



HSRP

This chapter contains the following sections:

- [About HSRP, on page 1](#)
- [About Cisco APIC and HSRP, on page 2](#)
- [HSRP Versions, on page 3](#)
- [Guidelines and Limitations, on page 3](#)
- [Default HSRP Settings , on page 5](#)
- [Configuring HSRP Using the GUI, on page 5](#)
- [Configuring HSRP in Cisco APIC Using Inline Parameters in NX-OS Style CLI, on page 7](#)
- [Configuring HSRP in Cisco APIC Using Template and Policy in NX-OS Style CLI, on page 8](#)
- [Configuring HSRP in APIC Using REST API, on page 9](#)

About HSRP

HSRP is a first-hop redundancy protocol (FHRP) that allows a transparent failover of the first-hop IP router. HSRP provides first-hop routing redundancy for IP hosts on Ethernet networks configured with a default router IP address. You use HSRP in a group of routers for selecting an active router and a standby router. In a group of routers, the active router is the router that routes packets, and the standby router is the router that takes over when the active router fails or when preset conditions are met.

Many host implementations do not support any dynamic router discovery mechanisms but can be configured with a default router. Running a dynamic router discovery mechanism on every host is not practical for many reasons, including administrative overhead, processing overhead, and security issues. HSRP provides failover services to such hosts.

When you use HSRP, you configure the HSRP virtual IP address as the default router of the host (instead of the IP address of the actual router). The virtual IP address is an IPv4 or IPv6 address that is shared among a group of routers that run HSRP.

When you configure HSRP on a network segment, you provide a virtual MAC address and a virtual IP address for the HSRP group. You configure the same virtual address on each HSRP-enabled interface in the group. You also configure a unique IP address and MAC address on each interface that acts as the real address. HSRP selects one of these interfaces to be the active router. The active router receives and routes packets destined for the virtual MAC address of the group.

HSRP detects when the designated active router fails. At that point, a selected standby router assumes control of the virtual MAC and IP addresses of the HSRP group. HSRP also selects a new standby router at that time.

HSRP uses a priority designator to determine which HSRP-configured interface becomes the default active router. To configure an interface as the active router, you assign it with a priority that is higher than the priority of all the other HSRP-configured interfaces in the group. The default priority is 100, so if you configure just one interface with a higher priority, that interface becomes the default active router.

Interfaces that run HSRP send and receive multicast User Datagram Protocol (UDP)-based hello messages to detect a failure and to designate active and standby routers. When the active router fails to send a hello message within a configurable period of time, the standby router with the highest priority becomes the active router. The transition of packet forwarding functions between the active and standby router is completely transparent to all hosts on the network.

You can configure multiple HSRP groups on an interface. The virtual router does not physically exist but represents the common default router for interfaces that are configured to provide backup to each other. You do not need to configure the hosts on the LAN with the IP address of the active router. Instead, you configure them with the IP address of the virtual router (virtual IP address) as their default router. If the active router fails to send a hello message within the configurable period of time, the standby router takes over, responds to the virtual addresses, and becomes the active router, assuming the active router duties. From the host perspective, the virtual router remains the same.




Note Packets received on a routed port destined for the HSRP virtual IP address terminate on the local router, regardless of whether that router is the active HSRP router or the standby HSRP router. This process includes ping and Telnet traffic. Packets received on a Layer 2 (VLAN) interface destined for the HSRP virtual IP address terminate on the active router.

About Cisco APIC and HSRP

HSRP in Cisco ACI is supported only on routed-interface or sub-interface. Therefore HSRP can only be configured under Layer 3 Out. Also there must be Layer 2 connectivity provided by external device(s) such as a Layer 2 switch between ACI leaf switches running HSRP because HSRP operates on leaf switches by exchanging Hello messages over external Layer 2 connections. An HSRP hello message does not pass through the spine switch.

The following is an example topology of an HSRP deployment in Cisco APIC.

Figure 1: HSRP Deployment Topology



HSRP Versions

Cisco APIC supports HSRP version 1 by default. You can configure an interface to use HSRP version 2.

