Nodes and add the device-specific configuration to each device in the cluster.

The following commands configure device-specific and routing settings for an ACOS device in an aVCS environment.

```
ACOS(config)# scaleout 64
ACOS(config-cluster:64)# device-context 1
```

The following commands enter the Scaleout local device configuration mode to configure the local device in a cluster. After configuring the local device in a cluster, disable the local device in order to add devices gracefully to the cluster.

```
ACOS(config-cluster:64)# local-device
ACOS(config-cluster:64-local-device)# priority 1
ACOS(config-cluster:64-local-device)# start-delay 30
ACOS(config-cluster:64-local-device)# disable

ACOS(config-cluster:64)# device-context 2
ACOS(config-cluster:64)# local-device
ACOS(config-cluster:64-local-device)# priority 2
ACOS(config-cluster:64-local-device)# start-delay 30
ACOS(config-cluster:64-local-device)# disable

ACOS(config-cluster:64)# device-context 3
```

```
ACOS(config-cluster:64)# local-device
ACOS(config-cluster:64-local-device)# priority 3
ACOS(config-cluster:64-local-device)# start-delay 30
ACOS(config-cluster:64-local-device)# disable
```

# Add a Device Gracefully

The following example explains how to add a device to a cluster:

```
ACOS(config)# scaleout 64
ACOS(config-cluster:64)# device-context 1
ACOS(config-cluster:64)# local-device
ACOS(config-cluster:64-local-device)# priority 1
ACOS(config-cluster:64-local-device)# start-delay 30
ACOS(config-cluster:64-local-device)# disable
ACOS(config-cluster:64-local-device)# exit
```

## Remove a Device Gracefully

Prior to removing a device from a cluster, you must first disable the local device. Once the local-device is disabled in `show scaleout`, the traffic map is updated and the node disappears from the list.

To remove more than one node, you must remove the node one by one from a cluster.

The following example shows how to remove a device gracefully from a cluster:

```
ACOS(config-cluster:64)# local-device
ACOS(config-cluster:64-local-device)# priority 1
ACOS(config-cluster:64-local-device)# disable
```

**NOTE:** Prior to deleting a scaleout cluster, you must disable CGN scaleout by using the **`no scaleout apps enable`** command.

# Enable Scaleout

Scaleout can be enabled for CGN, Gi/SGi Firewall, and standalone Firewall.

### Prerequisites

- A scaleout cluster must be configured.
- No application configuration must be present prior to enabling scaleout.

Use the following command to enable Scaleout:

```
ACOS(config)# scaleout apps enable
```

**NOTE:** If you disable Scaleout using the **no scaleout apps enable** command, reload the device for the change to take effect.

## Configure Route Redistribution

To access the BGP router configuration level, use the following command at the global configuration level:

```
ACOS(config)# router bgp 2
```

To specify the networks to be advertised by the ACOS device's BGP routing process, use the following command:

```
ACOS(config-bgp:2)# network 50.50.50.102/32
```

To specify each of the ACOS device's neighbor (peer) BGP routers, use the following command:

```
ACOS(config-bgp:2)# neighbor 50.50.50.105 remote-as 2
```

To redistribute routes into BGP for reaching translated NAT addresses allocated from a pool, use the following command:

```
ACOS(config-bgp:2)# redistribute ip-nat
```

To redistribute routes into BGP for reaching translated NAT address allocated from a range list, use the following command:

```
ACOS(config-bgp:2)# redistribute ip-nat-list
```

## Configuration Example

The following is the configuration example for CGN on a Scaleout cluster.

```
ACOS(config)# cgnv6 nat pool cp1 30.30.30.1 30.30.30.225 netmask /24
ACOS(config)# system ipv6-prefix-length 72
ACOS(config)# cgnv6 nat64 user-quota-prefix-length 72
ACOS(config)# cgnv6 lsn-lid 1
ACOS(config-lsn-lid)# user-quota-prefix-length 96
ACOS(config-lsn-lid)# source-nat-pool cp1
ACOS(config-lsn-lid)# exit
ACOS(config)# class-list ipv6-users
ACOS(config-class list)# 3001::/64 lsn-lid 1
ACOS(config-class list)# exit
```

```
ACOS(config)# cgnv6 nat64 inside source class-list ipv6-users
```

```
ACOS(config)# interface ve 118
ACOS(config-if:ve118)# ipv6 address 2001:db8::2:15/96
ACOS(config-if:ve118)# ipv6 nat inside
ACOS(config-if:ve118)# exit
```

```
ACOS(config)# interface ve 119
ACOS(config-if:ve119)# ip address 30.30.30.1 255.255.255.0
ACOS(config-if:ve119)# ip nat outside
ACOS(config-if:ve119)# exit
```

```
ACOS(config)# scaleout 64
ACOS(config-cluster:64)# device-context 1
ACOS(config-cluster:64)# local-device
ACOS(config-cluster:64-local-device)# priority 1
ACOS(config-cluster:64-local-device)# start-delay 30
ACOS(config-cluster:64-local-device)# l2-redirect interface ethernet 5
vlan 60
ACOS(config-cluster:64-local-device)# session-sync
ACOS(config-cluster:64-local-device-session...)# interface ve 20
```

```
ACOS(config-cluster:64)# device-context 2
ACOS(config-cluster:64)# local-device
```

```
ACOS(config-cluster:64-local-device)# priority 2
ACOS(config-cluster:64-local-device)# start-delay 30
ACOS(config-cluster:64-local-device)# l2-redirect interface ethernet 5
vlan 2000
ACOS(config-cluster:64-local-device)# session-sync
ACOS(config-cluster:64-local-device-session...)# interface ve 20
ACOS(config-cluster:64-local-device-session...)# exit
```

```
Scaleout apps enable
```

# Configuring Hairpinning Scaleout CGN

ACOS supports hairpin communication between internal CGN clients in Scaleout cluster using their public NAT addresses. Scaleout architecture used in large-scale CGN deployments distributes NAT sessions across multiple nodes.

The following topics are covered:



## Overview Hairpin Solution

The hairpin effect occurs when traffic from an internal source is destined for an internal service but first exits through a network device, such as a router, before being redirected back inside. This is required when policies, security checks, or NAT rules require the traffic to pass through an external-facing interface.

In a Scaleout architecture, a hairpin solution refers to a technique that enables direct communication between two clients that reside behind the same external-facing IP but need to route traffic through a firewall or router. To maintain connectivity and hairpin communication between NATed clients, hairpin NAT, routing, and session handling must be appropriately configured.

## Key Considerations

For successful hairpin communication, the following factors must be considered:

- Traffic received from **`ip dmz`**, **`ip nat inside`** and **`ip client`** interfaces are treated the same way. They follow the same hashing logic and rely on the source IP address to determine how packets are forwarded to a node.
- For traffic coming from interfaces tagged as **`ip server`** or **`ip nat outside`**, it selects the node based on the destination IP only if the interface does not have any **`ip client`**, **`ip dmg`** or **`ip nat inside`** tag configured.
- For interfaces having multiple tags, such as, **`ip client`** and **`ip nat outside`**, the source IP is considered. In such cases, the system prioritizes the tags **`ip dmz`**, **`ip nat inside`** and **`ip client`**.
- The interfaces on the ACOS scaleout node must be properly tagged. Proper tagging ensures that traffic flows through the correct paths without misclassification or unintended blocking.
  - The FW client facing interface should be tagged with **`ip client`**
  - CGN-client facing interface should be tagged with **`ip nat inside`**
  - The internet facing interface should be tagged with **`ip nat outside`**

## Skip URPF Check

If the **`fw urpf strict`** command is enabled, hairpin traffic redirected by the router must be configured to skip the URPF check. Since some interfaces would be tagged, while some others might not, packets arriving from an outside interface will be dropped. To ensure that traffic received from an unexpected interface is also allowed, configuring the skip URPF check is necessary.

For details, see (skip-urpf-check command in the Command Line Interface Reference guide).

However, considering the security issues that skipping URPF checks might bring in, it is recommended to use either of the following solutions if URPF skip check is configured:

- Apply URPF checks on the router.
- Use a Separate VRF (Virtual Routing and Forwarding) for hairpin traffic.

# Deployment Example

ACOS supports CGN to CGN hairpin using router direction. For details on the CGN to CGN hairpinning solution, see the Hairpining section in the Carrier Grade NAT Guide.

The following deployment diagram illustrates the CGN-to-FW Hairpinning scenario across different nodes in a Scaleout, where CGN and FW sessions handle traffic that must be routed back within the same cluster. The configuration ensures traffic flows correctly through NAT and firewall processing without session mismatches or looping issues.

Figure 9 : CGN to FW Hairpinning Deployment

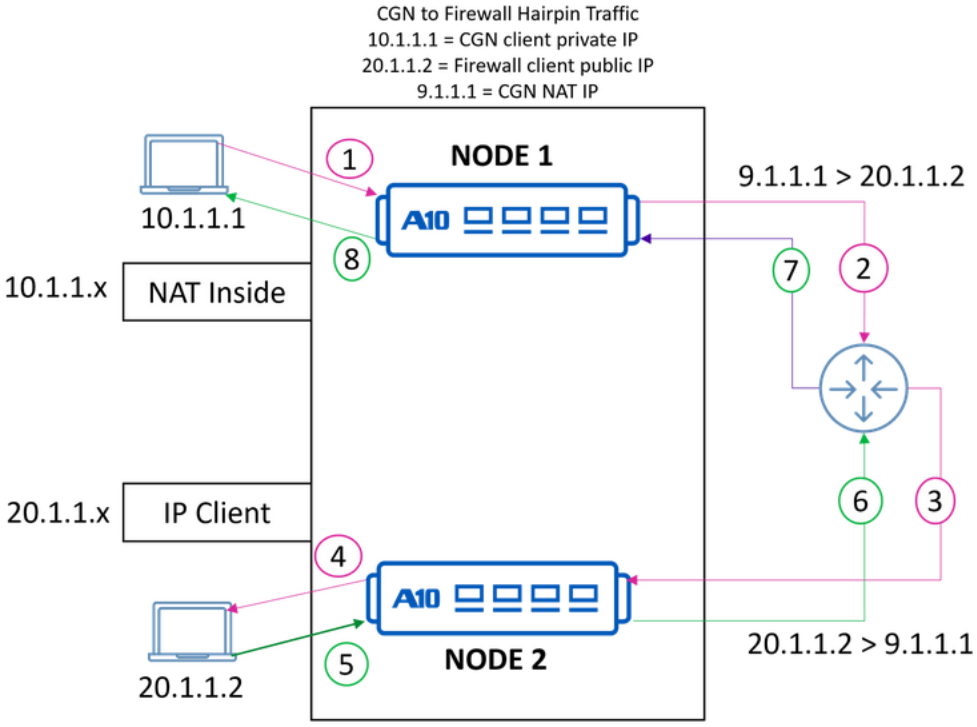


The CGN client is on Node 1, and the firewall client is on Node 2. The following is the high-level client 1 to client (CGN-to-FW) hairpining traffic flow:

1. The CGN client 1 (10.1.1.1) from the NAT inside network initiates a connection to the destination (20.1.1.2 FW client 2). The packet is sent to node 1.

