



Verified Scalability for Cisco Nexus 6000 Series NX-OS Release 7.0(2)N1(1)

This chapter contains the following sections:

- [Overview of Verified Scalability, page 1](#)
- [Verified Scalability for a Layer 2 Switching Deployment, page 2](#)
- [Verified Scalability for a Layer 2 Switching and Layer 3 Routing Deployment, page 3](#)
- [Verified Scalability for a Layer 3 Aggregation Routing Deployment, page 6](#)
- [Verified Scalability for a Layer 2 Switching and Unified Fabric \(FCoE\) Deployment, page 8](#)
- [Verified Scalability for Multicast Routing, page 10](#)
- [Verified Scalability for Unicast Routing, page 10](#)

Overview of Verified Scalability

This document lists the Cisco verified scalability limits.



Note

The following scaling capabilities apply to Cisco Nexus 6000 Series switches. For example, 6001 and 6004.

In the following tables, the Verified Topology column lists the verified scaling capabilities with all listed features enabled at the same time. The numbers listed here exceed those used by most customers in their topologies. The scale numbers listed here are not the maximum verified values if each feature is viewed in isolation.

The Verified Maximum column lists the maximum scale capability tested for the corresponding feature individually. This number is the absolute maximum currently supported by the Cisco NX-OS Release software for the corresponding feature. If the hardware is capable of a higher scale, future software releases may increase this verified maximum limit.

Verified Scalability for a Layer 2 Switching Deployment

This table lists the verified scalability for a Layer 2 switching deployment.

Table 1: Verified Scalability for a Layer 2 Switching Deployment

Feature	Verified Topology	Verified Maximum
Active VLANs/VSANs per switch	4000	4013 (31 are reserved for VSANs, and the remaining are for VLANs.)
VLAN/VSAN ID Space	4013	4013 Unreserved Space
Logical Interfaces ¹	96,000 (MST) ²	64,000 (96,000 MST) ³
VLAN ACLs (VACLs)	128 (10 Unique VACLs)	1024 (512 unique VACLs with up to 1024 ACE entries across all VACLs)
Maximum interfaces per EtherChannel	16	16
IGMP Snooping Groups	4000 (in FEX deployments) 8000 (in non-FEX deployments)	4000 (in FEX deployments) 16,000 (in non-FEX deployments)
Maximum FEXs per Switch	48 ⁴	48 ⁵
Maximum FEXs dual-homed to a vPC Switch Pair	48	48
MAC Table Size (Entries)	64,000	115,000
Number of Switchport Etherchannels	48	96 (Single member port-channel for 40G ports) 384 (Single member port-channel for 10G ports) 64 (Multi member port-channel)
Number of HIF FEX port channels/vPCs (across the maximum number of FEXs)	576	576
SPAN Sessions	4 active sessions 32 source VLANs as a RX source	16 active sessions 32 source VLANs as a RX source

Feature	Verified Topology	Verified Maximum
SVIs	2	256
FabricPath VLANs	4000 ⁶	4000
FabricPath Switch IDs	500	500
FabricPath Multicast Trees	1	1
Number of FabricPath Topologies	2	2
Number of FabricPath Core Port-Channels	4 core links with 4 ports each	16
FEX Host Interface Storm Control	1936 ⁷	1936
Segmentation ID	3000 (1000 global segments, 2000 local segments)	3000 (1000 global segments, 2000 local segments)