HSRP version 2 has the following enhancements to HSRP version 1:

- Expands the group number range. HSRP version 1 supports group numbers from 0 to 255. HSRP version 2 supports group numbers from 0 to 4095.
- For IPv4, uses the IPv4 multicast address 224.0.0.102 or the IPv6 multicast address FF02::66 to send hello packets instead of the multicast address of 224.0.0.2, which is used by HSRP version 1.
- Uses the MAC address range from 0000.0C9F.F000 to 0000.0C9F.FFFF for IPv4 and 0005.73A0.0000 through 0005.73A0.0FFF for IPv6 addresses. HSRP version 1 uses the MAC address range 0000.0C07.AC00 to 0000.0C07.ACFF.

Guidelines and Limitations

Follow these guidelines and limitations:

- The HSRP state must be the same for both HSRP IPv4 and IPv6. The priority and preemption must be configured to result in the same state after failovers.

- Currently, only one IPv4 and one IPv6 group is supported on the same sub-interface in Cisco ACI. Even when dual stack is configured, Virtual MAC must be the same in IPv4 and IPv6 HSRP configurations.
- BFD IPv4 and IPv6 is supported when the network connecting the HSRP peers is a pure layer 2 network. You must configure a different router MAC address on the leaf switches. The BFD sessions become active only if you configure different MAC addresses in the leaf interfaces.
- Users must configure the same MAC address for IPv4 and IPv6 HSRP groups for dual stack configurations.
- HSRP VIP must be in the same subnet as the interface IP.
- It is recommended that you configure interface delay for HSRP configurations.
- HSRP is only supported on routed-interface or sub-interface. HSRP is not supported on VLAN interfaces and switched virtual interface (SVI). Therefore, no VPC support for HSRP is available.
- Object tracking on HSRP is not supported.
- HSRP Management Information Base (MIB) for SNMP is not supported.
- Multiple group optimization (MGO) is not supported with HSRP.
- ICMP IPv4 and IPv6 redirects are not supported.
- Cold Standby and Non-Stop Forwarding (NSF) are not supported because HSRP cannot be restarted in the Cisco ACI environment.
- There is no extended hold-down timer support as HSRP is supported only on leaf switches. HSRP is not supported on spine switches.
- HSRP version change is not supported in APIC. You must remove the configuration and reconfigure with the new version.
- HSRP version 2 does not inter-operate with HSRP version 1. An interface cannot operate both version 1 and version 2 because both versions are mutually exclusive. However, the different versions can be run on different physical interfaces of the same router.
- Route Segmentation is programmed in Cisco Nexus 93128TX, Cisco Nexus 9396PX, and Cisco Nexus 9396TX leaf switches when HSRP is active on the interface. Therefore, there is no DMAC=router MAC check conducted for route packets on the interface. This limitation does not apply for Cisco Nexus 93180LC-EX, Cisco Nexus 93180YC-EX, and Cisco Nexus 93108TC-EX leaf switches.
- HSRP configurations are not supported in the Basic GUI mode. The Basic GUI mode has been deprecated starting with APIC release 3.0(1).
- Fabric to Layer 3 Out traffic will always load balance across all the HSRP leaf switches, irrespective of their state. If HSRP leaf switches span multiple pods, the fabric to out traffic will always use leaf switches in the same pod.
- This limitation applies to some of the earlier Cisco Nexus 93128TX, Cisco Nexus 9396PX, and Cisco Nexus 9396TX switches. When using HSRP, the MAC address for one of the routed interfaces or routed sub-interfaces must be modified to prevent MAC address flapping on the Layer 2 external device. This is because Cisco APIC assigns the same MAC address (00:22:BD:F8:19:FF) to every logical interface under the interface logical profiles.

Default HSRP Settings

Parameters	Default Value
Version	1
Delay	0
Reload Delay	0
Interface Control	No Use-Burned-in Address (BIA)
Group ID	0
Group Af	IPv4
IP Obtain Mode	admin
Priority	100
Hello Interval	3000 msec
Hold Interval	10000 msec
Group Control	Preemption disabled
Preempt Delay	0
Authentication Type	Plain Text
Authentication Key Timeout	0
VMAC	Derived (from HSRP groupID)

Configuring HSRP Using the GUI

HSRP is enabled when the leaf switch is configured.

Before you begin

- The tenant and VRF configured.
- VLAN pools must be configured with the appropriate VLAN range defined and the appropriate Layer 3 domain created and attached to the VLAN pool.
- The Attach Entity Profile must also be associated with the Layer 3 domain.
- The interface profile for the leaf switches must be configured as required.