2. The traffic is first processed by Node 1 (CGN processing), where NAT translation is applied. A CGN session is created and the packet is forwarded to the outside router based on the fw hairpin next-hop-follow configuration.
3. The router forwards the packet back to node 2 based on the route configured. Note that the router should have routes configured for redirecting the packet back to the Scaleout cluster.
4. Node 2 creates the FW session and forwards the packet back to client 2.
5. The destination FW client (20.1.1.2) sends a response back to node 2.
6. The packet matches the FW session created in step 4 on node 2 and node 2 forwards the packet to the outside router.
7. The outside router sends the response packet back to node 1.
8. The node 1 forwards the packet to client 1 on 10.1.1.1.

**NOTE:**

The network behavior differs based on whether:

- Zones are used (src NAT applied)
- Zones are not used (direct NAT translation without src NAT)

These conditions impact routing decisions and session persistence.

## CGN to FW Hairpin Configuration Example 1

The following CLI commands illustrate how to configure Scaleout CGN to FW Hairpin. No specific rule is needed to permit hairpin traffic between CGN and FW clients in this case. However, CGN and FW clients can access each other via EIF full-cone sessions and ALGs.

1. Configure Firewall Hairpin Next Hop

   The following commands configure the ACOS device to redirect the hairpin traffic back to the nodes. IP routes are also configured. It ensures that hairpin traffic is redirected to the correct next-hop within the 11.0.0.0/24 network. The **`next-hop-follow`** keyword allows dynamic updates of the next-hop route based on network changes.

```
ACOS(config)# fw hairpin next-hop-follow 11.0.0.0/24
ACOS(config)# ip route 11.0.0.0 /24 30.1.1.2
ACOS(config)# ip route 11.0.0.0 /24 30.1.1.5
```

2. Configure Endpoint-Independent Mapping (EIM) and Filtering.

```
ACOS(config)# cgnv6 lsn endpoint-independent-mapping tcp
ACOS(config-eim-tcp)# port 1 to 65535
ACOS(config)# cgnv6 lsn endpoint-independent-mapping udp
ACOS(config-eim-udp)# port 1 to 65535
ACOS(config)# cgnv6 lsn endpoint-independent-filtering tcp
ACOS(config-eif-tcp)# port 1 to 655353
```

3. Configure the firewall rule set with rules. The following command defines the rule set to allow traffic between internal zones.

```
ACOS(config)#rule-set rs1
```

4. Configure a firewall rule for the CGN zone under rule-set rs1.

```
ACOS(config-rule set:rs1)# rule cgn_zone
ACOS(config-rule set:rs1-rule:cgn_zone)# action permit cgnv6 lsn-lid 1
ACOS(config-rule set:rs1-rule:cgn_zone)# source ipv4-address
10.1.1.0/24
ACOS(config-rule set:rs1-rule:cgn_zone)# source zone cgn1
ACOS(config-rule set:rs1-rule:cgn_zone)# dest ipv4-address any
ACOS(config-rule set:rs1-rule:cgn_zone)# dest zone any
ACOS(config-rule set:rs1-rule:cgn_zone)# service any
ACOS(config-rule set:rs1-rule:cgn_zone)# exit
ACOS(config-rule set:rs1)# exit
```

These commands define the *cgn_zone* rule. It permits CGN IPv6 NAT (cgnv6) for source traffic originating from 10.1.1.0/24 within CGN zone (*cgn1*). Traffic is allowed to any destination.

5. Configure a firewall rule under the rule-set.

```
ACOS(config-rule set:rs1)# rule fw_zone
ACOS(config-rule set:rs1-rule:fw_zone)# action permit listen-on-port
ACOS(config-rule set:rs1-rule:fw_zone)# source ipv4-address 20.1.1.0/24
ACOS(config-rule set:rs1-rule:fw_zone)# source zone fw1
ACOS(config-rule set:rs1-rule:fw_zone)# dest ipv4-address any
ACOS(config-rule set:rs1-rule:fw_zone)# dest zone any
```

```
ACOS(config-rule set:rs1-rule:fw_zone)# service any
```

## CGN to FW Hairpin Configuration Example 2

The following CLI commands illustrate how to configure Scaleout CGN to FW Hairpin. A specific rule is added to permit hairpin traffic from CGN clients to FW clients. In this case, all CGN clients can access FW clients without EIF or ALGs. But for FW client to access CGN clients, CGN EIF full-cone session or ALGs are needed.

1. Configure Firewall Hairpin Next Hop.

   The following command configures ACOS device to redirect the hairpin traffic back to the nodes. It ensures that hairpin traffic is redirected to the correct next-hop within the 192.168.1.0/24 network. The **`next-hop-follow`** keyword allows dynamic updates of the next-hop route based on network changes.

   ```
   ACOS (config)# fw hairpin next-hop-follow 11.0.0.0/24
   ACOS (config)# ip route 11.0.0.0 /24 30.1.1.2
   ACOS (config)# ip route 11.0.0.0 /24 30.1.1.5
   ```

2. Configure rule set with rules.

   The following commands define the rule set to allow traffic between internal zones.

   ```
   ACOS(config)# rule-set rs1
   ```

3. Configure a firewall rule for the CGN zone under rule-set rs1.

   These commands define the *`cgn_zone`* rule. It permits CGN IPv6 NAT (cgnv6) for source traffic originating from 10.1.1.0/24 within CGN zone (cgn1). Traffic is allowed to any destination.

   ```
   ACOS(config-rule set:rs1)# rule cgn_zone
   ACOS(config-rule set:rs1-rule:cgn_zone)# action permit cgnv6 lsn-lid 1
   ACOS(config-rule set:rs1-rule:cgn_zone)# source ipv4-address
   10.1.1.0/24
   ACOS(config-rule set:rs1-rule:cgn_zone)# source zone cgn1
   ACOS(config-rule set:rs1-rule:cgn_zone)# dest ipv4-address any
   ACOS(config-rule set:rs1-rule:cgn_zone)# dest zone any
   ACOS(config-rule set:rs1-rule:cgn_zone)# service any
   ACOS(config-rule set:rs1-rule:cgn_zone)# exit
   ```

```
ACOS(config-rule set:rs1)# exit
```

4. Configure a firewall rule.

   These commands define a *fw_zone* rule that permits traffic from 20.1.1.0/24 within firewall zone (fw1). The traffic is allowed to any destination.

```
ACOS(config-rule set:rs1)# rule fw_zone
ACOS(config-rule set:rs1-rule:fw_zone)# action permit
ACOS(config-rule set:rs1-rule:fw_zone)# source ipv4-address 20.1.1.0/24
ACOS(config-rule set:rs1-rule:fw_zone)# source zone fw1
ACOS(config-rule set:rs1-rule:fw_zone)# dest ipv4-address any
ACOS(config-rule set:rs1-rule:fw_zone)# dest zone any
ACOS(config-rule set:rs1-rule:fw_zone)# service any
```

5. Configure a firewall rule for CGN-to-FW traffic.

   These commands define an explicit rule to permit CGN-to-FW traffic through an external router. The **skip-urpf-check** command disables Unicast Reverse Path Forwarding (URPF) checks to prevent unnecessary packet drops. This command must be configured only if **fw urpf strict** is configured.

   The source traffic comes from 9.1.1.0/24 within outside zone. The traffic is destined to fw1 zone without restriction on destination IP. CGN clients can thus access FW clients without requiring EIF full-cone NAT or ALGs.

```
ACOS(config-rule set:rs1)# rule cgn_to_fw_via_router
ACOS(config-rule set:rs1-rule:cgn_to_fw_v...)# action-group
ACOS(config-rule set:rs1-rule:cgn_to_fw_v...)# permit skip-urpf-check
ACOS(config-rule set:rs1-rule:cgn_to_fw_v...)# source ipv4-address
9.1.1.0/24
ACOS(config-rule set:rs1-rule:cgn_to_fw_v...)# source zone outside
ACOS(config-rule set:rs1-rule:cgn_to_fw_v...)# dest zone fw1
ACOS(config-rule set:rs1-rule:cgn_to_fw_v...)# dest ipv4-address any
```

# Scaleout for Gi/SGi Firewall and Standalone Firewall

This chapter describes how Scaleout is configured in Gi/SGi-Firewall and Standalone Firewall deployments.

The following topics are covered:



## Traffic Distribution

The following topics are covered:



### Traffic Distribution in Gi/SGi-Firewall Deployment

In the Gi/SGi-Firewall Scaleout deployment, the inside users and NAT IP addresses are hashed and mapped to the user groups similar to CGN scaleout.

For more information about the following topics, see Overview:

- Traffic Distribution
- Scaleout Mapping
- Scaleout Redirection

## Traffic Distribution in Standalone Firewall Deployment

Traffic from upstream devices is distributed across the nodes based on the user group associated with the packet. Each node is responsible for processing packets belonging to a subset of user groups. In a network, all the traffic coming from the same subscriber may have different IPv6 addresses. However, ACOS expects the traffic coming from a subscriber to have the same prefix (of length IPv6 prefix length) in the source IPv6 address and accordingly identifies the subscriber associated with the packet based on the configured IPv6 prefix length. All the packets with the same prefix (of length IPv6 prefix length) in the source IPv6 address will map to the same user group and hence be processed by the same node in the Scaleout cluster.

The `ipv6-prefix-length` command is a system-level attribute that indicates the length of the source and destination IPv6 prefix. This prefix is used to determine the user group for processing a packet.

The IPv6 prefix length can be configured only in the shared partition but is applicable to all L3V partitions, LIDs, and Class-lists. The default value of this attribute is set to 128 on all devices.

Use the following command to configure the IPv6 prefix length:

```
ACOS(config)#system ipv6-prefix-length length
```

You can select a value between 16 and 128 for length. By default, the IPv6 prefix length is 128.

**NOTE:** The IPv6 prefix length must be modified only when the Scaleout cluster is not in use. Else, unpredictable disruptions may occur to the traffic.

## Distributed Forwarding

When an ACOS device receives packets on a node that is not active for the IP source, it redirects the packets to the active node where a session is created, packets in this session processed, and sent to the destination. ACOS provides an ability through Distributed forwarding to reduce the number of redirection of packets in the Firewall Scaleout to improve packet handling performance.

When the packet count in the receiving node exceeds the given threshold, a shadow session is created in the receiving node. You can configure the packet threshold value to offload sessions with UDP or TCP protocol. The threshold value can be 4 – 63. By default, the value is set to 5.

You can configure the directions (uplink or downlink or both (by default)) to offload the sessions. The shadow sessions are mimics of the sessions on the Active node. After the shadow session creation, the packets that are received on the shadow node are processed locally and sent to the destination based on the routing table. This minimizes the redirection of packets.

The default idle timeout for shadow sessions is 120 sec. After 120 seconds if the node does not receive any packet, the shadow session is deleted.

**NOTE:** Shadow sessions are not created for short-lived sessions and are measured by the number of packets in the session.

The following command enables the Scaleout distributed-forwarding for the Firewall:

```
ACOS(config)# scaleout distributed-forwarding fw enable
```

The following command disables the Scaleout distributed-forwarding for the Firewall:

```
ACOS(config)# scaleout distributed-forwarding fw disable
```

The following command configures the session offload in the uplink direction:

```
ACOS(config)# scaleout distributed-forwarding fw session-offload-direction
uplink
```

The following command configures the session offload in the downlink direction:

```
ACOS(config)# scaleout distributed-forwarding fw session-offload-direction
downlink
```

The following command configures the session offload in both directions:

```
ACOS(config)# scaleout distributed-forwarding fw session-offload-direction
both
```

The following command configures the packet threshold value to offload the udp or tcp sessions:

```
ACOS(config)# scaleout distributed-forwarding fw threshold <4-63> <udp |
tcp>
```

The following command disables the packet threshold value to offload the udp or tcp sessions:

```
ACOS(config)# no scaleout distributed-forwarding fw threshold <4-63> <udp
| tcp>
```

Consider the following points to use the `scaleout distributed-forwarding fw` command:

- Scaleout app must be enabled.
- All the nodes must have the same distributed forwarding configuration.

