



## **Cisco Nexus 9000 Series NX-OS Verified Scalability Guide, Release 7.0(3)I7(4)**

[Verified Scalability Limits](#) 2

## Verified Scalability Limits

This document describes the Cisco NX-OS configuration limits for the Cisco Nexus 9000 Series switches.

### Introduction

This document describes the Cisco NX-OS configuration limits for the Cisco Nexus 9000 Series switches.

The values provided in this guide should not be interpreted as theoretical system limits for Cisco Nexus 9000 Series 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

The tables in this section list the verified scalability limits for Cisco NX-OS Release 7.0(3)I7(4). 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 currently supported by this Cisco NX-OS release for the corresponding feature. If the hardware is capable of a higher scale, future software releases might increase this verified maximum limit. Results might differ from the values listed here when trying to achieve maximum scalability with multiple features enabled.

**Table 1: Cisco Nexus 2000 Series Fabric Extenders (FEX) Straight Through Mode Verified Scalability Limits (Unidimensional)**

Feature	9500 Platform Verified Limit <sup>1</sup>	9300 Platform Verified Limit <sup>2</sup>	9200 Platform Verified Limit	9300-EX Platform Verified Limit	9300-FX Platform Verified Limit
Fabric Extenders <sup>3</sup> and Fabric Extender server interfaces	32 and 1536	16 and 768	Not applicable	16 and 768	16 and 768
VLANs across all Fabric Extenders	2000	2000	Not applicable	562	562
VLANs per Fabric Extender server interface <sup>4</sup>	75	75	Not applicable	75	75
Port channels	426	256	Not applicable	232	232
Unique Fabric Extenders per Cisco Nexus 9500 Series supported line card	12	Not applicable	Not applicable	Not applicable	Not applicable

- <sup>1</sup> The Cisco Nexus 2200 Series and B22 Series Fabric Extenders are supported with X9464PX and X9564PX line cards on Cisco Nexus 9500 Series switches. The Cisco Nexus 2300 Series Fabric Extenders are supported with X9432PQ, X9464PX, X9464TX, X9536PQ, X9564PX, X9564TX, and X9636PQ line cards on Cisco Nexus 9500 Series switches.
- <sup>2</sup> The Cisco Nexus 2200 Series and B22 Series Fabric Extenders are supported with the Cisco Nexus 9396PX, 9372PX, and 9372PX-E chassis. The Cisco Nexus 2300 Series Fabric Extenders are supported with the Cisco Nexus 9332PQ, 9396PX, 9372PX, and 9372PX-E chassis.
- <sup>3</sup> When FEX configured using "AA" mode, then the maximum number of 6 FEX on NFE base ToR and 16 FEX for LSE base ToR are supported.
- <sup>4</sup> For FEX HIF port channels, Cisco recommends that you enable STP port type edge using the **spanning tree port type edge [trunk]** command.

**Table 2: FCoE Verified Scalability Limits (Unidimensional)**

Feature <sup>5</sup>	Verified Limit <sup>6</sup>
Number of FLOGI per port	256
Number of FLOGI per switch	1000
Number of port channels	8 <sup>7</sup>
Maximum number of member ports in a port channel	16
Number of VFCs	544
Number of VSANs	32

<sup>5</sup> Feature Verified Limit is 5.

<sup>6</sup> All above numbers are verified against Platform — Nexus 93180YC-FX

<sup>7</sup> The number of SAN port channels and virtual FC port channels, together, can be only 8 on the Cisco Nexus 9000 Series switch.



**Note** For a list of platforms on which FCoE is supported, please refer to the Cisco Nexus 9000 Series NX-OS FCoE Configuration Guide, Release 7.x, at the following URL:  
[https://www.cisco.com/c/en/us/td/docs/switches/datacenter/nexus9000/sw/7-x/FCoE/configuration/guide/b\\_Cisco\\_Nexus\\_9000\\_Series\\_NX-OS\\_FCoE\\_Configuration\\_Guide\\_7x.html](https://www.cisco.com/c/en/us/td/docs/switches/datacenter/nexus9000/sw/7-x/FCoE/configuration/guide/b_Cisco_Nexus_9000_Series_NX-OS_FCoE_Configuration_Guide_7x.html)

**Table 3: FC Verified Scalability Limits (Unidimensional)**

Feature <sup>8</sup>	Verified Limit <sup>9</sup>
Number of FLOGI per port	256
Number of FLOGI per switch	1000
Number of port channels	8 <sup>10</sup>
Maximum number of member ports in a port channel	16
Max number of FC ports supported	48
Number of VSANs	32

<sup>8</sup> Feature Verified Limit is 6.

<sup>9</sup> All above numbers are verified against Platform — Nexus 93180YC-FX

<sup>10</sup> The number of SAN port channels and virtual FC port channels, together, can be only 8 on the Cisco Nexus 9000 Series switch.

**Table 4: Intelligent Traffic Director Verified Scalability (Unidimensional)**

Platform	Number of Device Groups	Number of Services	Nodes per Device Group	Buckets per Service	Include acl-ace entries
<b>Without include ACL (for IPv4)</b>					
Cisco Nexus 93xxx-EX with TCAM allocated 2048	8	8	32	64	None
Cisco Nexus 9372xX with TCAM allocated 2048	8	8	32	64	None
Cisco Nexus 9500 with FM with TCAM allocated 2048	8	8	32	64	None
Cisco Nexus 93xxx-FX with TCAM allocated 2304	8	32	32	128	None
<b>With include ACL (for IPv4)</b>					
Cisco Nexus 93xxx-EX with TCAM allocated 2048	2	2	32	32	14
Cisco Nexus 9372xX with TCAM allocated 2048	2	2	32	32	14
Cisco Nexus 9500 with FM with TCAM allocated 2048	2	2	32	32	14
Cisco Nexus 93xxx-FX with TCAM allocated 2304	2	2	32	32	16
<b>With VIP (for IPv4)</b>					