- ¹ Logical interfaces are a product of the number of VLANs times the number of ports. This parameter reflects the load of handling port programming, and is not dependent on the spanning-tree mode or configuration.
- ² 32,000 STP logical interfaces are verified in the unified fabric topology.
- ³ 16,000 Port-VLAN scaling number applies to Rapid PVST+ and non-STP modes. For MST mode only, the maximum verified limit for Port-VLAN scaling is 96,000. 64000 limit is verified when a switch is running in MST mode and performing Rapid PVST+ stimulation. 48,000 is for MST and 16,000 is for Rapid PVST+.
- ⁴ 48 maximum FEXs per switch applies to the Cisco Nexus 6004 only. The limit is 24 maximum FEXs per switch for the Cisco Nexus 6001.
- ⁵ 48 maximum FEXs per switch applies to the Cisco Nexus 6004 only. The limit is 24 maximum FEXs per switch for the Cisco Nexus 6001.
- ⁶ FabricPath VLANs are verified in the unified fabric topology
- ⁷ This is the target maximum number that HIF-SC can support. Beyond this number, NIF-SC is recommended for deployment.

Verified Scalability for a Layer 2 Switching and Layer 3 Routing Deployment

This table contains the verified scalability for a Layer 2 switching and Layer 3 routing deployment.



Note

The currently tested values do not provide an indication for the maximum scalability of the control plane. These numbers vary based on the load of the system in terms of routing protocols, timers settings, and other values. Proof of concept testing should be used to determine the scalability of a given feature for your environment.

Table 2: Verified Scalability for a Layer 2 Switching and Layer 3 Routing Deployment

Feature	Verified Topology	Verified Maximum
Active VLANs/VSANs per Switch	1000 8	4013 (31 are set reserved for VSANs and the remaining are for VLANs)
VLAN/VSAN ID Space	4013 Unreserved space	4013 Unreserved space
STP Instances	16,000	16,000
Maximum Interfaces per EtherChannel	16	16
IGMP Snooping Groups	4000 (in FEX deployments) 8000 (in non-FEX deployments)	4000 (in FEX deployments) 16,000 (in non-FEX deployments)
Maximum FEXs per Switch	24	24
Maximum FEXs Dual-homed to a vPC Switch Pair 9	24 10	24 11
MAC Table Size (Entries)	32,000 12	64,000 13
Number of FEX Port Channels/vPCs (across the maximum number of FEXs)	500	768
SPAN Sessions	4 active sessions 32 source VLANs as an RX source	16 active sessions 32 source VLANs as an RX source
Number of SVIs	256	256
Dynamic IPv4 Routes 14	8000	24,000
Dynamic IPv6 Routes 15	4000 16	6000
Multicast IPv4 Routes 17	4000	8000
ARPs (IPv4 Hosts) 18	32,000	64,000

Feature	Verified Topology	Verified Maximum
IPv6 Hosts	16,000	32,000
VRFs	25	1000
RACLs	64 Ingress RACLs with up to 1152 ACE entries across all the RACLs	896 Ingress RACLs with up to 1152 ACE entries across all the RACLs
HSRP Groups 19	256	500
VRRP Groups 20	256	500
BFD sessions over L3-intf for CE mode	8 sessions (250ms intvl, 750ms dead-intvl)	30 (250ms intvl, 750ms dead-intvl)
BFD sessions over SVI for L2MP mode	64 (250ms intvl, 750ms dead-intvl)	64 (250ms intvl, 750ms dead-intvl)
PBR IPv4	95	95
PBR IPv6	95	95