**NOTE:** Reboot is not required for the `distributed forwarding` command to work.

## Configuration Example

The following configuration example enables the scaleout distributed forwarding for the Firewall with the session offload in the uplink direction and the threshold value set to 6.

```
ACOS(config)# scaleout distributed-forwarding fw enable
ACOS(config)# scaleout distributed-forwarding fw session-offload-direction
uplink
ACOS(config)# scaleout distributed-forwarding fw threshold 6
```

## Limitations

Consider the following limitations:

- Distributed Forwarding is applicable for both TCP and UDP sessions and does not apply to IP-in-IP, ICMP, SCTP, GTP, or UDP-DNS protocols.
- In the L3 Scaleout mode, the Distributed Forwarding does not support fragmented packets (for both Outer IP and Inner IP).
- TCP Window Check (currently enabled by default) is not supported along with Distributed Forwarding.
- FW Rate Limiting is not supported with Distributed Forwarding.

# Configure Scaleout for Gi/SGi Firewall

This section describes the steps to configure scaleout for Gi/SGi Firewall.

Figure 10 illustrates a sample Gi/SGi Firewall Scaleout topology.

Figure 10 : Sample Gi/SGi Scaleout Topology

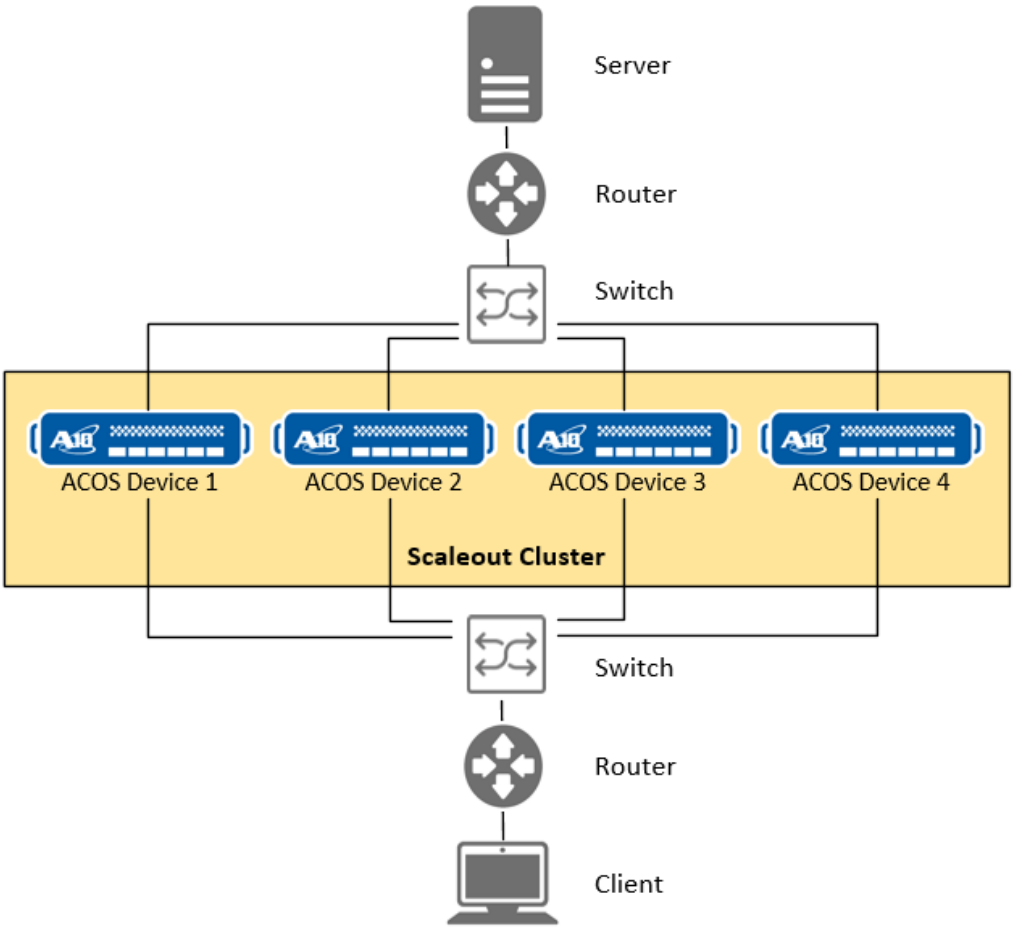


The steps to configure Gi/SGi Firewall:

1. Configure aVCS on each device. (See Configure aVCS on Each Device.)

   aVCS is not mandatory for Scaleout-related configuration. For Scaleout to function efficiently, the Scaleout-related configuration on all devices in the

cluster must be applied and synchronized. To accomplish this, you may use aVCS to automatically synchronize the configurations on all devices. Alternatively, if you choose not to use aVCS, then you must manually replicate the configuration on all devices.

**NOTE:** aVCS is not mandatory for Scaleout-related configuration. When aVCS is used, configuration synchronization and cluster management will be automatically done in a cluster.

2. Set up the Scaleout configuration on the vMaster. (See Configure Clusters.)

   With aVCS configured and enabled, configuration changes on the vMaster are automatically synchronized to the Service Nodes.

3. Enable Scaleout. (See Enable Scaleout.)
4. Set up Gi/SGi Firewall or standalone firewall configuration. (See Configure Gi/SGi Firewall in CGN Deployment or Configure Standalone Firewall.)

## Configuring IPv6 Prefix Length

The following command configures the IPv6 prefix length in the source IPv6 address:

```
ACOS(config)#system ipv6-prefix-length length
```

**NOTE:** The IPv6 prefix length can be configured only in the shared partition.

## Configure Standalone Firewall

Perform the following:

1. Configure the interfaces specified as the client and the server.

```
ACOS1(config)# interface ve 81
ACOS1(config-if:ve:81)# ip address 10.1.1.47 255.255.255.0
ACOS1(config-if:ve:81)# ip client
ACOS1(config-if:ve:81)# ipv6 address 1000::47/64
ACOS1(config-if:ve:81)# exit
```

```
ACOS1(config)# interface ve 83
ACOS1(config-if:ve:83)# ip address 30.1.1.47 255.255.255.0
ACOS1(config-if:ve:83)# ip server
ACOS1(config-if:ve:83)# ipv6 address 3000::47/64
ACOS1(config-if:ve:83)# exit
```

**NOTE:** `ip client` and `ip server` must be configured on different interfaces.

2. Configure a firewall rule-set. Rules should contain the match criteria and associated action.

```
ACOS(config)# rule-set rule2
ACOS(config-rule set: rule2)# rule r1
ACOS(config-rule set: rule2-rule:r1)# action permit
ACOS(config-rule set: rule2-rule:r1)# source ipv4-address any
ACOS(config-rule set: rule2-rule:r1)# source zone any
ACOS(config-rule set: rule2-rule:r1)# dest ipv4-address any
ACOS(config-rule set: rule2-rule:r1)# dest zone any
ACOS(config-rule set: rule2-rule:r1)# service any
ACOS(config-rule set: rule2-rule:r1)# exit

ACOS(config-rule set: rule2)# rule r2
ACOS(config-rule set: rule2-rule:r2)# action permit
ACOS(config-rule set: rule2-rule:r2)# ip-version v6
ACOS(config-rule set: rule2-rule:r2)# source ipv6-address any
ACOS(config-rule set: rule2-rule:r2)# source zone any
ACOS(config-rule set: rule2-rule:r2)# dest ipv6-address any
ACOS(config-rule set: rule2-rule:r2)# dest zone any
ACOS(config-rule set: rule2-rule:r2)# service any
ACOS(config-rule set: rule2-rule:r2)# exit
```

3. Configure a firewall logging server that specifies the IPv4 or IPv6 address, or hostname for the logging server, enables health monitoring of the server, and specifies the TCP or UDP port on which the server listens for traffic.

```
ACOS(config)# fw server s1 30.1.1.45
ACOS(config-real server)# health-check-disable
ACOS(config-real server)# port 514 udp
ACOS(config-real server-node port)# health-check-disable
```

```
ACOS(config-real server-node port)# port 514 tcp
ACOS(config-real server-node port)# health-check-disable
ACOS(config-real server-node port)# exit
```

4. Configure a service group for the firewall logging server and adds the external log server and port to the service group.

```
ACOS(config)# fw service-group fw_udp udp
ACOS(config-fw svc group)# member s1 514
ACOS(config-fw svc group)# exit
```

5. Configure a firewall logging template.

```
ACOS(config)# fw template logging log1
ACOS(config-logging)# log http-requests url
ACOS(config-logging)# rule http-requests dest-port 80
ACOS(config-logging)# service-group fw_udp
ACOS(config-logging)# exit
```

6. Bind a firewall logging template to the firewall.

```
ACOS(config)# fw logging log1
ACOS(config)# exit
```

7. Activate the firewall function using the specified rule-set.

```
ACOS(config)# fw active-rule-set rule2
```

For detailed information about Gi/SGi Firewall and standalone firewall configuration, see Scaleout for Gi/SGi Firewall and Standalone Firewall.

**NOTE:** A Firewall Scaleout cluster requires a minimum of 2 devices to be operational, up to a maximum of 8 devices. When a majority of the devices in the cluster is down, then all service nodes are removed from the cluster and the service goes down.

## Configure Gi/SGi Firewall in CGN Deployment

The following command configures a named set of IP addresses for use by CGN or LSN:

```
ACOS(config)# cgnv6 nat pool p1 9.9.9.0 9.9.9.255 netmask/24
```

The following commands configure a LID for NAT64 and add the pool to it:

```
ACOS(config)# cgnv6 lsn-lid 1
ACOS(config-lsn-lid)# source-nat-pool p1
ACOS(config-lsn-lid)# exit
```

The following command enables Fixed NAT:

```
ACOS(config)# cgnv6 fixed-nat inside 3201::100 3201::4ff netmask 96 nat
9.9.19.0 9.9.19.255 netmask /24
```

The following commands configure a firewall rule-set that contains a set of rules. In this example, rule 1 specifies that any packets matching this rule must be handled by LSN configurations, whereas packets matching rule 2 must be handled by Fixed NAT configurations. Packets matching rule 3 are permitted, and no CGN configurations are applied to them.

```
ACOS(config)# rule-set firewall
ACOS(config-rule set:firewall)# rule 1
ACOS(config-rule set:firewall-rule:1)# action permit cgnv6 lsn-lid 1
ACOS(config-rule set:firewall-rule:1)# dest ipv4-address any
ACOS(config-rule set:firewall-rule:1)# dest zone outside
ACOS(config-rule set:firewall-rule:1)# exit

ACOS(config)# rule-set firewall
ACOS(config-rule set:firewall)# rule 2
ACOS(config-rule set:firewall-rule:2)# action permit cgnv6 fixed-nat
ACOS(config-rule set:firewall-rule:2)# source ipv4-address 3201::172/128
ACOS(config-rule set:firewall-rule:2)# source zone inside
ACOS(config-rule set:firewall-rule:2)# exit
```

## Configure Route Redistribution

For more information, see Configure Route Redistribution.