<b>Platform</b>	<b>Number of Device Groups</b>	<b>Number of Services</b>	<b>Nodes per Device Group</b>	<b>Buckets per Service</b>	<b>Include acl-ace entries</b>
Cisco Nexus 93xxx-EX with TCAM allocated 2048	4	4	32	64	None
Cisco Nexus 9372TX with TCAM allocated 2048	4	4	32	64	None
Cisco Nexus 9500 with FM with TCAM allocated 2048	4	4	32	64	None
Cisco Nexus 93xxx-FX with TCAM allocated 2304	4	4	32	64	None
<b>Without include ACL (for IPv6)</b>					
Cisco Nexus 93xxx-EX with TCAM allocated 2048	8	8	32	32	None
Cisco NExus 93xxx-FX with TCAM allocated 2304	8	8	32	64	None
<b>With include ACL (for IPv6)</b>					
93xxx-EX with TCAM allocated 2048	2	2	32	16	10
Cisco Nexus 93xxx-FX with TCAM allocated 2304	2	2	32	16	16
<b>With VIP (For IPv6)</b>					
Cisco Nexus 93xxx-EX with TCAM allocated 2048	2	2	32	64	None

Platform	Number of Device Groups	Number of Services	Nodes per Device Group	Buckets per Service	Include acl-ace entries
Cisco Nexus 93xxx-FX with TCAM allocated 2304	4	4	32	32	None

**Table 5: Interfaces Verified Scalability Limits (Unidimensional)**

Feature	9500 Platform Verified Limit	9300 Platform Verified Limit	9200 Platform Verified Limit	9300-EX Platform Verified Limit	9300-FX Platform Verified Limit	9300-FX2 Platform Verified Limit	9500 with X9700-EX Verified Limit
DHCP clients per switch	10 (IPv4) + 10 (IPv6)	10 (IPv4) + 10 (IPv6)	10 (IPv4) + 10 (IPv6)	10 (IPv4) + 10 (IPv6)	10 (IPv4) + 10 (IPv6)	10 (IPv4) + 10 (IPv6)	10 (IPv4) + 10 (IPv6)
IP DHCP relay addresses (helper addresses) per switch	32 (IPv4) + 32 (IPv6)	32 (IPv4) + 32 (IPv6)	32 (IPv4) + 32 (IPv6)	32 (IPv4) + 32 (IPv6)	32 (IPv4) + 32 (IPv6)	32 (IPv4) + 32 (IPv6)	32 (IPv4) + 32 (IPv6)
Generic routing encapsulation (GRE) tunnels	8	8	8	8	8	8	8
Port channel links	32	32	32	32	32	32	32
SVIs	490 (with HSRP), 1500 (without HSRP)	450 (with HSRP)	490	450 (with HSRP)	450 (with HSRP)	450 (with HSRP)	490 (with HSRP), 1500 (without HSRP)
SVI Unnumbered	Primary (50) Secondary (450) 1 primary SVI can have a maximum of 50 secondary SVIs	Primary (50) Secondary (450) 1 primary SVI can have a maximum of 50 secondary SVIs	Primary (50) Secondary (450) 1 primary SVI can have a maximum of 50 secondary SVIs	Primary (50) Secondary (450) 1 primary SVI can have a maximum of 50 secondary SVIs	Primary (50) Secondary (450) 1 primary SVI can have a maximum of 50 secondary SVIs	Primary (50) Secondary (450) 1 primary SVI can have a maximum of 50 secondary SVIs	Primary (50) Secondary (450) 1 primary SVI can have a maximum of 50 secondary SVIs
vPCs	300	48	48	48	48	48	300

Feature	9500 Platform Verified Limit	9300 Platform Verified Limit	9200 Platform Verified Limit	9300-EX Platform Verified Limit	9300-FX Platform Verified Limit	9300-FX2 Platform Verified Limit	9500 with X9700-EX Verified Limit
Static network address translation (NAT)	Not applicable	1023	1023	1023	1023	1023	Not applicable
Dynamic network address translation (NAT)	Not applicable	1023	1023	1023	1023	1023	Not applicable
Static twice network address translation (NAT)	Not applicable	768	768	768	768	768	Not applicable
Dynamic twice network address translation (NAT)	Not applicable	1023	1023	1023	1023	1023	Not applicable

**Table 6: Label Switching Verified Scalability Limits (Unidimensional)**

Feature	9500 Platform Verified Limit	9300 Platform Verified Limit	9200 Platform Verified Limit	9300-EX Platform Verified Limit	9300-FX Platform Verified Limit	9300-FX2 Platform Verified Limit	9500 with X9700-EX Verified Limit
Forwarding Equivalence Classes (FECs)	128	128	MPLS Heavy Template: 512 Default: 128	MPLS Heavy Template: 2048 Default: 1024	MPLS Heavy Template: 2048 Default: 1024	Not supported	MPLS Heavy Template: 2048
Equal-cost multipaths (ECMPs)	16	16	16	32	32	Not supported	32
Equal-cost multipaths Groups (ECMPs)	Not applicable	Not applicable	Not applicable	MPLS Heavy Template: 2048 Default: 1024	MPLS Heavy Template: 2048 Default: 1024	Not supported	MPLS Heavy Template: 2048 Default: 1024
FECs * ECMPs	1000	1000	1000	Not applicable	Not applicable	Not supported	Not applicable

Feature	9500 Platform Verified Limit	9300 Platform Verified Limit	9200 Platform Verified Limit	9300-EX Platform Verified Limit	9300-FX Platform Verified Limit	9300-FX2 Platform Verified Limit	9500 with X9700-EX Verified Limit
Flex counters for static MPLS in egress direction	4000	4000	4000	Not applicable	Not applicable	Not supported	Not applicable
Flex counters per adjacency	2	2	2	Not applicable	Not applicable	Not supported	Not applicable
Adjacencies	1024	1024	1024	48k	48k	Not supported	48k
Egress Peer Engineering	64	64	64	64	64	Not supported	64
Label-switched paths (LSPs) for label stack imposition <sup>11</sup>	128 (with 4-way ECMP and 3 label stack push)	128 (with 4-way ECMP and 3 label stack push)	256 (with 32-way ECMP and 5 label stack push)	256 (with 32-way ECMP and 5 label stack push)	256 (with 32-way ECMP and 5 label stack push)	Not supported	256 (with 32-way ECMP and 5 label stack push)
Layer 3 EVPN	128	128	Not applicable	1000 (With MPLS Heavy Template)	1000 (With MPLS Heavy Template)	Not supported	Not applicable
<b>Private VLANs (PVLANS)</b>							
Primary VLANs	16	16	Not applicable	16	16	16	16
Secondary VLANs	20	20	Not applicable	20	20	20	20
Ports in Community host mode	40	40	Not applicable	40	40	40	40
Ports in isolated host mode	20	40	Not applicable	40	40	40	40
Ports in isolated trunk host mode	22	40	Not applicable	40	40	40	40
Ports in promiscuous mode	48	5	Not applicable	5	5	5	5