- 8 4,013 VLANs are verified in Layer 2 switching, Fibre Channel, and FCoE topologies.
- 9 FEXs are verified in the Layer 2 topology
- 10 24 FEXs are verified in the Layer 2 topology.
- 11 24 FEXs are verified in the Layer 2 topology.
- 12 128,000 entries are reserved for unicast MAC entries and 128,000 entries are reserved for IP host routes.
- 13 128,000 entries are reserved for unicast MAC entries and 128,000 entries are reserved for IP host routes.
- 14 The maximum number of entries that can be supported is 24,000. This table is shared between IPv4 and IPv6. An IPv4 route takes up one entry in the table and an IPv6 route takes up four entries.
- 15 The maximum number of entries that can be supported is 24,000. This table is shared between IPv4 and IPv6. An IPv4 route takes up one entry in the table and an IPv6 route takes up four entries.
- 16 With no IPv4, the number increases to 6000.
- 17 This includes (*,G) entries, (S,G) entries, and the entries required for vPC with bind-vrf configured. When bind-vrf is configured, each (*,G) and (S,G) entry is replicated.
- 18 The maximum number of hosts supported is listed under ARPs. This includes IPv4 and IPv6 hosts. IPv4 hosts take up one entry and IPv6 hosts take up 2 entries in hardware. So, for the verified limit, the switch supports one of the following: 64,000 IPv4 hosts and 0 IPv6 hosts, 32,000 IPv6 hosts and 0 IPv4 hosts, or a combination of IPv4 and IPv6 hosts.
- 19 The limit of the table that holds the Router MAC and Virtual MAC entries that determines whether the packet needs to be bridged or routed is 500 entries. The Virtual MAC entries can be shared across Layer 3 interfaces. So, we recommend that you configure the same group ID across all or multiple Layer 3 interfaces/SVIs. If multiple group IDs are configured on an Layer 3 interface, then we recommend that you configure the same set of group IDs across all or multiple Layer 3 interfaces. This way, HSRP/VRRP can be supported on more interfaces. Please refer to the unicast configuration guide for more information.
- 20 The limit of the table that holds the Router MAC and Virtual MAC entries for determining packet routing or switching is 500 entries. The Virtual MAC entries can be shared across Layer 3 interfaces. So, we recommend that you configure the same group ID across all or multiple Layer 3 interfaces/SVIs. If multiple group IDs are configured on an Layer 3 interface, then we recommend that you configure the same set of group IDs across all or multiple Layer 3 interfaces. This way, HSRP/VRRP can be supported on more interfaces. Please refer to the unicast configuration guide for more information.

Verified Scalability for a Layer 3 Aggregation Routing Deployment

This table lists the verified scalability for a Layer 3 aggregation routing deployment.

Table 3: Verified Scalability for a Layer 3 Aggregation Routing Deployment

Feature	Verified Topology	Verified Maximum
Active VLANs/VSANs per Switch	4000 21	4000
VLAN/VSAN ID Space	4013 unreserved space	4013 unreserved space
STP Instances	64,000	64,000
Maximum Interfaces per EtherChannel	16	16
IGMP Snooping Groups	8000 (in non-FEX deployments)	16,000 (in non-FEX deployments)
MAC Table Size	64,000 22	115,000 23
SPAN Sessions	4 active sessions 32 source VLANs as an RX source	16 active sessions 32 source VLANs as an RX source
SVIs	1000	4000
Dynamic IPv4 Routes 24	8000	24,000
Dynamic IPv6 Routes 25	4000 26	6000
Multicast IPv4 Routes 27	8000	16,000
RACLs	64 ingress RACLs with up to 1,152 ACE entries across all of the RACLs.	512 ingress RACLs with up to 1,152 ACE entries across all of the RACLs.
VRFs	25	1,000
ARPs (IPv4 Hosts) 28	64,000	64,000

Feature	Verified Topology	Verified Maximum
IPv6 Hosts 29	10,000 30	32,000
IGP Peers	64	100
HSRP Groups 31	1721 (911 IPv4, 810 IPv6)	1721 (911 IPv4, 810 IPv6)
VRRP Groups 32	500	500
FabricPath Switch IDs	500	500
FabricPath Multicast Trees	2	2
Number of FabricPath Topologies	2	2
Number of FabricPath Core Links	32	32
PBR IPv4	110	110
PBR IPv6	110	110
BFD sessions over L3-intf for CE Mode	100 sessions (250ms intvl, 750ms dead-intvl)	100 sessions (250ms intvl, 750ms dead-intvl)
BFD sessions over SVI for L2MP mode	64 sessions (250ms intvl, 750ms dead-intvl)	64 sessions (250ms intvl, 750ms dead-intvl)