# Configuring Hairpinning in Scaleout Firewall

ACOS supports hairpin communication between internal Firewall clients in Scaleout cluster. A router is used to redirect hairpin traffic back to the cluster to ensure that it reaches the correct node within the firewall. Interface tagging is used to maintain session consistency.

For details on hairpin traffic and solution offered, as well as for recommendations for interface configuration, see Overview Hairpin Solution and Key Considerations sections.

The following topics are covered:



# Deployment Example

The following deployment diagram illustrates the FW-to-FW Hairpinning scenario across different nodes in a Scaleout deployment, where two FW sessions handle traffic that must be routed back within the same cluster —one for client1 and one for client2.

Figure 11 : FW-to-FW Hairpinning Deployment

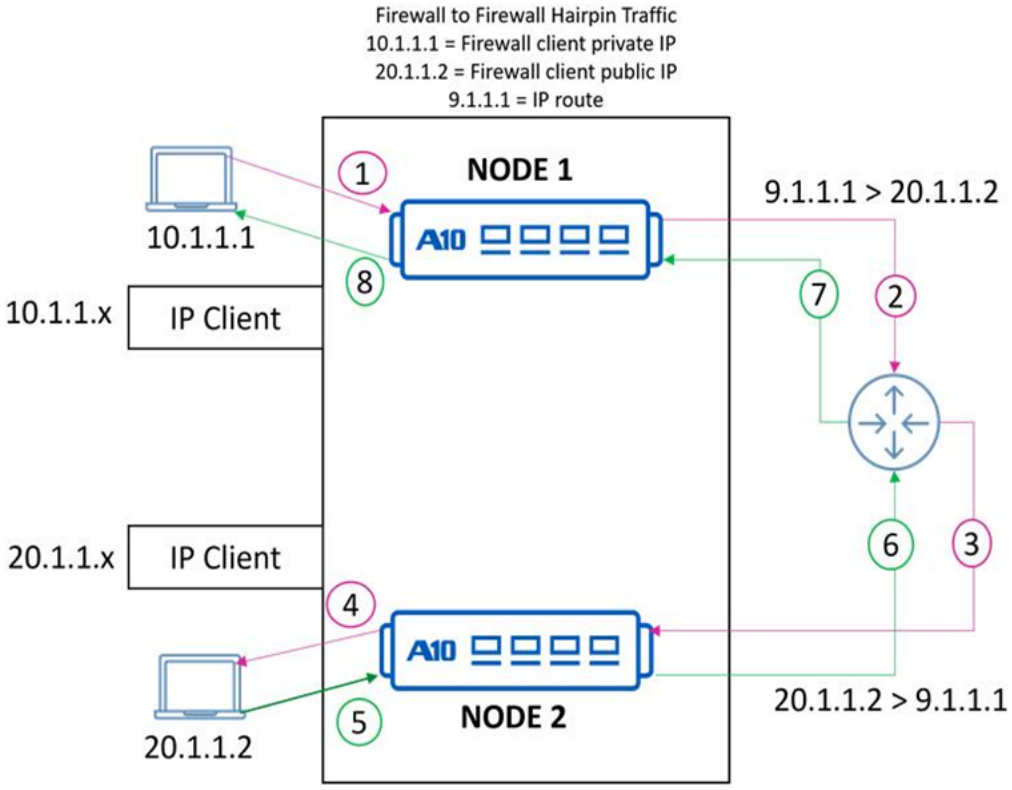


The Hairpin traffic flows from FW clients from zone 1 to FW clients zone 2 without full-cone session or ALGs. However, FW client zone 2 to FW clients zone 1 require fw full-cone session or ALGs.

## FW to FW Hairpin Configuration Example 1

The following CLI commands illustrate how to configure Scaleout Firewall Hairpin traffic from FW clients from zone 1 to FW clients zone 2. There is no need to configure a specific rule to permit hairpin traffic between FW and FW clients. However, FW clients can access each other via EIF full-cone sessions and ALGs.

1. Configure Firewall Hairpin Next Hop.

   These commands configure the ACOS device to redirect the hairpin traffic back to the firewall cluster. It ensures that hairpin traffic is redirected to the correct next-hop within the 11.0.0.0/24 network. The **next-hop-follow** keyword allows dynamic updates of the next-hop route based on network changes.

```
ACOS(config)# fw hairpin next-hop-follow 11.0.0.0/24
ACOS(config)# ip route 11.0.0.0 /24 30.1.1.2
ACOS(config)# ip route 11.0.0.0 /24 30.1.1.5
```

2. Configure rule set with rules.

   This command defines a new rule set named rs1, which contains firewall rules to control traffic between zones.

```
ACOS(config)#rule-set rs1
```

3. Configure a firewall rule.

   These commands configure a rule to permit all traffic from 10.1.1.0/24 (inside fw1 zone) to any destination. There are no restrictions on the destination IP, zone, or service (ports/protocols).

```
ACOS(config-rule set:rs1)# rule fw_zone1
ACOS(config-rule set:rs1-rule:fw_zone1)# action permit
ACOS(config-rule set:rs1-rule:fw_zone1)# action permit listen-on-port
ACOS(config-rule set:rs1-rule:fw_zone1)# source ipv4-address
10.1.1.0/24
ACOS(config-rule set:rs1-rule:fw_zone1)# source zone fw1
ACOS(config-rule set:rs1-rule:fw_zone1)# dest ipv4-address any
ACOS(config-rule set:rs1-rule:fw_zone1)# dest zone any
ACOS(config-rule set:rs1-rule:fw_zone1)# service any
```

4. Configure a firewall rule.

   These commands configure a rule to permit all traffic from 20.1.1.0/24 (inside fw2 zone) to any destination. Similar to fw_zone1, there are no restrictions on destination IP, zone, or service.

```
ACOS(config-rule set:rs1)#rule fw_zone2
ACOS(config-rule set:rs1-rule:fw_zone2)#action permit
ACOS(config-rule set:rs1-rule:fw_zone2)#action permit listen-on-port
```

```
ACOS(config-rule set:rs1-rule:fw_zone2)#source ipv4-address 20.1.1.0/24
ACOS(config-rule set:rs1-rule:fw_zone2)#source zone fw2
ACOS(config-rule set:rs1-rule:fw_zone2)#dest ipv4-address any
ACOS(config-rule set:rs1-rule:fw_zone2)#dest zone any
ACOS(config-rule set:rs1-rule:fw_zone2)#service any
```

## FW to FW Hairpin Configuration Example 2

The following CLI commands illustrate how to configure Scaleout Firewall Hairpin traffic from FW clients from zone 1 to FW clients zone 2. A specific rule is added to permit hairpin traffic from FW zone1 clients to FW zone2 clients. As a result, all the traffic initiated from FW zone 1 is allowed to FW zone 2. However, all FW zone 2 clients can only access FW zone 1 clients using EIF and ALGs.

1. Configure Firewall Hairpin Next Hop.

   These commands configure the ACOS device to redirect the hairpin traffic back to the firewall cluster. It ensures that hairpin traffic is redirected to the correct next-hop within the 11.0.0.0/24 network. The **`next-hop-follow`** keyword allows dynamic updates of the next-hop route based on network changes.

   ```
   ACOS (config)# fw hairpin next-hop-follow 11.0.0.0/24
   ACOS (config)# ip route 11.0.0.0 /24 30.1.1.2
   ACOS (config)# ip route 11.0.0.0 /24 30.1.1.5
   ```

2. Configure rule set with rules.

   This command defines a new rule set named rs1, which contains firewall rules to control traffic between zones.

   ```
   ACOS(config)# rule-set rs1
   ```

3. Configure a firewall rule, fw_zone1.

   These commands configure a rule to permit all traffic from 10.1.1.0/24 (inside fw1 zone) to any destination. There are no restrictions on the destination IP, zone, or service (ports/protocols).

   ```
   ACOS(config-rule set:rs1)# rule fw_zone1
   ACOS(config-rule set:rs1-rule:fw_zone1)# action permit
   ACOS(config-rule set:rs1-rule:fw_zone1)# source ipv4-address
   10.1.1.0/24
   ```

```
ACOS(config-rule set:rs1-rule:fw_zone1)# source zone fw1
ACOS(config-rule set:rs1-rule:fw_zone1)# dest ipv4-address any
ACOS(config-rule set:rs1-rule:fw_zone1)# dest zone any
ACOS(config-rule set:rs1-rule:fw_zone1)# service any
```

4. Configure a firewall rule, fw_zone2. These commands configure a rule to permit all traffic from 20.1.1.0/24 (inside fw2 zone) to any destination. Similar to fw_ zone1, there are no restrictions on destination IP, zone, or service.

```
ACOS(config-rule set:rs1)#rule fw_zone2
ACOS(config-rule set:rs1-rule:fw_zone2)#action permit
ACOS(config-rule set:rs1-rule:fw_zone2)#source ipv4-address 20.1.1.0/24
ACOS(config-rule set:rs1-rule:fw_zone2)#source zone fw2
ACOS(config-rule set:rs1-rule:fw_zone2)#dest ipv4-address any
ACOS(config-rule set:rs1-rule:fw_zone2)#dest zone any
ACOS(config-rule set:rs1-rule:fw_zone2)#service any
```

5. Configure a firewall rule for FW zone 1 to FW zone 2. These commands configure a rule to allow traffic from fw_zone1 to fw_zone2 through an external router. The **`permit skip-urpf-check`** command permits the traffic and bypasses URPF checks, ensuring that traffic flows correctly. This command must be configured only if **`fw urpf strict`** is configured.

   The source IP indicates that the traffic originates from 10.1.1.0/24. The outside source zone enables that traffic can traverse through an external network. The traffic is destined for fw2.

```
ACOS(config-rule set:rs1)# rule fw_zone1_to_zone2_via_router
ACOS(config-rule set:rs1-rule:fw_zone1_to...)# action-group
ACOS(config-rule set:rs1-rule:fw_zone1_to...)# permit skip-urpf-check
ACOS(config-rule set:rs1-rule:fw_zone1_to...)# source ipv4-address
10.1.1.0/24
ACOS(config-rule set:rs1-rule:fw_zone1_to...)# source zone outside
ACOS(config-rule set:rs1-rule:fw_zone1_to...)# dest zone fw2
ACOS(config-rule set:rs1-rule:fw_zone1_to...)# dest ipv4-address any
```

## FW to FW Hairpin Configuration Example 3

The following CLI commands illustrate how to configure Scaleout Firewall Hairpin traffic from FW clients from zone 1 to FW clients zone 2. Here, two specific rules are

added to permit hairpin traffic from FW zone 1 clients to FW zone 2 clients and from FW zone 2 to FW zone 1 clients. Clients in FW zone 1 and zone 2 can access each other without the need of full-cone session or ALGs.