Feature	9500 Platform Verified Limit	9300 Platform Verified Limit	9200 Platform Verified Limit	9300-EX Platform Verified Limit	9300-FX Platform Verified Limit	9300-FX2 Platform Verified Limit	9500 with X9700-EX Verified Limit
Ports in promiscuous trunk mode	80	5	Not applicable	5	5	5	5
PVLANs allowed on a PVLAN port	16	16	Not applicable	16	16	16	16

<sup>11</sup> For Cisco Nexus 9300 and 9500 Series switches, LSPs \*ECMP\* label stack push cannot exceed 1500.



**Note** For network scalability, Cisco recommends using a hierarchical routing design with multi-hop BGP for advertising the attached prefixes from a top-of-rack (ToR) or border leaf switch.

**Table 7: Layer 2 Switching Verified Scalability Limits (Unidimensional)**

Feature	9500 Platform Verified Limit	9300 Platform Verified Limit	9200 Platform Verified Limit	9300-EX Platform Verified Limit	9300-FX Platform Verified Limit	9300-FX2 Platform Verified Limit	9500 with X9700-EX Verified Limit
MAC addresses	90,000	90,000	92,000	92,000	92,000	92,000	92,000
MST instances	64	64	64	64	64	64	64
MST virtual ports	85,000	48,000	48,000	48,000	48,000	48,000	85,000
RPVST virtual ports	22,000	48,000	48,000	48,000	48,000	48,000	22,000
VLANs	3967 (the remaining 127 VLANs are reserved)	3967 (the remaining 127 VLANs are reserved)	3967 (the remaining 127 VLANs are reserved)	3967 (the remaining 127 VLANs are reserved)	3967 (the remaining 127 VLANs are reserved)	3967 (the remaining 127 VLANs are reserved)	3967 (the remaining 127 VLANs are reserved)
VLANs in RPVST mode	500	500	3967	3967	3967	3967	3967 <sup>12</sup>

Feature	9500 Platform Verified Limit	9300 Platform Verified Limit	9200 Platform Verified Limit	9300-EX Platform Verified Limit	9300-FX Platform Verified Limit	9300-FX2 Platform Verified Limit	9500 with X9700-EX Verified Limit
Total number of VLANs × ports with switchport isolated (3967 VLANs x 48 ports)	190,000	190,000	190,000	190,000	190,000	190,000	190,000
<b>Private VLANs (PVLANS)</b>							
Primary VLANs	16	16	Not applicable	16	16	16	16
Secondary VLANs	20	20	Not applicable	20	20	20	20
Ports in Community host mode	40	40	Not applicable	40	40	40	40
Ports in isolated host mode	20	40	Not applicable	40	40	40	40
Ports in isolated trunk host mode	22	40	Not applicable	40	40	40	40
Ports in promiscuous mode	48	5	Not applicable	5	5	5	5
Ports in promiscuous trunk mode	80	5	Not applicable	5	5	5	5
PVLANS allowed on a PVLAN port	16	16	Not applicable	16	16	16	16

<sup>12</sup> On EOR, support is for 12000 PV count with 3967 vlans and RPVST with default timers. If 22000 PV count is needed with 3968 vlans and RPVST, recommended hello timer value is 4 or higher. It is also recommended to tune forward delay and max age accordingly



**Note** The number of supported VLANs per vPC should be within the MST or RPVST virtual port count specified in this table, depending on the topology.



**Note** The number of supported STP VLAN port instances, for Fabric Extender host interface ports, should be less than 13,000.

**Table 8: Multicast Routing Verified Scalability Limits (Unidimensional)**

Feature	9500 Platform Verified Limit	9300 Platform Verified Limit	9200 Platform Verified Limit	9300-EX Platform Verified Limit	9300-FX Platform Verified Limit	9300-FX2 Platform Verified Limit	9500 with X9700-EX Verified Limit
IPv4 multicast routes	32,000 (Layer 2 + Layer 3)	8000 (Layer 2 + Layer 3)	8000 (Layer 2 + Layer 3) Cisco Nexus 9200 switches do not support the system routing template-lpm-heavy mode for IPv4 multicast routes. Make sure to reset the LPM maximum limit to 0.	8000 (Layer 2 + Layer 3); 32,000 (layer 2 + Layer 3 with system routing template-lpm-heavy mode)	8000 (Layer 2 + Layer 3); 32,000 (layer 2 + Layer 3 with system routing template-lpm-heavy mode)	8000 (Layer 2 + Layer 3); 32,000 (layer 2 + Layer 3 with system routing template-lpm-heavy mode)	8000 (Layer 2 + Layer 3); 32,000 (layer 2 + Layer 3 with system routing template-lpm-heavy mode); 8000 (with system routing template-lpm-heavy mode)
Outgoing interfaces (OIFs)	40 (SVI + physical Layer 3) or 256 (physical Layer 3)	40 (SVI + physical Layer 3)	40 (SVI + physical Layer 3)	40 (SVI + physical Layer 3)	40 (SVI + physical Layer 3)	40 (SVI + physical Layer 3)	40 (SVI + physical Layer 3) or 256 (physical Layer 3)
IGMP snooping groups	32,000	8000	8000	8000	8000	8000	32,000
PIM neighbors	500	250	250	250	250	250	500



**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.



**Note** High availability (graceful restart and stateful switchover) is not supported when unicast or multicast aggressive timers are configured at any scale.