Procedure

-
- Step 1** On the menu bar, click > **Tenants > Tenant_name**. In the **Navigation** pane, click **Networking > External Routed Networks > External Routed Network_name > Logical Node Profiles > Logical Interface Profile**.
An HSRP interface profile will be created here.
- Step 2** Choose a logical interface profile, and click **Create HSRP Interface Profile**.
- Step 3** In the **Create HSRP Interface Profile** dialog box, perform the following actions:
- In the **Version** field, choose the desired version.
 - In the **HSRP Interface Policy** field, from the drop-down, choose **Create HSRP Interface Policy**.
 - In the **Create HSRP Interface Policy** dialog box, in the **Name** field, enter a name for the policy.
 - In the **Control** field, choose the desired control.
 - In the **Delay** field and the **Reload Delay** field, set the desired values. Click **Submit**.
The HSRP interface policy is created and associated with the interface profile.
- Step 4** In the **Create HSRP Interface Profile** dialog box, expand **HSRP Interface Groups**.
- Step 5** In the **Create HSRP Group Profile** dialog box, perform the following actions:
- In the **Name** field, enter an HSRP interface group name.
 - In the **Group ID** field, choose the appropriate ID.
The values available depend upon whether HSRP version 1 or version 2 was chosen in the interface profile.
 - In the **IP** field, enter an IP address.
The IP address must be in the same subnet as the interface.
 - In the **MAC address** field, enter a Mac address.
 - In the **Group Name** field, enter a group name.
This is the name used in the protocol by HSRP for the HSRP MGO feature.
 - In the **Group Type** field, choose the desired type.
 - In the **IP Obtain Mode** field, choose the desired mode.
 - In the **HSRP Group Policy** field, from the drop-down list, choose **Create HSRP Group Policy**.
- Step 6** In the **Create HSRP Group Policy** dialog box, perform the following actions:
- In the **Name** field, enter an HSRP group policy name.
 - The **Key or Password** field is automatically populated.
The default value for authentication type is simple, and the key is "cisco." This is selected by default when a user creates a new policy.
 - In the **Type** field, choose the level of security desired.
 - In the **Priority** field choose the priority to define the active router and the standby router.
 - In the remaining fields, choose the desired values, and click **Submit**.
The HSRP group policy is created.
 - Create secondary virtual IPs by populating the **Secondary Virtual IPs** field.
This can be used to enable HSRP on each sub-interface with secondary virtual IPs. The IP address that you provide here also must be in the subnet of the interface.
 - Click **OK**.

- Step 7** In the **Create HSRP Interface Profile** dialog box, click **Submit**. This completes the HSRP configuration.
- Step 8** To verify the HSRP interface and group policies created, in the Navigation pane, click **Networking > Protocol Policies > HSRP**.

Configuring HSRP in Cisco APIC Using Inline Parameters in NX-OS Style CLI

HSRP is enabled when the leaf switch is configured.

Before you begin

- The tenant and VRF configured.
- VLAN pools must be configured with the appropriate VLAN range defined and the appropriate Layer 3 domain created and attached to the VLAN pool.
- The Attach Entity Profile must also be associated with the Layer 3 domain.
- The interface profile for the leaf switches must be configured as required.

Procedure

	Command or Action	Purpose
Step 1	configure Example: apicl# configure	Enters configuration mode.
Step 2	Configure HSRP by creating inline parameters. Example: <pre> apicl(config)# leaf 101 apicl(config-leaf)# interface ethernet 1/17 apicl(config-leaf-if)# hsrp version 1 apicl(config-leaf-if)# hsrp use-bia apicl(config-leaf-if)# hsrp delay minimum 30 apicl(config-leaf-if)# hsrp delay reload 30 apicl(config-leaf-if)# hsrp 10 ipv4 apicl(config-if-hsrp)# ip 182.16.1.2 apicl(config-if-hsrp)# ip 182.16.1.3 secondary apicl(config-if-hsrp)# ip 182.16.1.4 secondary apicl(config-if-hsrp)# mac-address 5000.1000.1060 apicl(config-if-hsrp)# timers 5 18 apicl(config-if-hsrp)# priority 100 </pre>	

	Command or Action	Purpose
	<pre> apic1(config-if-hsrp) # preempt apic1(config-if-hsrp) # preempt delay minimum 60 apic1(config-if-hsrp) # preempt delay reload 60 apic1(config-if-hsrp) # preempt delay sync 60 apic1(config-if-hsrp) # authentication none apic1(config-if-hsrp) # authentication simple apic1(config-if-hsrp) # authentication md5 apic1(config-if-hsrp) # authentication-key <mypassword> apic1(config-if-hsrp) # authentication-key-timeout <timeout> </pre>	

Configuring HSRP in Cisco APIC Using Template and Policy in NX-OS Style CLI

HSRP is enabled when the leaf switch is configured.