1. Configure Firewall Hairpin Next Hop. These commands configure the ACOS device to redirect the hairpin traffic back to the firewall cluster. It ensures that hairpin traffic is redirected to the correct next-hop within the 11.0.0.0/24 network. The **next-hop-follow** keyword allows dynamic updates of the next-hop route based on network changes.

```
ACOS (config)# fw hairpin next-hop-follow 11.0.0.0/24
ACOS (config)# ip route 11.0.0.0 /24 30.1.1.2
ACOS (config)# ip route 11.0.0.0 /24 30.1.1.5
```

2. Configure rule set with rules.

   This command defines a new rule set named rs1, which contains firewall rules to control traffic between zones.

```
ACOS(config)#rule-set rs1
```

3. Configure a firewall rule, fw_zone1.

   These commands configure a rule to permit all traffic from 10.1.1.0/24 (inside fw1 zone) to any destination. There are no restrictions on the destination IP, zone, or service (ports/protocols).

```
ACOS(config-rule set:rs1)# rule fw_zone1
ACOS(config-rule set:rs1-rule:fw_zone1)# action permit
ACOS(config-rule set:rs1-rule:fw_zone1)# source ipv4-address
10.1.1.0/24
ACOS(config-rule set:rs1-rule:fw_zone1)# source zone fw1
ACOS(config-rule set:rs1-rule:fw_zone1)# dest ipv4-address any
ACOS(config-rule set:rs1-rule:fw_zone1)# dest zone any
ACOS(config-rule set:rs1-rule:fw_zone1)# service any
```

4. Configure a firewall rule, fw_zone2. These commands configure a rule to permit all traffic from 20.1.1.0/24 (inside fw2 zone) to any destination. Similar to fw_ zone1, there are no restrictions on destination IP, zone, or service.

```
ACOS(config-rule set:rs1)#rule fw_zone2
ACOS(config-rule set:rs1-rule:fw_zone2)#action permit
ACOS(config-rule set:rs1-rule:fw_zone2)#source ipv4-address 20.1.1.0/24
```

```
ACOS(config-rule set:rs1-rule:fw_zone2)#source zone fw2
ACOS(config-rule set:rs1-rule:fw_zone2)#dest ipv4-address any
ACOS(config-rule set:rs1-rule:fw_zone2)#dest zone any
ACOS(config-rule set:rs1-rule:fw_zone2)#service any
```

5. Configure a firewall rule for FW zone 1 to FW zone 2. These commands configure a rule to allow traffic from fw_zone1 to fw_zone2 through an external router. The **`permit skip-urpf-check`** command permits the traffic and bypasses URPF checks, ensuring that traffic flows correctly. This command must be configured only if **`fw urpf strict`** is configured.

   The source IP indicates that the traffic originates from 10.1.1.0/24. The outside source zone enables that traffic can traverse through an external network. The traffic is destined for fw2.

```
ACOS(config-rule set:rs1)#rule fw_zone1_to_zone2_via_router
ACOS(config-rule set:rs1-rule:fw_zone1_to...)#action-group
ACOS(config-rule set:rs1-rule:fw_zone1_to...)#permit skip-urpf-check
ACOS(config-rule set:rs1-rule:fw_zone1_to...)#source ipv4-address
10.1.1.0/24
ACOS(config-rule set:rs1-rule:fw_zone1_to...)#source zone outside
ACOS(config-rule set:rs1-rule:fw_zone1_to...)#dest zone fw2
ACOS(config-rule set:rs1-rule:fw_zone1_to...)#dest ipv4-address any
```

6. Configure a firewall rule for FW zone 2 to FW zone 1. These commands configure a rule to allow traffic from fw_zone2 to fw_zone1 through an external router. The permit skip-urpf-check command permits the traffic and bypasses URPF checks, ensuring that traffic flows correctly.

   The source IP indicates that the traffic originates from 20.1.1.0/24. The outside source zone enables that traffic can traverse through an external network. The traffic is destined for fw1.

```
ACOS(config-rule set:rs1)# rule fw_zone2_to_zone1_via_router
ACOS(config-rule set:rs1-rule:fw_zone2_to...)# action-group
ACOS(config-rule set:rs1-rule:fw_zone2_to...)# permit skip-urpf-check
ACOS(config-rule set:rs1-rule:fw_zone2_to...)# source ipv4-address
20.1.1.0/24
ACOS(config-rule set:rs1-rule:fw_zone2_to...)# source zone outside
ACOS(config-rule set:rs1-rule:fw_zone2_to...)# dest zone fw1
ACOS(config-rule set:rs1-rule:fw_zone2_to...)# dest ipv4-address any
```

# Firewall Scaleout Limitations

The Firewall Scaleout has the following limitations:

- The features that depend on the physical connectivity of the device are not supported. For example, a zone containing physical interfaces is not supported whereas a zone containing virtual interfaces is supported.
- The **`fw respond-to-user-mac`** command is not supported.
- The CGN traffic direction and the firewall traffic direction should be the same. The IP server cannot be configured under the interface with IP NAT inside configuration.

# Upgrading Scaleout Cluster from ACOS 5.2.1-Px to ACOS 6.0.x and Later Releases

This chapter describes a sequence of steps to migrate a 3-device Scaleout cluster from ACOS 5.2.1-Px to an ACOS 6.0.x-based release. The procedure may be duly adapted to a larger or smaller cluster.

There are configuration differences between the Scaleout configuration on 5.2.1-Px and 6.0.x. A 5.2.1-Px cluster cannot communicate or interoperate with a 6.0.x cluster. The steps below detail a method of removing devices from a 5.2.1-Px cluster and adding to a new 6.0.x cluster while minimizing traffic loss.

The following topics are covered:

# Summary of Steps

The summary of the migration procedure is as follows:

1. Disable and remove device 3 from the 5.2.1-Px cluster.
2. Migrate the configuration of device 3 to 6.0.x based configuration.
3. Divert the traffic from the 2-device 5.2.1-Px cluster to the 6.0.x single device cluster using BGP AS-path manipulation.
4. Disable and remove device 2 from the 5.2.1-Px cluster and add to the 6.0.x cluster. The 6.0.x cluster is now a 2-device cluster and is carrying traffic.
5. Disable and remove device 1 from the 5.2.1-Px cluster and add to the 6.0.x cluster.
6. All the devices are migrated to the 3-device 6.0.x cluster and the cluster is fully functional.

# Migration Procedure

Consider a scenario of 5.2.1-Px based 3-device Scaleout cluster is up and running.

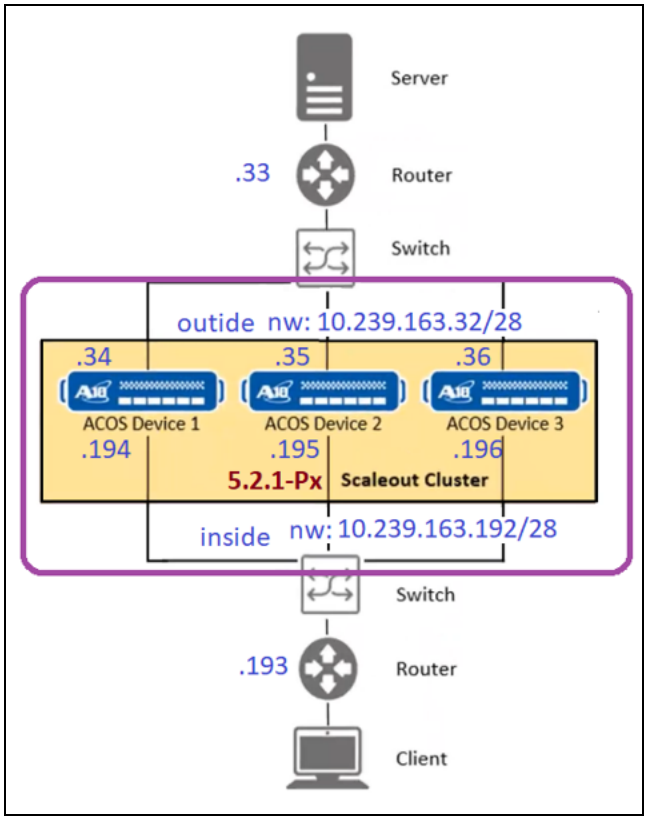


To migrate a 3-device Scaleout cluster from ACOS 5.2.1-Px to ACOS 6.0.x:

1. Check the running configuration of Scaleout in 5.2.1-Px.

```
ACOS(config)#show  running-config  scaleout
!Section configuration: 228 bytes
!
scaleout 64
 local-device
  priority 100
  id 3
 cluster-devices
  device-id 1
   ip 10.239.163.194
  device-id 2
   ip 10.239.163.195
  device-id 3
   ip 10.239.163.196
```

```
!
scaleout apps enable
!
```

a. Check the inside router's route table.

```
Gateway of last resort:
B E  0.0.0.0/0 [200/0] via 10.239.163.194, Vlan1300
                       via 10.239.163.195, Vlan1300
                       via 10.239.163.196, Vlan1300
```

2. Disable the device-id 3 from Scaleout.

```
ACOS3(config)#scaleout 64
ACOS3(config-cluster:64)#local-device
ACOS3(config-cluster:64-local-device)#disable
ACOS3(config-cluster:64-local-device)#write memory
Building configuration...

ACOS3(config-cluster:64-local-device)#show  scaleout

Device Role  : Service Node
Cluster Mode : Layer-2

Device 1 - Active (Master)
Device 2 - Active
Device 3 - Disabled (Local)
```

a. Check the inside router's route table.

```
Gateway of last resort:
B E      0.0.0.0/0 [200/0] via 10.239.163.194, Vlan1300
                           via 10.239.163.195, Vlan1300
```

3. Wait for the Scaleout status to update on the remaining two devices.

a. Check the Scaleout status on device 1.

```
ACOS1#show scaleout
Device Role  : Cluster Master
Cluster Mode : Layer-2

Device 1 - Active (Local) (Master)
```

```
Device 2 - Active
Device 3 - Disabled

ACOS1#show scaleout
Device Role  : Cluster Master
Cluster Mode : Layer-2

Device 1 - Active (Local) (Master)
Device 2 - Active
```

4. Remove the device-id 3 from the cluster-devices configuration on the remaining two devices.
   a. Remove the device-id configuration from the cluster on device 2.

```
ACOS2#config
ACOS2(config)#scaleout 64
ACOS2(config-cluster:64)#cluster-devices
ACOS2(config-cluster:64-cluster-devices)#no device-id 3
Please configure on each node in the cluster
ACOS2(config-cluster:64-cluster-devices)#show  scaleout
Device Role  : Service Node
Cluster Mode : Layer-2

Device 1 - Active (Master)
Device 2 - Active (Local)

ACOS2(config-cluster:64-cluster-devices)#write memory
Building configuration...
```

   b. Remove the device-id configuration from the cluster on device 1.

```
ACOS1#config
ACOS1(config)#scaleout 64
ACOS1(config-cluster:64)#cluster-devices
ACOS1(config-cluster:64-cluster-devices)#no device-id 3
Please configure on each node in the cluster
ACOS1(config-cluster:64-cluster-devices)#write memory
Building configuration...
```

5. Upgrade only the device 3 to 6.0.x.
6. Shut down the BGP neighbors on device 3.

```
ACOS3#config
ACOS3(config)#router bgp 2400
ACOS3(config-bgp:2400)#neighbor 10.239.163.193 shutdown
ACOS3(config-bgp:2400)#neighbor 10.239.163.33 shutdown
ACOS3(config-bgp:2400)#write memory
Building configuration...
```

7. Configure Scaleout on device 3 running 6.0.x code to create a new Scaleout cluster.

   a. In 6.0.x, configure `vrrp-a common`, `vcs database-distribution enable`, and `vcs device` before enabling the Scaleout.

