



Configuring a Forwarding Scale Profile Policy

- [Forwarding Scale Profile Policy Overview, on page 1](#)
- [Supported Platforms, on page 3](#)
- [Guidelines and Limitations, on page 4](#)
- [Configuring the Forwarding Scale Profile Policy Using the REST API, on page 5](#)

Forwarding Scale Profile Policy Overview

The Forwarding Scale Profile policy provides different scalability options. For example:

- **Dual Stack**—provides scalability of up to 12,000 endpoints for IPv6 configurations and up to 24,000 endpoints for IPv4 configurations.
- **High LPM**—provides scalability similar to the dual-stack profile, except that the longest prefix match (LPM) scale is 128,000 and the policy scale is 8,000.
- **IPv4 Scale**—enables systems with no IPv6 configurations to increase scalability to 48,000 IPv4 endpoints.
- **High Dual Stack**—provides scalability of up to 64,000 MAC endpoints and 64,000 IPv4 endpoints. IPv6 endpoint scale can be 24,000/48,000, depending on the switch hardware model.



Note With Cisco Application Policy Infrastructure Controller (APIC) release 3.2(1), depending on your TOR switch hardware, a Forwarding Scale Profile with the **High Dual Stack** option has different scales; for example:

- For Cisco Nexus 9000 Series TOR switches with FX in the switch name, the high dual-stack option has scalability of 48,000 IPv6 endpoints instead of 24,000 and 128,000 policies instead of 8,000.
- For Cisco Nexus 9000 Series TOR switches with EX in the switch name, the high dual-stack option has the same scale values as with earlier Cisco APIC releases.

See the following table for more details.

Table 1: Forwarding Scale Profile Policy Scalability

Forwarding Scale Profile Policy Options	TOR Switches with EX Names	TOR Switches with FX Names
Dual Stack	<ul style="list-style-type: none"> • EP MAC: 24,000 • EP IPv4: 24,000 • EP IPv6: 12,000 • LPM: 20,000 • Policy: 64,000 • Multicast: 8,000 	Has the same scalability numbers as Dual Stack scale on earlier switches.
High Dual Stack	<ul style="list-style-type: none"> • EP MAC: 64,000 • EP IPv4: 64,000 • EP IPv6: 24,000 • LPM: 38,000 • Policy: 8,000 • Multicast: 0 	<ul style="list-style-type: none"> • EP MAC: 64,000 • EP IPv4: 64,000 • EP IPv6: 48,000 • LPM: 38,000 • Policy: 128,000 • Multicast: 512
High LPM	Provides scalability similar to the dual-stack profile, except that the longest prefix match (LPM) scale is 128,000 and the policy scale is 8,000.	Has the same scalability numbers as on earlier switches.
IPv4 Scale	<ul style="list-style-type: none"> • EP MAC: 48,000 • EP IPv4: 48,000 • EP IPv6: 0 • LPM: 38,000 • Policy: 64,000 • Multicast: 8,000 	Has the same scalability numbers as IPv4 scale on earlier switches.

**Note**

- Because the IPv4 forwarding scale profile policy does not support IPv6 configurations, all IPv6 configurations must be removed from switches configured with the IPv4 forwarding scale profile policy.
- Because the high dual stack profile has reduced-scale support for contract policies (8,000), the contracts scale must be reduced accordingly prior to deploying that profile.
- Before migrating to minimal tenant multicast scale leaf profiles, such as high dual stack, we recommend that you first disable Layer 2 IGMP snooping-, Layer 3 IGMP-, and PIM-related configurations to prevent having a stale multicast state in your hardware.
- Applying a scale profile to a node requires a manual reload of that node. Any unsupported switches are ignored. For a list of supported switches, see the supported platforms section.
- VPCs associated with different scale profile settings are not supported. The VPC members must be configured with the same scale profile settings.

Supported Platforms

This section provides forwarding scale profiles hardware support information for Release 4.1(1).

The following table summarizes platform support for each forwarding scale profile.

Table 2: Supported Switches

Scale Profile	Supported on EX and FX2 Switches	Supported on FX Switches
Dual Stack	Yes	Yes
High Dual Stack	Yes	Yes Different scale numbers than EX/FX2.
High LPM	Yes	Yes
IPv4 Scale	Yes	Yes

- Forwarding scale profiles are supported on the following switches:
 - Cisco Nexus 9300-EX series switches
 - Cisco Nexus 9300-FX series switches
 - Cisco Nexus 9300-FX2 series switches
- Switches not listed here do not support forwarding scale profiles in this release.