- 21 4,013 VLANs are verified in Layer 2 switching, Fibre Channel, and FCoE topologies.
- 22 128,000 entries are reserved for Unicast MAC entries, and 128,000 entries are reserved for IP host routes.
- 23 128,000 entries are reserved for Unicast MAC entries, and 128,000 entries are reserved for IP host routes.
- 24 The maximum number of entries that can be supported is 24,000. This table is shared between IPv4 and IPv6. An IPv4 route takes up one entry in the table and an IPv6 route takes up four entries.
- 25 The maximum number of entries that can be supported is 24,000. This table is shared between IPv4 and IPv6. An IPv4 route takes up one entry in the table and an IPv6 route takes up four entries.
- 26 Entries shared between IPv4, IPv6 network routes.
- 27 This includes (*,G) entries, (S,G) entries, and the entries required for vPC with bind-vrf configured. When bind-vrf is configured, each (*,G) and (S,G) entry is replicated. This includes (*,G) entries, (S,G) entries, and IGMP-snooping entries combined.
- 28 The maximum number of hosts supported is listed under ARPs. This includes IPv4 and IPv6 hosts. IPv4 hosts take up one entry and IPv6 hosts take up 2 entries in hardware. So, for the verified limit, the switch supports one of the following: 64,000 IPv4 hosts and 0 IPv6 hosts, 16,000 IPv6 hosts and 0 IPv4 hosts, or a combination of IPv4 and IPv6 hosts.
- 29 The maximum number of hosts supported is listed under ARPs. This includes IPv4 and IPv6 hosts. IPv4 hosts take up one entry and IPv6 hosts take up 2 entries in hardware. So, for the verified limit, the switch supports one of the following: 64,000 IPv4 hosts and 0 IPv6 hosts, 16,000 IPv6 hosts and 0 IPv4 hosts, or a combination of IPv4 and IPv6 hosts.
- 30 Entries shared between IPv4 multicast, IPv4, IPv6 host routes .
- 31 The limit of the table that holds the Router MAC and Virtual MAC entries that determine whether the packet needs to be bridged or routed is 500 entries. The Virtual MAC entries can be shared across Layer 3 interfaces. So we recommend that you configure the same group ID across all or multiple Layer 3 interfaces/SVIs. If multiple group IDs are configured on a Layer 3 interface, then we recommend that you configure the same set of group IDs across all or multiple Layer 3 interfaces. This way, HSRP/VRRP can be supported on more interfaces. Please refer to the Unicast Routing Configuration Guide for more information.

- ³² The limit of the table that holds the Router MAC and Virtual MAC entries for determining packet routing or switching is 500 entries. The Virtual MAC entries can be shared across Layer 3 interfaces. So we recommend that you configure the same group ID across all or multiple Layer 3 interfaces/SVIs. If multiple group IDs are configured on a Layer 3 interface, then we recommend that you configure the same set of group IDs across all or multiple Layer 3 interfaces. This way, HSRP/VRRP can be supported on more interfaces. Please refer to the Unicast Routing Configuration Guide for more information.

Verified Scalability for a Layer 2 Switching and Unified Fabric (FCoE) Deployment

This table lists the verified scalability for a Layer 2 switching and unified fabric (FCoE) deployment.

Table 4: Verified Scalability for a Layer 2 Switching and Unified Fabric (FCoE) Deployment

Feature	Verified Topology	Verified Maximum
Active VLANs/VSANs per switch	4,013	4,013 (31 are set reserved for VSANs and the remaining are for VLANs)
VLAN/VSAN ID Space	4,013 Unreserved Space	4,013 Unreserved Space
Logical Interfaces 33	32,000	32,000
IGMP Snooping Groups	4,000 (in FEX deployments) 8,000 (in non-FEX deployments)	4,000 (in FEX deployments) 16,000 (in non-FEX deployments)
Maximum FEXs per Cisco Nexus 6000 Series Switch	17 per DUT	24
Maximum FEXs Dual Homed to a vPC Cisco Nexus 6000 Series Switch Pair	34	34
MAC Table Size	32,000 unicast entries 34	115,000 entries 35
Number of Switchport EtherChannels	8	384
SPAN Sessions	4 active sessions 32 source VLANs as a RX source	16 active sessions 32 source VLANs as a RX source
Number of FEX Port Channels/vPCs (across the maximum number of FEXs)	544	768
FabricPath VLANs	4,000	4,000

