Migration Steps

> ✎ **Note** It is recommended not to use VLAN 1 as control segment.

This chapter contains the following sections:

1) Upgrading and Configuring the Spine Switch Software

You must first upgrade all the spine switch software.

Before You Begin

The following prerequisites must be met before you upgrade the Cisco Nexus 6000 Series spine switch software:

• The Cisco Nexus 6000 Series switch must be running Cisco NX-OS Release 6.0(x) or later

> **Note** If you have anything other than a Cisco Nexus 6000 Series switch, you must physically replace the switch with Release 7.0(0)N1(1); the configuration remains the same as the previous image.

Step 1 On the Cisco Nexus 6000 Series spine switches, perform a nondisruptive in-service software upgrade (ISSU) upgrade to Cisco NX-OS Release 7.0(0)N1(1).
See the Cisco Nexus 6000 Series NX-OS Software Upgrade and Downgrade, Release 7.0 for instructions on performing the ISSU upgrade.

No impact to traffic should occur as a result of the ISSU upgrade.

Step 2 Add Cisco Dynamic Fabric Automation-specific configuration on the spine.
For more information, see Migration Configuration.

2) Upgrading the Border Leaf Software

You can perform a disruptive in-service software upgrade (ISSU) for the first border leaf pair.

Before You Begin

Before you perform an ISSU upgrade, move any existing switch virtual interface (SVI) configurations from other switches in the spine-leaf topology to the border-leaf pair nodes.

> **Note** You must perform configuration on the border-leaf nodes manually, not through power-on auto-provisioning (POAP).

Step 1 Upgrade the first border leaf node from Cisco NX-OS Release 6.0(2)N2 to Cisco NX-OS Release 7.0(0)N1(1) using an ISSU disruptive upgrade procedure. See the Cisco Nexus 6000 Series NX-OS Software Upgrade and Downgrade, Release 7.0 for information on performing an ISSU upgrade.

Step 2 Verify that the first border leaf comes up fully and becomes operational again.

Step 3 Verify that the traffic streams are already running intra-VLAN, inter-VLAN, across pods, and that north-bound traffic remains unaffected.

Step 4 Repeat Steps 1 to 3 for the second border leaf node in the pair and also for additional border leaf pairs.
Although the disruptive upgrade has some effect on traffic, traffic flow is not changed

3) Configuring the Border Leaf Pair

> ✎
>
> **Note** For specific configuration commands and examples, see Migration Configuration.

Step 1 On the first border leaf switch, do the following:

    a) Configure the Hot Standby Router Protocol (HSRP) per VLAN with the anycast gateway MAC address and an unused IP address.

    b) (Optional) route reflector (RR) normally runs on the spine but they can be configured to run on the leaf or border leaf.

    c) Configure the anycast gateway MAC address.

    d) Add a vrf-tenant-profile and configure the virtual network identifier (VNI) under the virtual routing and forwarding (VRF) instance.

    e) Create switch virtual interfaces (SVIs), if they are not present.

    f) Enable anycast-gateway on the SVIs.

    g) Enable traditional forwarding on the switch virtual interfaces (SVIs).

    h) Configure the BGP so that host routes are advertised to the BGP route reflector.

Step 2 Repeat Step 1 for the second border leaf switch in the pair and also for additional border leaf pairs.

4) Upgrading the FabricPath Leaf Pair

You can perform an in-service software upgrade (ISSU) for the FabricPath leaf pair.

Step 1 Upgrade the leaf node from Cisco NX-OS Release 6.0(2)N2 to Cisco NX-OS Release 7.0(0)N1(1) using an ISSU upgrade procedure. See the Cisco Nexus 6000 Series NX-OS Software Upgrade and Downgrade, Release 7.0. for information on performing an ISSU upgrade.

Step 2 Verify that the first leaf comes up fully and becomes operational again.

Step 3 Make sure traffic streams are already running intra-VLAN, inter-VLAN, across pods, and that north-bound traffic remains unaffected.