Before you begin

- The tenant and VRF configured.
- VLAN pools must be configured with the appropriate VLAN range defined and the appropriate Layer 3 domain created and attached to the VLAN pool.
- The Attach Entity Profile must also be associated with the Layer 3 domain.
- The interface profile for the leaf switches must be configured as required.

Procedure

	Command or Action	Purpose
Step 1	<pre> configure Example: apic1# configure </pre>	Enters configuration mode.
Step 2	<pre> Configure HSRP policy templates. Example: apic1(config)# leaf 101 apic1(config-leaf)# template hsrp interface-policy hsrp-intfPol1 tenant t9 apic1(config-template-hsrp-if-pol)# hsrp use-bia apic1(config-template-hsrp-if-pol)# hsrp </pre>	

	Command or Action	Purpose
	<pre> delay minimum 30 apicl(config-template-hsrp-if-pol)# hsrp delay reload 30 apicl(config)# leaf 101 apicl(config-leaf)# template hsrp group-policy hsrp-groupPoll tenant t9 apicl(config-template-hsrp-group-pol)# timers 5 18 apicl(config-template-hsrp-group-pol)# priority 100 apicl(config-template-hsrp-group-pol)# preempt apicl(config-template-hsrp-group-pol)# preempt delay minimum 60 apicl(config-template-hsrp-group-pol)# preempt delay reload 60 apicl(config-template-hsrp-group-pol)# preempt delay sync 60 </pre>	
Step 3	<p>Use the configured policy templates</p> <p>Example:</p> <pre> apicl(config)# leaf 101 apicl(config-leaf)# interface ethernet 1/17 apicl(config-leaf-if)# hsrp version 1 apicl(config-leaf-if)# inherit hsrp interface-policy hsrp-intfPoll apicl(config-leaf-if)# hsrp 10 ipv4 apicl(config-if-hsrp)# ip 182.16.1.2 apicl(config-if-hsrp)# ip 182.16.1.3 secondary apicl(config-if-hsrp)# ip 182.16.1.4 secondary apicl(config-if-hsrp)# mac-address 5000.1000.1060 apicl(config-if-hsrp)# inherit hsrp group-policy hsrp-groupPoll </pre>	

Configuring HSRP in APIC Using REST API

HSRP is enabled when the leaf switch is configured.

Before you begin

- The tenant and VRF must be configured.
- VLAN pools must be configured with the appropriate VLAN range defined and the appropriate Layer 3 domain created and attached to the VLAN pool.
- The Attach Entity Profile must also be associated with the Layer 3 domain.

- The interface profile for the leaf switches must be configured as required.

Procedure

Step 1 Create port selectors.

Example:

```
<polUni>
  <infraInfra dn="uni/infra">
    <infraNodeP name="TenantNode_101">
      <infraLeafS name="leafselector" type="range">
        <infraNodeBlk name="nodeblk" from_"101" to_"101">
          </infraNodeBlk>
        </infraLeafS>
      <infraRsAccPortP tDn="uni/infra/accportprof-TenantPorts_101"/>
    </infraNodeP>
    <infraAccPortP name="TenantPorts_101">
      <infraHPortS name="portselector" type="range">
        <infraPortBlk name="portblk" fromCard="1" toCard="1" fromPort="41" toPort="41">
          </infraPortBlk>
        <infraRsAccBaseGrp tDn="uni/infra/funcprof/accportgrp-TenantPortGrp_101"/>
      </infraHPortS>
    </infraAccPortP>
    <infraFuncP>
      <infraAccPortGrp name="TenantPortGrp_101">
        <infraRsAttEntP tDn="uni/infra/attentp-AttEntityProfTenant"/>
        <infraRsHIfPol tnFabricHIfPolName="default"/>
      </infraAccPortGrp>
    </infraFuncP>
  </infraInfra>
</polUni>
```

Step 2 Create a tenant policy.