   **NOTE:** In the ACOS 6.0.x release, the default aVCS multicast IP address has been changed from 224.0.0.210 to 224.0.1.210. If you want to continue using the old 224.0.0.210 as the multicast IP address in the aVCS deployment after upgrading to ACOS 6.0.x, you may need to manually add a line of configuration: `vcs multicast-ip 224.0.0.210`.

```
ACOS3#config
ACOS3(config)#vrrp-a common
ACOS3(config-common)#set-id 14
ACOS3(config-common)#device-id 3
ACOS3(config-common)#exit
ACOS3(config)#vcs database-distribution enable
The changed configuration of aVCS will take effect only after 'vcs
reload'

ACOS3(config)#vcs device 3
The changed configuration of aVCS will take effect only after 'vcs
reload'.
ACOS3(config-device:3)#interfaces management
The changed configuration of aVCS will take effect only after 'vcs
reload'.
Run 'vcs reload cluster-discovery' if only aVCS interface was added
or removed.
ACOS3(config-device:3)#enable
```

```
The changed configuration of aVCS will take effect only after 'vcs
reload'.
ACOS3(config-device:3)#priority 148
The changed configuration of aVCS will take effect only after 'vcs
reload'.
ACOS3(config-device:3)#vcs reload

System configuration has been modified. Save? [yes/no]:yes
Building configuration...
```

b. Verify the running configuration of Scaleout on 6.0.x.

```
ACOS3#show running-config scaleout
!Section configuration: 86 bytes
!
scaleout 64
 local-device
  priority 100
  disable
!
scaleout apps enable
!
```

c. Enable Scaleout.

```
ACOS3#config
ACOS3(config)#scaleout 64
ACOS3(config-cluster:64)#local-device
ACOS3(config-cluster:64-local-device)#enable
ACOS3(config-cluster:64-local-device)#end
ACOS3#write memory
Building configuration...

ACOS3#show running-config scaleout
!Section configuration: 86 bytes
!
scaleout 64
 local-device
  priority 100
!
```

```
scaleout apps enable
!
```

d. Check the Scaleout status.

```
ACOS3#show scaleout
Device Role  : Unknown Node
Cluster Mode : Layer-2

ACOS3#
ACOS3#show scaleout
Device Role  : Cluster Master
Cluster Mode : Layer-2

Device 3 - Active (Local) (Master)
ACOS3#
```

Now there are two clusters, one is a 2-device Scaleout cluster running 5.2.1-Px (devices 1 and 2) and the other running 6.0.x (device 3).

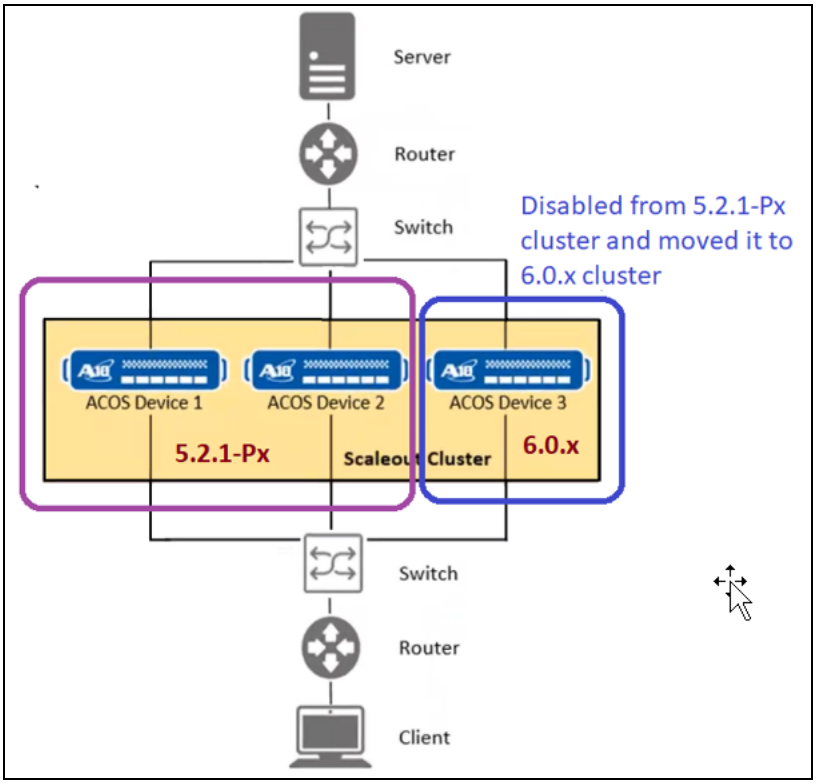

The 6.0.x cluster has BGP peering down and cannot handle traffic.

8. Bring up the 6.0.x cluster to involve in the BGP routing and ensure that traffic goes through the older cluster. Use the `as-path` attribute to influence the routes in both inside and outside directions.
9. Configure the BGP `as-path prepend` attribute to the default routes and NAT IPs on the 6.0.x cluster with a longer path than the 5.2.1-Px cluster.
   a. Modify the route-map on device 3 (6.0.x) to advertise routes with longer `as-path` set.

```
ACOS3#config
ACOS3(config)#route-map 151.254.128.0 permit 1
ACOS3(config-route-map:1)#set as-path prepend 1000 1001
ACOS3(config-route-map:1)#route-map DG permit 1
ACOS3(config-route-map:1)#set as-path prepend 23 24
ACOS3(config-route-map:1)#end
ACOS3(config-bgp:2400)#write memory
Building configuration...
```

10. Bring up the BGP neighbors on the 6.0.x cluster. The traffic continues to the 5.2.1-Px cluster.
    a. Use no shutdown bgp neighbors on device 3.

```
ACOS3#config
ACOS3(config)#router bgp 2400
ACOS3(config-bgp:2400)#no neighbor 10.239.163.193 shutdown
ACOS3(config-bgp:2400)#no neighbor 10.239.163.33 shutdown
ACOS3(config-bgp:2400)#write memory
Building configuration...
```

    b. Check the inside router's route table.

```
Gateway of last resort:
B E      0.0.0.0/0 [200/0] via 10.239.163.194, Vlan1300
                           via 10.239.163.195, Vlan1300
Network            Next Hop        Metric  LocPref Weight  Path
* >Ec   0.0.0.0/0  10.239.163.194  0       100     0       2400 1 ?
*  ec   0.0.0.0/0  10.239.163.195  0       100     0       2400 1 ?
*       0.0.0.0/0  10.239.163.196  0       100     0       2400 23
24 1 ?
```

c. Check the outside router's route table.

```
B E      151.254.128.3/32 [200/0] via 10.239.163.34, Vlan1301
B E      151.254.128.4/32 [200/0] via 10.239.163.35, Vlan1301
Network                Next Hop        Metric  LocPref Weight  Path
* > 151.254.128.3/32   10.239.163.34   0       100     0       2400
?
*   151.254.128.3/32   10.239.163.36   0       100     0       2400
1000 1001 ?
* > 151.254.128.4/32   10.239.163.36   0       100     0       2400
1000 1001 ?
*   151.254.128.4/32   10.239.163.35   0       100     0       2400
?
```

11. Migrate traffic to the 6.0.x cluster by adjusting the BGP `as-path` attribute.
    a. Change the `as-path` on devices 1 and 2 and make device 3 to take the traffic.
        i. Use the same `as-path` attribute to influence the incoming traffic to select the device-3 as next-hop so that the traffic gets migrated to the 6.0.x cluster.
        ii. Add `as-path length` on devices 1 and 2 longer than the device 3's routes.
            i. Modify the route-map on device 1 (5.2.1-Px) to advertise routes with longer `as-path` set than device 3 (6.0.x).

```
ACOS1#config
ACOS1(config)#route-map DG permit 1
ACOS1(config-route-map:1)#set as-path prepend 25 26 27
ACOS1(config-route-map:1)#route-map 151.254.128.0 permit 1
ACOS1(config-route-map:1)#set as-path prepend 1002 1003 1004
ACOS1(config-route-map:1)#write memory
```

            ii. Modify the route-map on device 2 (5.2.1-Px) to advertise routes with longer `as-path` set than device 3 (6.0.x).

```
ACOS2#config
ACOS2(config)#route-map 151.254.128.0 permit 1
ACOS2(config-route-map:1)#set as-path prepend 1002 1003 1004
ACOS2(config-route-map:1)#route-map DG permit 1
ACOS2(config-route-map:1)#set as-path prepend 25 26 27
ACOS2(config-route-map:1)#write memory
```

```
Building configuration...
```

i. Check the inside router's route table.

```
Gateway of last resort:
B E      0.0.0.0/0 [200/0] via 10.239.163.196, Vlan1300
Network         Next Hop         Metric  LocPref Weight
Path
* > 0.0.0.0/0  10.239.163.195  0       100     0
2400 25 26 27 1 ?
*   0.0.0.0/0  10.239.163.196  0       100     0
2400 23 24 1 ?
*   0.0.0.0/0  10.239.163.194  0       100     0
2400 25 26 27 1 ?
```

ii. Check the outside router's route table.

```
B E      151.254.128.3/32 [200/0] via 10.239.163.36,
Vlan1301
B E      151.254.128.4/32 [200/0] via 10.239.163.36,
Vlan1301
B E      151.254.128.5/32 [200/0] via 10.239.163.36,
Vlan1301
Network               Next Hop       Metric  LocPref
Weight  Path
* > 151.254.128.3/32  10.239.163.36  0       100     0
  2400 1000 1001 ?
*   151.254.128.3/32  10.239.163.34  0       100     0
  2400 1002 1003 1004 ?
* > 151.254.128.4/32  10.239.163.36  0       100     0
  2400 1000 1001 ?
*   151.254.128.4/32  10.239.163.35  0       100     0
  2400 1002 1003 1004 ?
```

iii. Remove the `as-path` on device 3, which continues to take the traffic.

```
ACOS3#config
ACOS3(config)#route-map DG permit 1
```

```
ACOS3(config-route-map:1)#no set as-path prepend 23 24
ACOS3(config-route-map:1)#route-map 151.254.128.0 permit 1
ACOS3(config-route-map:1)#no set as-path prepend 1000 1001
ACOS3(config-route-map:1)#write memory
Building configuration...
```

**NOTE:** Traffic drop is expected for the existing sessions. So, the source must reinitiate the traffic. This is because the session synchronization between the 2 different Scaleout clusters (5.2.1-Px and 6.0.x) is not supported.

12. Disable the device-id 2 from the 5.2.1-Px scaleout.

```
ACOS2#config
ACOS2(config)#scaleout 64
ACOS2(config-cluster:64)#local-device
ACOS2(config-cluster:64-local-device)#disable
ACOS2(config-cluster:64-local-device)#end

ACOS2#show scaleout

Device Role  : Service Node
Cluster Mode : Layer-2

Device 1 - Active (Master)
Device 2 - Disabled (Local)

ACOS2#show scaleout
Device Role  : Scaleout is not Active

ACOS2#write memory
Building configuration...
```

13. Upgrade the device-id 2 to 6.0.x, add to the 6.0.x cluster, and adjust BGP `as-path`.
    a. Remove the `as-path` on device 2.
    b. Device 2 is upgraded to 6.0.x. Enable Scaleout configurations.