Step 4 Repeat Steps 1 to 3 for the second border leaf node in the pair.

There is no change in the traffic flow.

5) Adding DFA Configuration to FabricPath Leaf Pair

You can configure the FabricPath leaf nodes in the network.

Note   For specific configuration commands and examples, see Migration Configuration.

Before You Begin

Before you configure the FabricPath leaf, you should upgrade the software.

Step 1   On the first switch in the pair, do the following:

a)   Install a Layer-3 license.

b)   Enable Cisco DFA.

c)   Configure the iBGP route reflector client.

d)   Add the segment ID and virtual routing and forwarding (VRF) instance.

e)   Add a vrf-tenant-profile and configure the virtual network identifier (VNI) under the virtual routing and forwarding (VRF) instance.

f)   Create switch virtual interfaces (SVIs) for all VLANs.

g)   Enable anycast-gateway mode for all VLANs.

Step 2   Repeat Step 1 for the second border leaf switch in the pair.

If you are migrating a fabric that includes both Cisco Nexus 5000 Series and Cisco Nexus 6000 Series Switches, do the following:

- Migration is completed if you have upgraded all Cisco Nexus 6000 Series software and enabled Cisco DFA forwarding.

- HSRP/VRRP remains if there are Cisco Nexus 5000 Series leaf nodes in the network.

- In Cisco Nexus 5000 Series-involved VLANs and SVIs, the VLANs are global, non-segment-ID-enabled, and the forwarding mode is anycast gateway mode.

- In upgraded and only Cisco Nexus 6000 Series-involved VLANs and SVIs, the VLANs and SVIs can be segment ID enabled, and the forwarding mode can be either proxy or anycast gateway mode.

- Multicast continues to run in the legacy multicast mode. Cisco DFA multicast should not be enabled.

What to Do Next

If you are migrating a fabric that includes both Cisco Nexus 5000 Series and Cisco Nexus 6000 Series Switches, the migration is completed if you have upgraded all Cisco Nexus 6000 Series software and enabled Cisco DFA forwarding.

6) Upgrading and Configuring All Remaining Leaf Switches

You should perform this procedure on all of the remaining leaf switches in the network.

Note For specific configuration commands and examples, see Migration Configuration.

Step 1 Upgrade the software on all of the remaining leaf switches.

Step 2 Add Cisco DFA-related configuration on all of the remaining leaf switches.

Step 3 Enable anycast-gateway mode on leaf switches for all VLANs.

7) Removing the HSRP Configurations on Border Leaf Pairs

Note This step is performed only when there are no Cisco Nexus 5000 Series leaf nodes in the Cisco Dynamic Fabric (DFA) topology.

During the migration, some hosts learn the anycast gateway MAC address as its MAC address for the default gateway. Some hosts learn the HSRP Virtual Mac Address (VMAC) as the MAC address for the default gateway. We recommend that you wait a couple of hours to make sure that the HSRP VMAC address is aged out on all hosts.

Note For specific configuration commands and examples, see Migration Configuration.

Before You Begin

You should have completed migration on all leaf switches.

Step 1 Remove the HSRP configuration on each border leaf switch.

Step 2 Change the SVI IP address to the Virtual IP (VIP) address.

After you remove the HSRP configurations, migration is complete.

Note
- If there are two or more HSRP groups, then there will be more than one HSRP VIP configured on the attached host. But only one VIP can be used as the DFA anycast gateway IP. This requires changes on the attached hosts to use the selected VIP as default gateway.

- Only one v4 or v6 HSRP dummy group is needed per SVI.

- You can move to Cisco DFA multicast, if preferred. There is no vPC on the border leaf connecting to external multicast routers.

- SVIs can be switched to proxy forwarding mode, if preferred.

- New VLANs can be segment ID enabled.

- In an all Cisco Nexus 6000 Series topology, you can move to Cisco DFA multicast, if preferred.