Cisco N9300 Series Smart Switches NX-OS Verified Scalability Guide, Release 10.6(1s)

Revised: September 16, 2025,

Introduction

This document describes the Cisco NX-OS configuration limits for Cisco N9300 Series smart switches.

The values provided in this guide should not be interpreted as theoretical system limits for Cisco NX-OS hardware or Cisco NX-OS software. These limits refer to values that have been validated by Cisco. They can increase over time as more testing and validation is done.

Verified Scalability Limits - Unidimensional

The following tables in this section list the verified scalability limits of the Cisco N9324C-SE1U, Cisco N9348Y2C6D-SE1U switches for Cisco NX-OS Release 10.6(1s).

These limits are validated with a unidimensional configuration. The values provided in these tables focus on the scalability of one particular feature at a time.

Each number is the absolute maximum that is currently supported by this Cisco NX-OS release for the corresponding feature. If the hardware is capable of a higher scale, future software releases could increase this verified maximum limit. Results might differ from the values that are listed in this guide when you try to achieve maximum scalability with multiple features enabled.

Note

1. If only one number is provided, the verified limit applies to all supported platforms and line cards.
2. Verified limits are provided only for supported platforms.
3. If a feature is not supported for a particular platform, the verified limit is not provided.

Note You can deploy up to 500 commands under config-profile.

Interfaces Verified Scalability Limits

Feature	Supported Platforms	Verified Limits
IP DHCP relay addresses (helper addresses) per interface	Cisco N9324C-SE1U, Cisco N9348Y2C6D-SE1U switches	32 (IPv4) + 32 (IPv6)
Port channel links	Cisco N9324C-SE1U switch	96
	Cisco N9348Y2C6D-SE1U switch	80
SVIs	Cisco N9324C-SE1U, Cisco N9348Y2C6D-SE1U switches	1000 (with HSRP) 1000 HSRP groups
vPCs	Cisco N9324C-SE1U switch	96
	Cisco N9348Y2C6D-SE1U switch	80
Sub-interfaces	Cisco N9324C-SE1U, Cisco N9348Y2C6D-SE1U switches	1000

Note

- The scale for vPC convergence is tested with an LACP supported device connected to the vPC pair link.
- For **interface port-channel** configuration, LACP vPC convergence must be enabled. For information on vPC, see Configuring vPCs.
- For **vPC domain** configuration, the tested time for the delay restore is 150 seconds, the delay restore for interface VLAN is 150 seconds, and the delay restore for orphan ports is 140 seconds. For information on vPC, see Configuring vPCs.

Label Switching Verified Scalability Limits

Feature	Verified Limits
Layer 3 VPN routes	400,000 (IPv4 routes) 90,000 (IPv6 routes)
Hierarchical ECMP [1]	
Node SID	4000
Adjacency SID	8
VRF	1000
VPN label	1000

Feature	Verified Limits
Level 1 ECMP groups	12
Level 2 ECMP groups	10
Level 2 ECMP members	8
VPN decap statistics	1000 VRF

[1] Hierarchical ECMP provides enhanced ECMP scale and convergence, with two level route resolution.

Layer 2 Switching Verified Scalability Limits

Feature	Verified Limits
MAC addresses	120,000 (including control plane MACs)
MST instances	64
RPVST logical ports (logical ports * VLANs)	22,000
VLANs in MST mode	3967 (the remaining 127 VLANs are reserved)
VLANs in RPVST mode	3967

Multicast Routing Verified Scalability Limits

Feature	Verified Limits
IPv4 multicast routes **Note** The limits are for a combination of IPv4 and IPv6 multicast routes. Layer 2 multicast entries are a part of the total 120K limits. For example, 110K IPv4 + 2K IPv6 multicast routes + 8K Layer 2 multicast entries.	32,768 (Layer 3)
IGMP snooping groups	16,000

Note The IPv4 multicast routes and the IPv4/IPv6 host routes share the same hardware table. Limits are provided for both the default line card mode and the max host line card mode.