Feature	Verified Topology	Verified Maximum
FabricPath Switch IDs	500	500
FabricPath Multicast Trees	2	2
Number of FabricPath Topologies	2	2
Number of FabricPath Core Links	2	16
FLOGIs or FDISCs per NPV Port Group	512	512
Zone Sets per Switch	32	500
Zone Members per Physical Fabric (includes all VSANs)	16,000	16,000
Zones per Switch (includes all VSANs)	8000	8000
Maximum Diameter of a SAN Fabric	7 hops	12
FSPF Interface Instances per Switch	192	3,072
ISL Instances per Switch	6	96
VFC Interfaces	544	544
Maximum FCIDs Allocated	1024	4000
Fibre Channel Flows	32	32
FLOGIs or FDISCs per switch	1024	4000
The maximum number of vFCs that can be bound to a port-channel	24 for the Nexus 6001 and 48 for the other(s)	24 for the Nexus 6001 and 48 for the other(s)

³³ Logical interfaces are a product of the number of VLANs times the number of ports. This parameter reflects the load of handling port programming, and is not dependent on the spanning-tree mode or configuration.

³⁴ 128,000 entries are reserved for unicast MAC entries and 128,000 entries are reserved for IP host routes.

³⁵ 128,000 entries are reserved for unicast MAC entries and 128,000 entries are reserved for IP host routes.

Verified Scalability for Multicast Routing

This table lists the verified scalability for multicast routing.

Table 5: Verified Scalability for Multicast Routing

Feature	Parameter	Verified Maximum
Protocol Independent Multicast (PIM)	Number of neighbors	500
	Number of neighbors/total routes per system with aggressive hello timers (5 seconds)	16/4,000
Multicast Source Discovery Protocol (MSDP)	Number of MSDP Source-Active (SA) cache entries	6,000



Note

- In vPC setup, TCAM exhaustion failure will lead to some routes not getting programmed in the hardware. Hence, there might exist a condition where mrib will show the route exists but mfib may not have it programmed.
- In bind-vrf configuration, for every mroute, additional mroute is programmed in the hardware and this could lead to TCAM exhaustion. Hence, ensure that the mroute count does not exceed $((\text{max-limit}/2) - 4)$ default routes.

For example: If the hardware profile multicast max-limit is 8000, then mroute count should not exceed $((8000/2)-4)$ default routes.

Verified Scalability for Unicast Routing

Guidelines and Limitations for Unicast Routing

- You can have up to four instances of OSPFv2.
- You can have up to four instances of OSPFv3.

This table lists the verified scalability for unicast routing.

Table 6: Verified Scalability for Unicast Routing

Feature	Parameter	Verified Maximum
OSPFv2	Number of active interfaces	256
	Number of passive interfaces	256
	Number of neighbors/total routes with aggressive timers (1 sec/ 3 sec)	16/6,000
OSPFv3	Number of active interfaces	256
	Number of passive interfaces	256
EIGRP	Number of active interfaces	50
BGP	Number of peers (iBGP and eBGP, active)	256
	Number of AS path entries	512
	Number of prefix-list entries in a single prefix-list	10,000
HSRP	Number of groups with aggressive timers (1 sec/3 sec)	500
L3 ISIS	Number of adjacencies	100
Unicast Adjacencies	Number of regular adjacencies	16374
	Number of ECMP adjacencies. The regular adjacencies are also part of the maximum ECMP adjacencies	28470