Example:

```
<polUni>
  <fvTenant name="t9" dn="uni/tn-t9" descr="">
    <fvCtx name="t9_ctx1" pcEnfPref="unenforced">
      </fvCtx>
    <fvBD name="t9_bd1" unkMacUcastAct="flood" arpFlood="yes">
      <fvRsCtx tnFvCtxName="t9_ctx1"/>
      <fvSubnet ip="101.9.1.1/24" scope="shared"/>
    </fvBD>
    <l3extOut dn="uni/tn-t9/out-l3extOut1" enforceRtctrl="export" name="l3extOut1">
      <l3extLNodeP name="Node101">
        <l3extRsNodeL3OutAtt rtrId="210.210.121.121" rtrIdLoopBack="no"
tDn="topology/pod-1/node-101"/>
      </l3extLNodeP>
      <l3extRsEctx tnFvCtxName="t9_ctx1"/>
      <l3extRsL3DomAtt tDn="uni/l3dom-dom1"/>
      <l3extInstP matchT="AtleastOne" name="extEpg" prio="unspecified"
targetDscp="unspecified">
        <l3extSubnet aggregate="" descr="" ip="176.21.21.21/21" name=""
scope="import-security"/>
      </l3extInstP>
    </l3extOut>
  </fvTenant>
</polUni>
```

Step 3 Create an HSRP interface policy.

Example:

```
<polUni>
  <fvTenant name="t9" dn="uni/tn-t9" descr="">
    <hsrpIfPol name="hsrpIfPol" ctrl="bfd" delay="4" reloadDelay="11"/>
  </fvTenant>
</polUni>
```

Step 4 Create an HSRP group policy.

Example:

```
<polUni>
  <fvTenant name="t9" dn="uni/tn-t9" descr="">
    <hsrpIfPol name="hsrpIfPol" ctrl="bfd" delay="4" reloadDelay="11"/>
  </fvTenant>
</polUni>
```

Step 5 Create an HSRP interface profile and an HSRP group profile.

Example:

```
<polUni>
  <fvTenant name="t9" dn="uni/tn-t9" descr="">
    <l3extOut dn="uni/tn-t9/out-l3extOut1" enforceRtctrl="export" name="l3extOut1">
      <l3extLNodeP name="Node101">
        <l3extLIIfP name="eth1-41-v6" ownerKey="" ownerTag="" tag="yellow-green">
          <hsrpIfP name="eth1-41-v6" version="v2">
            <hsrpRsIfPol tnHsrpIfPolName="hsrpIfPol"/>
            <hsrpGroupP descr="" name="HSRPV6-2" groupId="330" groupAf="ipv6" ip="fe80::3"
mac="00:00:0C:18:AC:01" ipObtainMode="admin">
              <hsrpRsGroupPol tnHsrpGroupPolName="G1"/>
            </hsrpGroupP>
          </hsrpIfP>
          <l3extRsPathL3OutAtt addr="2002::100/64" descr="" encap="unknown" encapScope="local"
ifInstT="l3-port" llAddr="::" mac="00:22:BD:F8:19:FF" mode="regular" mtu="inherit"
tDn="topology/pod-1/paths-101/pathep-[eth1/41]" targetDscp="unspecified">
            <l3extIp addr="2004::100/64"/>
          </l3extRsPathL3OutAtt>
        </l3extLIIfP>
        <l3extLIIfP name="eth1-41-v4" ownerKey="" ownerTag="" tag="yellow-green">
          <hsrpIfP name="eth1-41-v4" version="v1">
            <hsrpRsIfPol tnHsrpIfPolName="hsrpIfPol"/>
            <hsrpGroupP descr="" name="HSRPV4-2" groupId="51" groupAf="ipv4" ip="177.21.21.21"
mac="00:00:0C:18:AC:01" ipObtainMode="admin">
              <hsrpRsGroupPol tnHsrpGroupPolName="G1"/>
            </hsrpGroupP>
          </hsrpIfP>
          <l3extRsPathL3OutAtt addr="177.21.21.11/24" descr="" encap="unknown"
encapScope="local" ifInstT="l3-port" llAddr="::" mac="00:22:BD:F8:19:FF" mode="regular"
mtu="inherit" tDn="topology/pod-1/paths-101/pathep-[eth1/41]" targetDscp="unspecified">
            <l3extIp addr="177.21.23.11/24"/>
          </l3extRsPathL3OutAtt>
        </l3extLIIfP>
      </l3extLNodeP>
    </l3extOut>
  </fvTenant>
</polUni>
```