IP Fabric for Media Solution Policer Verified Scalability Limits

Feature	Supported Platforms
NBM Flow Policers (Slice/System)	Cisco N9324C-SE1U, Cisco N9348Y2C6D-SE1U switches

Programmability Verified Scalability Limits

Feature	Verified Limits
DME	
Note • If all the DME features are configured together, it may cause performance issues. • Model to CLI Conversion of payload is not supported.	
Telemetry	4 telemetry receivers can be streamed in parallel

QoS Verified Scalability Limits

Feature	Verified Limits
Class maps per policy map	128

Security Verified Scalability Limits

Feature	Verified Limits
Access Control Lists (ACL)	• Ingress RACLv4 - 7168 Access Control Entries (ACE) per slice • Ingress RACLv6 - 3584 ACE per slice • Egress RACLv4 - 7168 ACE per slice • Egress RACLv6 - 3584 ACE per slice • Ingress PACLv4 - 7168 ACE per slice • Ingress PACLv6 - 3584 ACE per slice • Ingress SVI RACLv4 - 7168 ACE per slice • Ingress SVI RACLv6 - 3584 ACE per slice • Egress SVI RACLv4 - 7168 ACE per slice • Egress SVI RACLv6 - 3584 ACE per slice
ACL Labels	• Ingress IPv4 - 126 unique labels • Ingress IPv6 - 126 unique labels • Egress IPv4 - 14 unique labels • Egress IPv6 - 14 unique labels
Ingress SUP IPv4 TCAM entries	360
Ingress SUP IPv6 TCAM entries	180

Note

- For Cisco N9324C-SE1U, Cisco N9348Y2C6D-SE1U switches, ACL scalability is defined per slice. Each ASIC typically contains two slices. To understand how interfaces are organized across these slices, use the **show interface hardware-mappings** command.
- The TCAM entries scalability limits also apply to policy-based TCAM entries (PBACLs).
- Only 62 unique ACLs can be configured. Each ACL takes one label. If the same ACL is configured on multiple interfaces, the same label is shared. If each ACL has unique entries, the ACL labels are not shared, and the label limit is 62.

System Management Verified Scalability Limits

Feature	Verified Limits
sFlow	
sFlow ports	64
SPAN and ERSPAN	
Active SPAN or ERSPAN sessions[2]	10
Destination interfaces per SPAN session	1 Physical only (no PO support).

[2] A single forwarding engine instance supports four SPAN or ERSPAN sessions. For Cisco Nexus 9300 Series switches, if the first three sessions have bidirectional sources, the fourth session has hardware resources only for Rx sources. This limitation might also apply to Cisco Nexus 9500 Series switches, depending on the SPAN or ERSPAN source's forwarding engine instance mappings.

Unicast Routing Verified Scalability Limits

Feature	Verified Limits
IPv4 ARP and IPv6 ND	
IPv4 ARP (Default routing template)	98,304
IPv6 ND (Default routing template)	98,304
IPv4 and IPv6 Routes	
Default Routing Template	
IPv4 host routes [3]	1,146,880
IPv6 host routes[4]	212,992
IPv4 LPM routes	1,048,576
IPv6 LPM routes	524,288
Unicast Protocols	

Feature	Verified Limits
Bidirectional Forwarding Detection (BFD)	
BFD sessions (echo mode)	128 when the BFD intervals are set to default, which is 50 ms 1000 (IPv4 and IPv6) sessions when the BFD intervals are relaxed to 300 ms **Note** For Nexus 9800 switches, the maximum session limit per L3 port channel and its subinterfaces is 128.
Border Gateway Protocol	
BGP neighbors (IPv4 and IPv6 combined)	1024
HSRP	
HSRP groups	1000 (virtual MAC address support)[5]
IS-IS	
IS-ISv4 adjacencies (either L1, L2, or sum of L1 and L2 with default timers)	255
IS-ISv4 BFD sessions (with default timers)	255
IS-ISv4 routes	10,000
OSPFv2 and OSPFv3	
OSPFv2/OSPFv3 LSA/LSDB size	100,000
OSPFv2/OSPFv3 areas	100
OSPFv2/OSPFv3 neighbors	256
Policy Based Routing	
Configured sequences per policy	128
Next-hop addresses per policy	32
IPv4 ACEs (unidimensional)	3582 (per network forwarding engine)
IPv6 ACEs (unidimensional)	1792 (per network forwarding engine)
Interfaces with PBR policy [6]	512
VRRP	
VRRPv3 groups per interface	250
VRRPv3 groups with default timers (1 s)	250
VRRPv3 groups with relaxed timers (3 s)	250