       Verify the running configuration of Scaleout on device 2.

```
ACOS2#show running-config  scaleout
```

```
!Section configuration: 86 bytes
!
scaleout 64
 local-device
  priority 100
  disable
!
scaleout apps enable
!
```

c. Enable the `vrrp-a common` and `vcs database distribution` configurations on device 2.

```
ACOS2#config
ACOS2(config)#vrrp-a common
ACOS2(config-common)#set-id 14
ACOS2(config-common)#device-id 2
ACOS2(config-common)#exit
ACOS2(config)#vcs database-distribution enable
The changed configuration of aVCS will take effect only after 'vcs
reload'
```

d. Enable the `vcs device 2` configuration on device 2.

```
ACOS2(config)#vcs device 2
The changed configuration of aVCS will take effect only after 'vcs
reload'.
ACOS2(config-device:2)#interfaces management
The changed configuration of aVCS will take effect only after 'vcs
reload'.
Run 'vcs reload cluster-discovery' if only aVCS interface was added
or  removed.
ACOS2(config-device:2)#priority 149
The changed configuration of aVCS will take effect only after 'vcs
reload'.
ACOS2(config-device:2)#enable
The changed configuration of aVCS will take effect only after 'vcs
reload'.
```

e. Enable `scaleout` and `vcs reload` on device 2.

```
ACOS2(config-device:2)#scaleout 64
ACOS2(config-cluster:64)#local-device
ACOS2(config-cluster:64-local-device)#enable
ACOS2(config-cluster:64-local-device)#vcs reload
System configuration has been modified. Save? [yes/no]:yes
Building configuration...
```

14. The 2-device 6.0.x Scaleout cluster takes traffic.

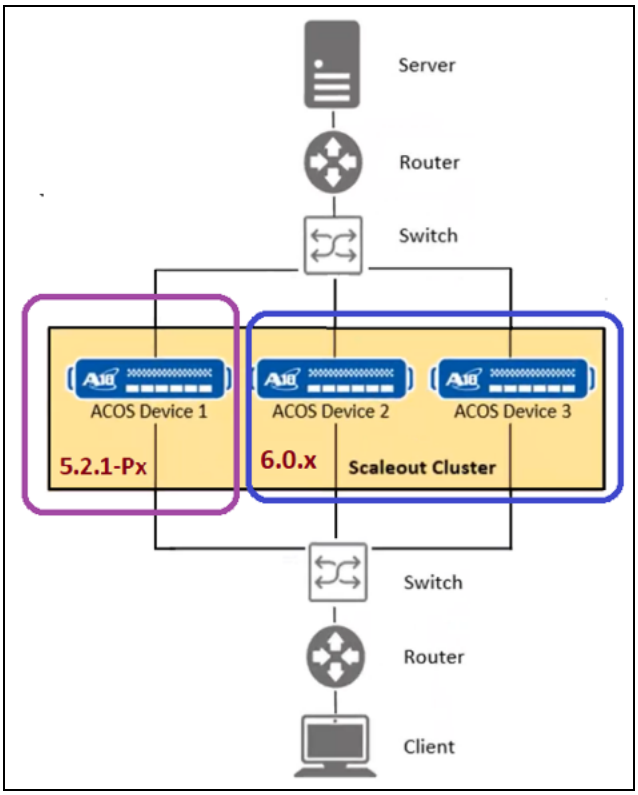


a. Check the Scaleout status on device 2 after adding it to the 6.0.x cluster.

```
ACOS2#show scaleout

Device Role  : Unknown Node
Cluster Mode : Layer-2

Device 3 - Active (Master)
ACOS2#
ACOS2#show scaleout
```

```
Device Role   : Service Node
Cluster Mode : Layer-2

Device 2 - Active (Local)
Device 3 - Active (Master)
ACOS2#
```

b. Remove the `as-path prepend` list on device 2 after adding it to the 6.0.x cluster.

```
ACOS2#config
ACOS2(config)#route-map DG permit 1
ACOS2(config-route-map:1)#no   set as-path prepend 25 26 27
ACOS2(config-route-map:1)#route-map 151.254.128.0 permit 1
ACOS2(config-route-map:1)#no   set as-path prepend 1002 1003 1004
ACOS2(config-route-map:1)#write memory
Building configuration...
```

15. Follow the Step 13 and Step 14 to add device 1 into the 6.0.x cluster.
16. Disable the device-id 1 from the 5.2.1-Px Scaleout, upgrade the device-id 1 to 6.0.x, add to the 6.0.x cluster, and remove the BGP `as-path` prepend list.
    a. Disable device 1 from the 5.2.1-Px Scaleout cluster.

```
ACOS1#config
ACOS1(config)#scaleout 64
ACOS1(config-cluster:64)#local-device
ACOS1(config-cluster:64-local-device)#disable

ACOS1(config-cluster:64-local-device)#show  scaleout
Device Role   : Scaleout is not Active

ACOS1(config-cluster:64-local-device)#end
ACOS1#write memory
Building configuration...
```

    b. After upgrading to 6.0.x, on device 1, configure the following, and then add it into the 6.0.x Scaleout cluster.
        i. Verify the running configuration of Scaleout on device 1.

```
ACOS1#show running-config  scaleout
```

```
!Section configuration: 86 bytes
!
scaleout 64
 local-device
  priority 101
  disable
!
scaleout apps enable
!
```

ii. Enable Scaleout on device 1.

```
ACOS1(config)#scaleout 64
ACOS1(config-cluster:64)#local-device
ACOS1(config-cluster:64-local-device)#enable
ACOS1(config-cluster:64-local-device)#end
```

iii. Enable the `vrrp-a common` and `vcs database distribution` configurations on device 1.

```
ACOS1#config
ACOS1(config)#vrrp-a common
ACOS1(config-common)#set-id 14
ACOS1(config-common)#device-id 1
ACOS1(config-common)#exit
ACOS1(config)#vcs database-distribution enable
The changed configuration of aVCS will take effect only after
'vcs reload'
```

iv. Enable the `vcs device 1` configuration on device 1.

```
ACOS1(config)#vcs device 1
The changed configuration of aVCS will take effect only after
'vcs reload'.
ACOS1(config-device:1)#interfaces management
The changed configuration of aVCS will take effect only after
'vcs reload'.
Run 'vcs reload cluster-discovery' if only aVCS interface was
added or removed.
ACOS1(config-device:1)#priority 150
```

```
The changed configuration of aVCS will take effect only after
'vcs reload'.
ACOS1(config-device:1)#enable
The changed configuration of aVCS will take effect only after
'vcs reload'.
ACOS1(config-device:1)#vcs reload
System configuration has been modified. Save? [yes/no]:yes
Building configuration...
```

The 3-device 6.0.x Scaleout cluster takes traffic.

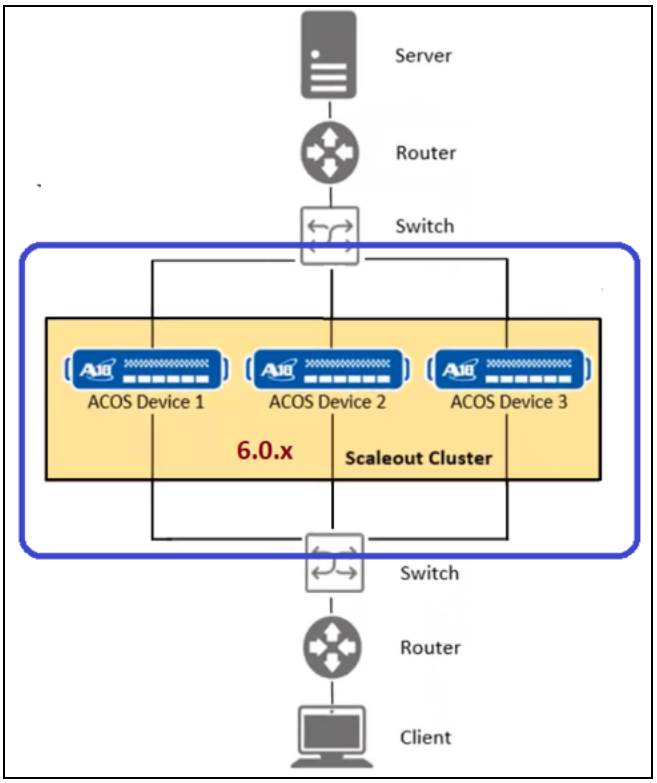


v. Check the Scaleout status on device 1 after adding it to the 6.0.x cluster.

```
ACOS1#show scaleout

Device Role   : Unknown Node
Cluster Mode : Layer-2

Device 2 - Active
```

```
Device 3 - Active (Master)

ACOS1#show scaleout

Device Role  : Service Node
Cluster Mode : Layer-2

Device 1 - Active (Local)
Device 2 - Active
Device 3 - Active (Master)
ACOS1#
```

c. Remove the `as-path` prepend list on device 2 after adding it to the 6.0.x cluster.

```
ACOS1#config
ACOS1(config)#route-map DG permit 1
ACOS1(config-route-map:1)#no   set as-path prepend 25 26 27
ACOS1(config-route-map:1)#route-map 151.254.128.0 permit 1
ACOS1(config-route-map:1)#no   set as-path prepend 1002 1003 1004
ACOS1(config-route-map:1)#end
ACOS1#write memory
Building configuration...
```

The 3-device Scaleout cluster is successfully migrated from ACOS 5.2.1-Px to ACOS 6.0.x.

**NOTE:** Save the configuration on each device after changing every configuration.

## To migrate a 4-device Scaleout cluster from ACOS 6.0.x to ACOS 6.0.3-P1

Before upgrading consider the following:

- The scaleout/traffic under the L3V partition is not supported.
- The traffic with the new IPv6 traffic map on the new node is not supported in any partition.

To upgrade the Scaleout cluster with 4-devices having CGN or FW deployment to 6.0.3-P1:

1. Use the `vcs vmaster-take-over` command to switch the aVCS vMaster to be the same node as the current Scaleout Master (maybe device 1).
2. If the upgrade is performed on device-id 4, use the `disable` command to disable device-4 from Scaleout.
3. Configure the `vcs force-wait-interval` command with a greater time length to prevent device-4 from preempting the vMaster after upgrade or reboot.
4. Execute the `write memory` command to save the configuration.

**NOTE:** During the upgrade process, the vMaster or Scaleout Master device must remain stable and prevent preemption in the cluster.

5. Upgrade device-id 4 with the new image and wait for aVCS to join completely.
6. Enable Scaleout on device-id 4 using the `enable` command. Wait for the Scaleout to join completely.
7. Perform the steps from 1 to 6 for other devices. Retain device-id 1 ( vMaster or Scaleout Master) as the last one.

# Limitation

Due to the incompatibility, there will be no session synchronization between the 2 clusters. It implies that the existing connections going through the ACOS 5.2.1-Px cluster could reset when you switch traffic to the ACOS 6.0.x cluster.

# Scaleout Configuration Migration from ACOS 5.2.1-Px to ACOS 6.0.x Using Script

This section describes a sequence of steps to migrate the configuration of devices in a Scaleout cluster running on the ACOS 5.2.1-Px version to the ACOS 6.0.x version.

**NOTE:** The `sovcs_transition_create_startupcfg.py` and `sovcs_transition_overwrite_startupcfg.py` scripts are provided to migrate the configuration to the latest ACOS 6.0.x version with Scaleout and VCS. Contact the A10 Sales team for procuring the scripts.

The configuration migration can be performed using any one of the following modes:

- Online Mode - In this mode, the script `sovcs_transition_create_startupcfg.py` connects to the device and converts the existing **startup.cfg** file to a new **startup.cfg** file that is compatible with the 6.0.x. The new **startup.cfg** file is saved in a folder in the same directory where the script is running.
- Offline Mode - In this mode, the script `sovcs_transition_create_startupcfg.py` runs locally and does not connect to the device. You must provide the **startup.cfg** file path, which will be converted to a new **startup.cfg** file that is compatible with 6.0.x. The new **startup.cfg** file is saved in a folder in the same directory where the script is running.