**Table 9: Programmability Verified Scalability Limits (Unidimensional)**

<b>Feature</b>	<b>9500 Platform Verified Limit</b>	<b>9300 Platform Verified Limit</b>	<b>9200 Platform Verified Limit</b>	<b>9300-EX Platform Verified Limit</b>	<b>9300-FX Platform Verified Limit</b>	<b>9300-FX2 Platform Verified Limit</b>	<b>9500 with X9700-EX Verified Limit</b>
<b>OpenFlow</b>							
OpenFlow ports	Not applicable	96	Not applicable	Not applicable	Not applicable	Not applicable	Not applicable
OpenFlow Layer 2 flows	Not applicable	32,000	Not applicable	Not applicable	Not applicable	Not applicable	Not applicable
OpenFlow Layer 3 flows	Not applicable	3000	Not applicable	Not applicable	Not applicable	Not applicable	Not applicable
OpenFlow IPv6 Layer 3 flows	Not applicable	1500	Not applicable	Not applicable	Not applicable	Not applicable	Not applicable

**Table 10: Security Verified Scalability Limits (Unidimensional)**

<b>Feature</b>	<b>9500 Platform Verified Limit</b>	<b>9300 Platform Verified Limit</b>	<b>9200 Platform Verified Limit</b>	<b>9300-EX Platform Verified Limit</b>	<b>9300-FX Platform Verified Limit</b>	<b>9300-FX2 Platform Verified Limit</b>	<b>9500 with X9700-EX Verified Limit</b>
DHCP snooping bindings	2048	2048	2048	2048	2048	2048	2048
IPv4 ingress TCAM entries	3072 (per network forwarding engine)	3072 (per network forwarding engine)	3582 (per slice of the forwarding engine)	3582 (per slice of the forwarding engine)	3582 (per slice of the forwarding engine)	3582	3582 (per slice of the forwarding engine)
IPv4 egress TCAM entries	768 (per network forwarding engine)	768 (per network forwarding engine)	1792 (per slice of the forwarding engine)	1792 (per slice of the forwarding engine)	1792 (per slice of the forwarding engine)	1792 (per slice of the forwarding engine)	1792 (per slice of the forwarding engine)
IPv6 ingress TCAM entries	1536 (per network forwarding engine)	1536 (per network forwarding engine)	1792 (per slice of the forwarding engine)	1792 (per slice of the forwarding engine)	1792 (per slice of the forwarding engine)	1792 (per slice of the forwarding engine)	1792 (per slice of the forwarding engine)
IPv6 egress TCAM entries	256 (per network forwarding engine)	256 (per network forwarding engine)	896 (per slice of the forwarding engine)	896 (per slice of the forwarding engine)	896 (per slice of the forwarding engine)	896 (per slice of the forwarding engine)	896 (per slice of the forwarding engine)



**Note** The TCAM entries scalability limits also apply to policy-based TCAM entries (PBACLs).

**Table 11: System Management Verified Scalability Limits (Unidimensional)**

Feature	9500 Platform Verified Limit	9300 Platform Verified Limit	9200 Platform Verified Limit	9300-EX Platform Verified Limit	9300-FX Platform Verified Limit	9300-FX2 Platform Verified Limit	9500 with X9700-EX Verified Limit
<b>MPLS Stripping</b>							
Labels	12,000	12,000	No limit	Not applicable	Not applicable	Not applicable	Not applicable
Ingress interfaces	400	48	48	Not applicable	Not applicable	Not applicable	Not applicable
Egress interfaces	64	16	16	Not applicable	Not applicable	Not applicable	Not applicable
<b>PTP</b>							
PTP ports <sup>13</sup>	44	44	44	44	44	44	1305
<b>sFlow</b>							
sFlow ports	256	64	64	64	64	64	256
<b>SPAN and ERSPAN</b>							
Configurable SPAN or ERSPAN sessions	32	32	32	32	32	32	32
Active SPAN or ERSPAN sessions <sup>14</sup>	4 to 32, based on the number of line cards and the session configuration	4	4	4	4	4	4 to 32, based on the number of line cards and the session configuration
Active localized SPAN or ERSPAN sessions per line card <sup>15</sup>	4	4	4	4	4	4	4

<b>Feature</b>	<b>9500 Platform Verified Limit</b>	<b>9300 Platform Verified Limit</b>	<b>9200 Platform Verified Limit</b>	<b>9300-EX Platform Verified Limit</b>	<b>9300-FX Platform Verified Limit</b>	<b>9300-FX2 Platform Verified Limit</b>	<b>9500 with X9700-EX Verified Limit</b>
Source interfaces per SPAN or ERSPAN session (Rx and Tx, Rx, or Tx)	48	48	48	48	48	48	48
Destination interfaces per SPAN session	1 (physical/PO interface)	1 (physical/PO interface)	1 (physical/PO interface)	1 (physical/PO interface)	1 (physical/PO interface)	1 (physical/PO interface)	1 (physical/PO interface)
Source VLANs per SPAN or ERSPAN session	32	32	32	32	32	32	32
<b>Tap Aggregation</b>							
Redirect interfaces in the redirect port list	12	12	12	12	12	12	Not applicable
Redirect port lists (or fan outs) per system	100	100	50	50	50	50	Not applicable
<b>NetFlow</b>							
Flow monitors	Not applicable	Not applicable	Not applicable	2 exporters and 2 flow monitors per type (2 IPv4 flow monitors and 2 IPv6 flow monitors)	2 exporters and 32 flow monitors per type (32 Layer 2 flow monitors, 32 IPv4 flow monitors, and 32 IPv6 flow monitors)	2 exporters and 32 flow monitors per type (32 Layer 2 flow monitors, 32 IPv4 flow monitors, and 32 IPv6 flow monitors)	Not applicable