Feature	Verified Limits
ECMP Scale	
ECMP Paths (IPv4 and IPv6 Unicast Address-family)	64
ECMP Groups	8192

[3] The hash table is subject to collisions. Depending on the host route pattern, collisions might occur.
[4] The hash table is subject to collisions. Depending on the host route pattern, collisions might occur.
[5] If you have more than 490 groups, then only one group per SVI. SVIs cannot have a user defined MAC or any VRRP group with it.
[6] When using PBR with the "set vrf" option, if the user attempts to shut down the recircular port, it will trigger a modify PPF session from RPM. If the maximum labels (510) have already been consumed and atomic update is enabled, a "label allocation failure" will occur. To avoid this, the user needs to disable atomic update.

Note

- With IPv6 scale, traffic loss could be there for a few seconds during switchover.
- The maximum number of PBR next-hops based on 4 FM-E supported is 192 per slice of the forwarding engine
  - The IPv4/IPv6 host routes and the IPv4 multicast routes share the same hardware table. Limits are provided for both the default line card mode and the max host line card mode.
  - The IPv4 and IPv6 unicast routes share the same hardware table. Limits are provided for both the default line card mode and the max host line card mode.
  - High availability (graceful restart and stateful switchover) is not supported when unicast or multicast aggressive timers are configured at any scale.

Guidelines and Limitations for OSPF Verified Scalability Limits

- To achieve the highest scale, we recommend that you use a single OSPF instance instead of multiple instances.
- Each OSPFv2 and OSPFv3 scale value might vary when combined with other parameters.
- The graceful restart timeout value might be increased in multidimensional scenarios.

VXLAN Verified Scalability Limits

Feature	Verified Limits
VTEP Peers[7]	350 vPC pairs or 1000 standalone VTEPs with IR 512 standalone VTEPs with multicast underlay
Underlay multicast groups	512
Overlay EVPN ECMP	

Feature	Verified Limits
ECMP Paths **Note** An ECMP entry is created for each tunnel. (There may be multiple tunnels for each VXLAN peer).	64
Underlay EVPN ECMP	
ECMP Paths	64
Multi-Site [8]	
Number of sites	128
Number of sites for TRM	16 sites
Tenant Route Multicast Layer 3 Mode with VXLAN BGP eVPN	
VXLAN Layer 3 VNI/VRFs	250
Total Multicast routes (PIM ASM & PIM SSM)	32,000
VXLAN Flood and Learn	
Ingress replication peers [9]	512
VXLAN BGP eVPN	
Layer 2 VNIs	1000
Layer 3 VNIs / VRFs[10]	1000
Underlay multicast groups	512
MAC addresses	96,000 **Note** To scale MAC addresses, use **system routing template-dual-stack-host-scale** command and reload the switch.
	90,000
Overlay IPv4 LPM routes	1,000,000
Overlay IPv6 LPM routes	512,000
VXLAN BGP eVPN Ingress Replication	
Layer 2 VNIs	1000
Layer 3 VNIs / VRFs[11]	1000
VTEPs	350 vPC pairs or 1000 standalone VTEPs with IR

Feature	Verified Limits
MAC addresses	96,000
IPv4 host routes	512,000
IPv6 host routes	112,000

[7] In case of IR, each VNI can have a max of 64 peers.
[8] All the other BGW numbers (number of supported L2VNIs, L3VNIs, MAC addresses, IP addresses, and so on) match the values supported on a generic VXLAN EVPN VTEP node.
[9] In case of IR, each VNI can have a maximum number of 64 peers; 512 peers supported on 100 VNIs only.
[10] ECMP objects are not shared across multiple VRFs.
[11] ECMP objects are not shared across multiple VRFs.

Americas Headquarters
Cisco Systems, Inc.
San Jose, CA 95134-1706
USA

Asia Pacific Headquarters
CiscoSystems(USA)Pte.Ltd.
Singapore

Europe Headquarters
CiscoSystemsInternationalBV
Amsterdam,TheNetherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at www.cisco.com/go/offices.