The following topics are covered:



## Prerequisites

The following are the setup requirements for the Online and Offline mode migration:

- 5.2.1-Px Scaleout cluster must be present in the setup
- VCS must not be enabled in the 5.2.1-Px setup
- Python version must be >= 3.2

# Online Mode

To migrate the configuration to the latest ACOS 6.0.x in the online mode, perform the following:

1. Log in to any environment (Windows or Linux) where Python is installed.
2. Open the command or shell prompt.
3. Change the directory to the folder where the script is saved.
4. Execute the following command to retrieve the existing **startup-cfg** file and generate the new **startup-cfg**:

   **python sovcs_transition_create_startupcfg.py** `<online> <device IP> <username> <password>`

   Where,

   - `online` – Online mode
   - `device IP` – IP address of the device. For example, 10.2.2.2
   - `username` – Username to connect the device
   - password – Password to connect the device

   For example, **`python sovcs_transition_create_startupcfg.py online 10.2.2.2 admin xxx`**.

   If the current **startup-cfg** file name on the device is **sovcs-demo**, a new **startup-cfg** file is created in the folder named **TH-< device-id>_<ip_addess>_<timestamp>**. The device-id is from the device's Scaleout configuration.

5. Upgrade the device to the ACOS 6.0.x version manually. For more information about the upgrade, see *Release Notes*.

   After the upgrade, Scaleout on this device will be in a DOWN state.

6. Execute the following command to overwrite the current **startup-cfg** file with the

new **startup-cfg** file on the device:

```
python sovcs_transition_overwrite_startupcfg.py <IP/hostname>
<username> <password> <file/file path>
```

Where,

- `IP/hostname` – IP address or hostname of the device. For example, 10.2.2.2.
- `username` – Username to connect the device
- `password` – Password to connect the device
- `file/file path` – Full path name of the new **startup-cfg** file on the server

For example, **`python sovcs_transition_overwrite_startupcfg.py 10.2.2.2 admin xxx sovcs-demo`**.

7. Reload the device. If prompted to save the configurations, select **No**.
8. After the device comes up, check if the Scaleout status is UP.
9. Perform the steps from 1 to 5 for other devices.

The configuration of the devices in the Scaleout cluster is successfully migrated from ACOS 5.2.1-Px to ACOS 6.0.x.

# Offline Mode

To migrate the configuration to the latest ACOS 6.0.x in the offline mode, perform the following:

1. Log in to any environment (Windows or Linux) where Python is installed.
2. Open the command or shell prompt.
3. Change the directory to the folder where the script is saved.
4. Execute the following command to generate a new file:

```
python sovcs_transition_create_startupcfg.py <offline> <filepath>
```

Where,

- `offline` – Offline mode
- `filepath` – Full path name of the old **startup-cfg** file on the server

For example, **`python sovcs_transition_create_startupcfg.py offline sovcs.demo`**.

The new file is updated with the 6.0.x compatible configurations and stored locally. If the current **startup-cfg** file name on the device is **sovcs-demo**, a new **startup-cfg** file is created in the folder named **TH-< device-id>_<timestamp>**. The device-id is from the device's Scaleout configuration.

5. To create the current (old) **startup-cfg** file from the device, execute the `show startup-config` command, copy the entire output till the last line that ends with `end`, and save.
6. Save the configuration in a file with the same name as the configuration profile name (first line of the output).

   In the example, the file name is `ipv6-so-l3mode` without any file extension.

   An example of the `startup-config` command output is shown below.

```
ACOS-vMaster[13/2]#show startup-config
Show configuration profile "ipv6-so-l3mode"
Building configuration...

!Current configuration: 13058 bytes
!Configuration last updated at 16:01:00 +03 Wed Feb 21 2024
!Configuration last saved at 16:27:20 +03 Wed Feb 21 2024
!64-bit Advanced Core OS (ACOS) version 6.0.3-P1, build 60 (Feb-20-
2024,04:46)
!
vrrp-a common
  device-id 2
  set-id 13
  exit-module
!
vcs enable
!
.....
.....
```

```
  rule deny
    action deny
    source ipv4-address any
    source zone any
    dest ipv4-address any
    dest zone any
    service any
    exit-module
  exit-module
!
fw active-rule-set firewall
!
  End
```

7. Use the old **startup-cfg** file to generate the ACOS 6.0.x compatible **startup-cfg** file.
8. Upgrade the device to the ACOS 6.0.x version manually. For more information about the upgrade, see *Release Notes*.

   After the upgrade, Scaleout on this device will be in a DOWN state.
9. Execute the following command to overwrite the current **startup-cfg** file with the new **startup-cfg** file:

   **python sovcs_transition_overwrite_startupcfg.py** <IP/hostname> <username> <password> <file/file path>

   Where,

   - `IP/hostname` – IP address or hostname of the device. For example, 10.2.2.2
   - `username` – Username to connect the device
   - `password` – Password to connect the device
   - `file/file path` – Full path name of the new **startup-cfg** file on the server

   For example, **sovcs_transition_overwrite_startupcfg.py 10.2.2.2 admin xxx sovcs-demo**.

   It takes the backup of all the profiles if they are already not backed up by the **sovcs_transition_create_startupcfg.py** script.
10. Reload the device. If prompted to save the configurations, select **No**.

11. After the device comes up, check if the Scaleout status is UP.
12. Perform the steps from 1 to 5 for other devices.

The configuration of the devices in the Scaleout cluster is successfully migrated from ACOS 5.2.1-Px to 6.0.x.

# Upgrading Scaleout/aVCS Cluster from pre-ACOS 6.0.6 to ACOS 6.0.6 and Later Releases

This chapter describes a sequence of steps to upgrade a two device Scaleout and Virtual Chassis Systems (aVCS) cluster from a pre-ACOS 6.0.6 image to ACOS 6.0.6 under different scenarios for the following platforms:

- Multi-PU Platform
- Non-Multi-PU Platform

**NOTE:** These upgrade steps apply only to upgrading from pre-ACOS 6.0.6 to ACOS 6.0.6 and are not intended for upgrades from ACOS 6.0.6 to later releases.

The following topics are covered:



## Multi-PU Platform

The upgrade procedure for different scenarios on a multi-PU platform is as follows. This upgrade procedure must be performed during a maintenance window.

**Scenario 1** - Upgrade a cluster running with the default pre-ACOS 6.0.6 aVCS multicast IP as 224.0.1.210 to a new ACOS 6.0.6 image.

**Pre-Requisites:**

- Two multi-PU devices in a Scaleout/aVCS cluster setup.
- aVCS database-distribution enabled for Scaleout, aVCS multicast discovery mode is set to multicast by default, and the multicast IP is 224.0.1.210.

  You may verify the current aVCS multicast IP address using the `show vcs summary` command.

**Upgrade Steps:**

1. Disable Scaleout on device 2.

```
ACOS2(config)#show  scaleout
Device Role  : Service Node
Cluster Mode : Layer-2

Device 1 - Active (Master)
Device 2 - Active (Local)
ACOS2(config)#

ACOS2(config)#scaleout 1
ACOS2(config-cluster:1)#local-device
ACOS2(config-cluster:1-local-device)#disable
```

2. Wait for 30 seconds until device 2 exits the Scaleout cluster.
3. Execute the `write memory` command, upgrade device 2 with the ACOS 6.0.6 image, and reboot.
4. Wait for device 2 to boot up and confirm that aVCS is ready and stable by executing the `show vcs summary` command.
5. On device 2, execute the `vcs multicast-ip 224.0.1.210` command followed by the `write memory` command.
6. Execute the `vcs reload` command on device 2 and wait for aVCS to become stable.

   Device 2 rejoins the aVCS cluster.
7. Enable Scaleout on device 2 (service-node).

```
ACOS2(config)#scaleout 1
ACOS2(config-cluster:1)#local-device
ACOS2(config-cluster:1-local-device)#enable
```

8. Wait for device 2 to rejoin the Scaleout cluster and confirm the traffic-map is stable.
9. On device 1, perform Steps 1 to 8.

   The upgrade to ACOS 6.0.6 is successfully completed.

**Scenario 2** - Upgrade to ACOS 6.0.6 in a configuration where the current multicast IP is not the default multicast IP 224.0.1.210

**Pre-Requisites:**

- Two multi-PU devices in a Scaleout/aVCS cluster setup.
- aVCS database-distribution enabled for Scaleout and aVCS discovery mode is set to multicast by default and the multicast IP is not equal to 224.0.1.210.

  You may verify the current aVCS multicast IP address using the `show vcs summary` command.

**Upgrade Steps:**

1. Upgrade device 2 with the ACOS 6.0.6 image and reboot.
2. Wait for device 2 to boot up, rejoin the Scaleout cluster, and confirm the traffic-map is stable.
3. On device 1, perform Steps 1 and 2.

   The upgrade to ACOS 6.0.6 is successfully completed.

# Non-Multi-PU Platform

The upgrade procedure for a Non-Multi-PU platform is as follows:

**Pre-Requisites:**

- Two devices in a Scaleout/aVCS cluster setup.
- aVCS database-distribution enabled for Scaleout, configuration-sync, and aVCS discovery mode is set to multicast by default, and multicast IP is 224.0.1.210.

**Upgrade Steps:**

1. On the pre-ACOS 6.0.6 cluster, change the aVCS multicast IP address from 224.0.1.210 to 224.0.0.211 on vMaster (device 1) using the `vcs multicast-ip 224.0.0.211` command. Ensure that the configurations sync to the vBlade (device 2).

**NOTE:** You may use any new multicast IP address. Here, 224.0.0.211 is shown as an example.

2. Save the configurations and perform VCS reload using the `vcs reload db-safe start` command on vMaster.
3. Verify that both vMaster and vBlade in the pre-ACOS 6.0.6 cluster are up and functional.
4. Start the staggered upgrade on vBlade from vMaster with the ACOS 6.0.6 image and reboot.

   vMaster handles all traffic.
5. Check that device 2, now running ACOS 6.0.6 is up and rejoined the Scaleout/aVCS cluster as vBlade and started handling traffic.
6. In a larger cluster with more than two devices, follow steps 4 and 5 for all vBlades and migrate them to the ACOS 6.0.6 image one by one. This must be done using the staggered upgrade from the vMaster.
7. Disable Scaleout on aVCS vMaster.

   After upgrading all vBlades, vMaster is the last device in the pre-ACOS 6.0.6 cluster.

```
ACOS1-vMaster[1/1](config:1)#scaleout 64
ACOS1-vMaster[1/1](config:1-cluster:64)#local-device
ACOS1-vMaster[1/1](config:1-cluster:64-local-device)#disable
This operation applied to device 1
```

8. Wait for device 1 to completely leave the Scaleout cluster and the traffic-map to become stable on the ACOS 6.0.6 cluster (device 2).
9. On vMaster (device 1), execute the `vcs disable` command, upgrade it with the ACOS 6.0.6 image, and reboot without saving the configurations.

   vBlade (device 2) must become vMaster when aVCS is disabled on device 1 and handles traffic.
10. Verify that device 1 rejoins the Scaleout/aVCS cluster as vBlade, now running ACOS 6.0.6, and starts handling traffic.
11. Confirm that both devices in the Scaleout/aVCS cluster are running ACOS 6.0.6

and functional.

The upgrade to ACOS 6.0.6 is successfully completed.

## Contact Us