Feature	9500 Platform Verified Limit	9300 Platform Verified Limit	9200 Platform Verified Limit	9300-EX Platform Verified Limit	9300-FX Platform Verified Limit	9300-FX2 Platform Verified Limit	9500 with X9700-EX Verified Limit
Number of Layer 3 interfaces (Layer 3 ports, port channels, and SVIs) to which IPv4 flow monitors can be applied	Not applicable	Not applicable	Not applicable	1016 (with members on just one ASIC slice) or 508 (with members on both ASIC slices)  <b>Note</b> You can use the show interface hardware-mappings command to check if the interface belongs to ASIC slice 0 or slice 1.	Not applicable	Not applicable	Not applicable
Number of Layer 3 interfaces (Layer 3 ports, port channels, and SVIs) to which IPv6 flow monitors can be applied	Not applicable	Not applicable	Not applicable	252 (with members on just one ASIC slice) or 126 (with members on both ASIC slices)  <b>Note</b> You can use the show interface hardware-mappings command to check if the interface belongs to ASIC slice 0 or slice 1.	Not applicable	Not applicable	Not applicable

<sup>13</sup> With PTP offload enabled.

<sup>14</sup> 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.

<sup>15</sup> The number of SPAN or ERSPAN sessions per line card reduces to two if the same interface is configured as the bidirectional source in more than one session.




---

**Note** Beginning with Cisco NX-OS Release 7.0(3)I1(2), PTP is supported for all Cisco Nexus 9000 Series hardware except for the 100G 9408PC line card and the 100G M4PC generic expansion module (GEM).

---

**Table 12: Unicast Routing Verified Scalability Limits (Unidimensional)**

<b>Feature</b>	<b>9500 Platform Verified Limit</b>	<b>9300 Platform Verified Limit</b>	<b>9200 Platform Verified Limit</b>	<b>9364C Platform Verified Limit</b>	<b>9300-EX Platform Verified Limit</b>	<b>9300-FX Platform Verified Limit</b>	<b>9300-FX2 Platform Verified Limit</b>	<b>9500 with X9700-EX /FX Verified Limit</b>
<b>Unicast Routing</b>								
BFD sessions (echo mode)	512	128	128	128	128	128	128	512 <sup>16</sup>
BGP neighbors	2000	512	512 (IPv4), 512 (IPv6), or 256 (IPv4 + IPv6)	512	512	512	512	2000
EIGRP routes	20,000	20,000	20,000	20,000	20,000	20,000	20,000	20,000
EIGRP neighbors	512	256	256	256	256	256	256	512
HSRP groups	490	490	490	490	490	490	490	490
IPv4 ARP	48,000	48,000	32,000	32,000	48,000	48,000 (in default routing mode, Hash Table: Shared between IPv6 ND, IPv4 ARP) 32,000 (with system routing template -lpm-heavy mode, Hash Table: Shared between IPv6 ND, IPv4 ARP)	64,000 (in default routing mode, Hash Table: Shared between IPv6 ND, IPv4 ARP)	48,000 (default), 32,000 (LPM heavy)



Feature	9500 Platform Verified Limit	9300 Platform Verified Limit	9200 Platform Verified Limit	9364C Platform Verified Limit	9300-EX Platform Verified Limit	9300-FX Platform Verified Limit	9300-FX2 Platform Verified Limit	9500 with X9700-EX /FX Verified Limit
IPv4 host routes <a href="#">17</a>	Default System Routing Mode: 208,000 (hash table and there will be more collisions after 80%)  ALPM Routing Mode: 128,000 with host Routes Programmed in the LPM Table	Default System Routing Mode: 208,000 (hash table and there will be more collisions after 80%)  ALPM Routing Mode: 128,000 with host Routes Programmed in the LPM Table	96,000 (hash table and there will be more collisions after 80%)	96,000	458,000 (default); 786,000 (with system routing template-lpm-heavy mode)	471,000 (default); 786,000 (with system routing template-lpm-heavy mode)	471,000 (default); 786,000 (with system routing template-lpm-heavy mode)	589,000 (default); 786,000 (with system routing template-lpm-heavy mode)
IPv6 host routes <a href="#">18</a>	Default System Routing Mode: 104,000 (hash table and there will be more collisions after 80%)  ALPM Routing Mode: 16000 with host Routes Programmed in the LPM Table	Default System Routing Mode: 104,000 (hash table and there will be more collisions after 80%)  ALPM Routing Mode: 16000 with host Routes Programmed in the LPM Table	48,000 (hash table and there will be more collisions after 80%)	48,000	24,000	265,000 (default) 442,000 (with system routing template-lpm-heavy mode)	265,000 (default) 442,000 (with system routing template-lpm-heavy mode)	32,000 (FM-E), 235000 (FM-E2)

<b>Feature</b>	<b>9500 Platform Verified Limit</b>	<b>9300 Platform Verified Limit</b>	<b>9200 Platform Verified Limit</b>	<b>9364C Platform Verified Limit</b>	<b>9300-EX Platform Verified Limit</b>	<b>9300-FX Platform Verified Limit</b>	<b>9300-FX2 Platform Verified Limit</b>	<b>9500 with X9700-EX /FX Verified Limit</b>
IPv6 ND	48,000	48,000	32,000 (default), 16,000 (lpm heavy)	32,000 (default), 16,000 (lpm heavy)	32,000 (default), 16,000 (lpm heavy)	32,000 (in default routing mode, Hash Table: Shared between IPv6 ND, IPv4 ARP)	32,000 (default), 16,000 (lpm heavy)	32,000

<b>Feature</b>	<b>9500 Platform Verified Limit</b>	<b>9300 Platform Verified Limit</b>	<b>9200 Platform Verified Limit</b>	<b>9364C Platform Verified Limit</b>	<b>9300-EX Platform Verified Limit</b>	<b>9300-FX Platform Verified Limit</b>	<b>9300-FX2 Platform Verified Limit</b>	<b>9500 with X9700-EX /FX Verified Limit</b>
IPv4 unicast routes (LPM)*	128,000 (default system routing mode) 16,000 (max-host routing mode) 128,000 with no IPv6 routes (64-bit ALPM routing mode)	12,000 (default system routing mode) 128,000 (ALPM routing mode)			458,000 (default)	471,000 (default)	471,000 (default)	589,000 (default)
IPv6 unicast routes (LPM)*	20,000 (default system routing mode) 4000 (max-host routing mode) 80,000 with no IPv4 routes (64-bit ALPM routing mode)	7000 (6000 routes < /64, 1000 routes > /64) (default system routing mode) 20,000 (ALPM routing mode)			206,000 (/64 prefix length); 1900 (non /64 prefix length)	265,000 (default)	265,000 (default)	176,000 (/64 prefix length); 3900 (non /64 prefix length) ( FM-E ) 235,000 ( any prefix len) ( FM-E2)