Guidelines and Limitations

- When downgrading to a release that does not support one or more switches in your current fabric, keep the following in mind:
 - If you downgrade your fabric to a release where one or more of your current switches are not supported, those switches will become inactive in the fabric.
 - If you later upgrade the fabric to a release where the switch is supported again, the APIC will not regain complete details about the switch. In this case, you will need to explicitly remove the switch from the APIC and then re-add it to the fabric.
- When downgrading to a release that does not support one or more of your current forwarding scale profiles, the default forwarding scale profile will be configured on the switch. You must reduce the configurations on the switch to fit the default profile before the upgrade.
- Because the IPv4 Scale forwarding scale profile does not support IPv6 configurations, you must remove all IPv6 configurations from the switches that need to be configured with the IPv4 Scale profile.
- Before switching between forwarding scale profiles, the configurations on the switch must be reduced appropriately and thoroughly verified so that scale parameters of the target profile are not exceeded.

For example, for switch models with EX at the end of the switch name, because the High Dual Stack profile has reduced scale support for contract policies, you must reduce the contracts scale accordingly before deploying that profile.
- Before migrating to minimal tenant multicast scale leaf profiles, such as High Dual Stack, we recommend that you first disable Layer 2 IGMP snooping, Layer 3 IGMP, and PIM-related configurations to prevent having a stale multicast state in your hardware.
- Applying a forwarding scale profile to a node requires a manual reload of that node. Any unsupported switches are ignored.
- vPCs associated with different forwarding scale profile settings are not supported. You must configure the vPC members with the same profile settings.
- With the default deny model in Cisco ACI, the configured tenant or VRFs have implicit rules that consume several TCAM entries for each VRF. With an increase in the number of VRFs configured on a single switch, these TCAM entries that are used per VRF also count toward the overall policy TCAM usage.
- Beginning with Release 4.1(1), the policy count updates that are reported by the leaf switches to the Cisco APIC through the MO `actrlRuleHit5min` is set to 0 for the High Dual Stack profile for all platforms.
- Beginning with Release 4.2(1), the policy count updates that are reported by the leaf switches to the Cisco APIC through the MO `actrlRuleHit5min` is set to 0 for the High Policy profile on the Cisco Nexus 93180YC-FX switch.
- Beginning with Release 4.2(2), the policy count updates that are reported by the leaf switches to the Cisco APIC through the MO `actrlRuleHit5min` is set to 0 for the High Policy profile on the Cisco Nexus 93600CD-GX switch.
- Beginning with Release 4.2(31), the policy count updates that are reported by the leaf switches to the Cisco APIC through the MO `actrlRuleHit5min` is set to 0 for the High Policy profile on the Cisco Nexus 9364C-GX switch.

- If you need to clear the configurations on the switch, we recommend using the `setup-clean-config.sh -k` command. The command will clear all configurations on the switch, except the forwarding scale profile and port profile configurations.

Configuring the Forwarding Scale Profile Policy Using the REST API

The forwarding scale profile policy requires supported switches. For a list of supported switches, see the supported platforms section. The switches that support the forwarding scale profile policy must be manually reloaded after the forwarding scale profile policy is applied.

The Forwarding Scale Profile policy provides different scalability options. For more information on the scalability options, see [Forwarding Scale Profile Policy Overview, on page 1](#).

This section explains how to create a forwarding scale profile policy and apply it to a leaf profile using the REST API.

To apply a forwarding scale profile policy with IPv4 scaling, send a post with XML similar to the following example:

Example:

```
<polUni>
  <infraInfra>
    <topoctrlFwdScaleProfilePol name="sampleFwdScaleProf" profType="ipv4"/>
    <infraAccNodePGrp name="sampleNodePolGrp">
      <infraRsTopoctrlFwdScaleProfilePol tnTopoctrlFwdScaleProfilePolName="sampleFwdScaleProf"/>
    </infraAccNodePGrp>
    <infraNodeP name="nodeProf_101">
      <infraLeafS name="leafS_101" type="range">
        <infraNodeBlk name="test" from_"101" to_"101"/>
        <infraRsAccNodePGrp tDn="uni/infra/funcprof/accnodepgrp-sampleNodePolGrp "/>
      </infraLeafS>
    </infraNodeP>
  </infraInfra>
</polUni>
```

To take effect, the switches that support the forwarding scale profile policy must be manually reloaded after the forwarding scale profile policy is applied.