Feature	9500 Platform Verified Limit	9300 Platform Verified Limit	9200 Platform Verified Limit	9364C Platform Verified Limit	9300-EX Platform Verified Limit	9300-FX Platform Verified Limit	9300-FX2 Platform Verified Limit	9500 with X9700-EX /FX Verified Limit
			<p>Default values: 6000 (IPv4), 1900 (IPv6), and 2000 (multicast)</p> <p>With hardware profile multicast max-limit lpm-entries 0 configured: 8000 (IPv4), 1900 (IPv6), and 0 (multicast)</p> <p>With hardware profile ipv6 lpm-entries maximum 0 configured: 14,000 (IPv4), 0 (IPv6), and 2000 (multicast)</p> <p>With hardware profile ipv6 lpm-entries maximum 4096 and hardware profile multicast max-limit lpm-entries 0 configured: 0 (IPv4), 4096 (IPv6), and 0 (multicast)</p> <p>When you allocate the entire table for IPv4 or IPv6 LPM unicast</p>	<p>Default values: 8000 (IPv4), 1900 (IPv6), and 2000 (multicast)</p> <p>With hardware profile multicast max-limit lpm-entries 0 configured: 10000 (IPv4), 1900 (IPv6), and 0 (multicast)</p> <p>With hardware profile ipv6 lpm-entries maximum 0 configured: 14,000 (IPv4), 0 (IPv6), and 2000 (multicast)</p> <p>With hardware profile ipv6 lpm-entries maximum 4096 and hardware profile multicast max-limit lpm-entries 0 configured: 4000 (IPv4), 4096 (IPv6), and 0 (multicast)</p> <p>When you allocate the entire table for IPv4 or IPv6</p>				

Feature	9500 Platform Verified Limit	9300 Platform Verified Limit	9200 Platform Verified Limit	9364C Platform Verified Limit	9300-EX Platform Verified Limit	9300-FX Platform Verified Limit	9300-FX2 Platform Verified Limit	9500 with X9700-EX /FX Verified Limit
			routes, the other address family cannot be used	LPM unicast routes, the other address family cannot be used				
IPv4 and IPv6 unicast routes (LPM) in 64-bit ALPM routing mode	128,000 (IPv4) 80,000 (IPv6)	Not applicable	Not applicable	Not applicable	Not applicable	Not applicable	Not applicable	Not applicable
IPv4 host routes (LPM heavy mode)	Not applicable	Not applicable	Cisco Nexus 9236C, 9272Q, and 92304QC switches: 262,000  Cisco Nexus 92160YC-X switches: 650,000	262,000	786,000	786,000	786,000	786,000
IPv6 host routes (LPM heavy mode)	Not applicable	Not applicable	16,000	131,000	24,000 (protocol learned host)	442,000 (protocol learned host)	442,000 (protocol learned host)	32,000 (shared between IPv6 ND and protocol learned host) ( FM-E ) 235,000 ( FM-E2)
IPv4 LPM routes (LPM heavy mode)	Not applicable	Not applicable	Cisco Nexus 9236C, 9272Q, and 92304QC switches: 262,000  Cisco Nexus 92160YC-X switches: 650,000	262,000	786,000	786,000	786,000	786,000

<b>Feature</b>	<b>9500 Platform Verified Limit</b>	<b>9300 Platform Verified Limit</b>	<b>9200 Platform Verified Limit</b>	<b>9364C Platform Verified Limit</b>	<b>9300-EX Platform Verified Limit</b>	<b>9300-FX Platform Verified Limit</b>	<b>9300-FX2 Platform Verified Limit</b>	<b>9500 with X9700-EX /FX Verified Limit</b>
IPv6 LPM routes (LPM heavy mode)	Not applicable	Not applicable	Cisco Nexus 9236C, 9272Q, and 92304QC switches: 131,000 (/64 prefix length); 1900 (non /64 LPM scale)  Cisco Nexus 92160YC-X switches: 294,000 (/64 prefix length); 1900 (non /64 LPM scale)	131,000	353,000 (/64 prefix length); 1900 (non /64 prefix length)	442,000	442,000	235,000 (/64 prefix length); 3900 (non /64 prefix length) (FM-E)  235,000 (any prefix len) (FM-E2)
IPv4 host routes (dual-host mode)	Not applicable	Not applicable	163,000	163,000	262,000	262,000	262,000	Not applicable
IPv6 host routes (dual-host mode)	Not applicable	Not applicable	81,000	81,000	131,000	131,000	131,000	Not applicable
IPv4 LPM routes (dual-host mode)	Not applicable	Not applicable	6000	8000	6000	7000	7000	Not applicable
IPv6 LPM routes (dual-host mode)	Not applicable	Not applicable	1900	1900	1900	1900	1900	Not applicable
IPv4 ARP (dual-host mode)	Not applicable	Not applicable	64,000	64,000	64,000	64,000	64,000	Not applicable
IPv6 ND (dual-host mode)	Not applicable	Not applicable	64,000	64,000	64,000	64,000	64,000	Not applicable

<b>Feature</b>	<b>9500 Platform Verified Limit</b>	<b>9300 Platform Verified Limit</b>	<b>9200 Platform Verified Limit</b>	<b>9364C Platform Verified Limit</b>	<b>9300-EX Platform Verified Limit</b>	<b>9300-FX Platform Verified Limit</b>	<b>9300-FX2 Platform Verified Limit</b>	<b>9500 with X9700-EX /FX Verified Limit</b>
IPv4 host routes (internet peering mode)	Not applicable	Not applicable	Not applicable	Not applicable	1 Million (protocol learned host)	1 Million (protocol learned host)	1 Million (protocol learned host)	1 Million (protocol learned host)
IPv6 host routes (internet peering mode)	Not applicable	Not applicable	Not applicable	Not applicable	16,000 (Hash Table: Shared between IPv6 ND, IPv4 ARP, and protocol learned IPv6 host)	5000,000	5000,000	16,000 (Hash Table: Shared between IPv6 ND, IPv4 ARP, and protocol learned IPv6 host)
IPv4 LPM routes (internet peering mode)	Not applicable	Not applicable	Not applicable	Not applicable	1 Million	1 Million	1 Million	1 Million
IPv6 LPM routes (internet peering mode)	Not applicable	Not applicable	Not applicable	Not applicable	500,000 (Prefix length 0-83) 1900 (Prefix length /84-127)	500,000	500,000	176,947 (Prefix 0-47) 500,000 (Prefix length 48-83) 1900 (Prefix length /84-127)
IPv4 ARP (internet peering mode)	Not applicable	Not applicable	Not applicable	Not applicable	32,000 (Hash Table: Shared between IPv6 ND, IPv4 ARP, and protocol learned IPv6 host)	32,000 (Hash Table: Shared between IPv6 ND, IPv4 ARP)	32,000 (Hash Table: Shared between IPv6 ND, IPv4 ARP)	32,000 (Hash Table: Shared between IPv6 ND, IPv4 ARP, and protocol learned IPv6 host)

<b>Feature</b>	<b>9500 Platform Verified Limit</b>	<b>9300 Platform Verified Limit</b>	<b>9200 Platform Verified Limit</b>	<b>9364C Platform Verified Limit</b>	<b>9300-EX Platform Verified Limit</b>	<b>9300-FX Platform Verified Limit</b>	<b>9300-FX2 Platform Verified Limit</b>	<b>9500 with X9700-EX /FX Verified Limit</b>
IPv6 ND (internet peering mode)	Not applicable	Not applicable	Not applicable	Not applicable	16,000 (Hash Table: Shared between IPv6 ND, IPv4 ARP, and protocol learned IPv6 host)	16,000 (Hash Table: Shared between IPv6 ND, IPv4 ARP)	16,000 (Hash Table: Shared between IPv6 ND, IPv4 ARP)	16,000 (Hash Table: Shared between IPv6 ND, IPv4 ARP, and protocol learned IPv6 host)
IS-ISv4 adjacencies (either L1, L2, or sum of L1 and L2 with default timers)	255	255	255	255	255	255	255	255
IS-ISv4 BFD sessions (with default timers)	255	255	Not applicable	255	255	255	255	255
IS-ISv4 routes	10,000	10,000	10,000	10,000	10,000	10,000	10,000	10,000
IS-ISv4 network type	Point to point, broadcast	Point to point, broadcast	Point to point, broadcast	Point to point, broadcast	Point to point, broadcast	Point to point, broadcast	Point to point, broadcast	Point to point, broadcast
OSPFv2 neighbors	1000	256	256	256	256	256	256	1000
OSPFv3 neighbors	1000	256	256	256	256	256	256	1000
OSPF/OSPFv3 LSA/LSDB size	100,000	100,000	100,000	100,000	100,000	100,000	100,000	100,000
OSPF/OSPFv3 areas	100	100	100	100	100	100	100	100
Static routes	4000	4000	4000	4000	4000	4000	4000	4000
VRFs	1000	1000	1000	1000	1000	1000	1000	1000



Feature	9500 Platform Verified Limit	9300 Platform Verified Limit	9200 Platform Verified Limit	9364C Platform Verified Limit	9300-EX Platform Verified Limit	9300-FX Platform Verified Limit	9300-FX2 Platform Verified Limit	9500 with X9700-EX /FX Verified Limit
VRRP groups per interface or I/O module	250	250	250	250	250	250	250	250
<b>Policy-based routing (PBR)</b>								
Configured sequences per policy	256	256	128	Not applicable	128	128	128	128
Next-hop addresses per policy	32	32	32	Not applicable	32	32	32	32
IPv4 ACEs (unidimensional)	3072 (per network forwarding engine)	3072 (per network forwarding engine)	3582 (per network forwarding engine)	Not applicable	3582 (per network forwarding engine)	3582 (per network forwarding engine)	3582 (per network forwarding engine)	3582 (per network forwarding engine)
IPv6 ACEs (unidimensional)	1536 (per network forwarding engine)	1536 (per network forwarding engine)	1792 (per network forwarding engine)	Not applicable	1792 (per network forwarding engine)	1792 (per network forwarding engine)	1792 (per network forwarding engine)	Not applicable
IPv4 and IPv6s ACEs	2048 IPv4 + 256 IPv6	2048 IPv4 + 256 IPv6	1024 IPv4 + 128 IPv6	Not applicable	1024 IPv4 + 128 IPv6	1024 IPv4 + 128 IPv6	1024 IPv4 + 128 IPv6	1024 IPv4, IPv6 not applicable
Interfaces with PBR policy	512	512	256	Not applicable	256	256	256	256
<b>VRRPv3</b>								
VRRPv3 groups per interface	255	255	255	Not applicable	255	255	255	255
VRRPv3 groups with default timers (1 s)	490	490	490	Not applicable	490	490	490	490
VRRPv3 groups with relaxed timers (3 s)	490	490	490	Not applicable	490	490	490	490

Feature	9500 Platform Verified Limit	9300 Platform Verified Limit	9200 Platform Verified Limit	9364C Platform Verified Limit	9300-EX Platform Verified Limit	9300-FX Platform Verified Limit	9300-FX2 Platform Verified Limit	9500 with X9700-EX /FX Verified Limit
Pathways with one VRRPv3 group with default timer (1 s)	489	489	489	Not applicable	489	489	489	489
VRRPv3 groups and pathways combined	490	490	490	Not applicable	490	490	490	490

<sup>16</sup> The limit of supported BFD sessions for each EoR line card is 75.

<sup>17</sup> The hash table is subject to collisions. Depending on the host route pattern, collisions might occur.

<sup>18</sup> The hash table is subject to collisions. Depending on the host route pattern, collisions might occur.

\*For the Cisco Nexus 9200 Platform switches, the default value for LPM unicast routes is 6000 (IPv4) or 1900 (IPv6). You can use the **hardware profile multicast max-limit lpm-entries 0** command to increase the number of IPv4 LPM unicast routes to 8000. The **hardware profile ipv6 lpm-entries maximum 0** command reserves the entire LPM table for IPv4. With this configuration, the IPv4 LPM scale is 14,000 (with 2000 reserved for multicast by default). This value can be increased to 16,000 with the **hardware profile multicast max-limit lpm-entries 0** command. The **hardware profile ipv6 lpm-entries maximum 4096** command reserves the entire LPM table for IPv6. With this configuration, the IPv6 LPM scale is 3900. When you allocate the entire table for IPv4 or IPv6 LPM unicast routes, the other address family cannot be used.



**Note**

- 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 need to be increased in multi-dimensional scenarios.

**Table 13: VXLAN Verified Scalability Limits (Unidimensional)**

Feature	9500 Platform Verified Limit	9300 Platform Verified Limit	9200 Platform Verified Limit	9300-EX Platform Verified Limit	9300-FX Platform Verified Limit	9300-FX2 Platform Verified Limit	9500 with X97xx-EX/FX Verified Limit
<b>IGMP snooping over VXLAN</b>							
VXLAN VLANs	Not applicable	1000	Not applicable	1000	1000	1000	Not applicable
VTEP Peers <sup>19</sup>	Not applicable	256	Not applicable	256	256	256	Not applicable
Underlay multicast groups	Not applicable	128	Not applicable	128	128	128	Not applicable
<b>VXLAN Flood and Learn</b>							
Virtual network identifiers (VNIs) or VXLAN-mapped VLANs	1000	2000	2000	2000	2000	2000	1000
Underlay multicast groups	128	128	128	128	128	128	128
Overlay MAC addresses	64,000	64,000	64,000	90,000	90,000	60,000	90,000
Remote VXLAN tunnel endpoints (VTEPs) <sup>20</sup>	256	256	256	256	256	256	256
Ingress replication peers	256	256	256	256	256	256	256
Ingress replication Layer 2 VNIs	1000	1000	1000	1000	1000	1000	1000
MAC addresses for ingress replication	64,000	64,000	64,000	90,000	90,000	90,000	90,000

Feature	9500 Platform Verified Limit	9300 Platform Verified Limit	9200 Platform Verified Limit	9300-EX Platform Verified Limit	9300-FX Platform Verified Limit	9300-FX2 Platform Verified Limit	9500 with X97xx-EX/FX Verified Limit
Port VLAN translations under an interface	100	100	Not applicable	500	500	500	100
Port VLAN translations in a switch	2000	2000	Not applicable	6000	6000	6000	2000
Static MAC addresses pointing to a remote VTEP	1000	1000	1000	1000	1000	1000	1000
VXLAN VLAN logical port VP count	7000	7000	Not applicable	Not applicable	Not applicable	Not applicable	Not applicable
VXLAN VLANs per FEX port (host interface)	75	75	Not applicable	75 <sup>21</sup>	Not applicable	Not applicable	Not applicable
Layer 2 routed VNIs for vPC-centralized gateway	450	450	450	450	450	450	450
IGMP groups	8192	8192	8192	8192	8192	8192	8192
<b>VXLAN BGP eVPN</b>							
Layer 2 VNIs	1000	2000	2000	2000, 4000 (with no Layer 3 VNIs)	2000, 4000 (with no Layer 3 VNIs)	2000, 4000 (with no Layer 3 VNIs)	1000
SVI with Distributed Anycast Gateway; Layer 2 VNI extended	1000	2000	2000 <sup>22</sup>	2000 <sup>23</sup>	2000	2000	1000
Layer 3 VNIs / VRFs <sup>24</sup>	750	900	900	900	900	900	750
Underlay multicast groups	128	128	128	128	128	128	128

Feature	9500 Platform Verified Limit	9300 Platform Verified Limit	9200 Platform Verified Limit	9300-EX Platform Verified Limit	9300-FX Platform Verified Limit	9300-FX2 Platform Verified Limit	9500 with X97xx-EX/FX Verified Limit
VTEPs	256	256	256	256	256	256	256
MAC addresses	64,000	64,000	64,000	90,000	90,000	90,000	90,000
IPv4 host routes	60,000	60,000	60,000	530,500	530,500	530,500	656,000
IPv6 host routes	7000	7000	7000	24,000	24,000	24,000	34,000
Overlay IPv4 LPM routes	12,000	12,000	8000	530,500	530,500	530,500	656,000
Overlay IPv6 LPM routes	7000	7000	2000	266,000 <sup>25</sup>	266,000 <sup>26</sup>	266,000 <sup>27</sup>	174,000 <sup>28</sup>
VXLAN VLAN logical port VP count	7000	10000	Not applicable	Not applicable	Not applicable	Not applicable	Not applicable
VXLAN VLANs per FEX port (host interface)	75	75	Not applicable	Not applicable <sup>29</sup>	Not applicable	Not applicable	Not applicable
IGMP groups	8192	8192	8192	8192	8192	8192	8192
<b>VXLAN BGP eVPN Ingress Replication</b>							
Layer 2 VNIs	1000	2000	2000	2000	2000	2000	1000
SVI with Distributed Anycast Gateway; Layer 2 VNI extended	1000	2000	2000 <sup>30</sup>	2000 <sup>31</sup>	2000	2000	1000
Layer 3 VNIs / VRFs <sup>32</sup>	750	900	900	900	900	900	750
VTEPs	256	256	256	256	256	256	256
MAC addresses	64,000	64,000	64,000	90,000	90,000	90,000	90,000
IPv4 host routes	32,000	32,000	32,000	530,500	530,500	530,500	656,000

Feature	9500 Platform Verified Limit	9300 Platform Verified Limit	9200 Platform Verified Limit	9300-EX Platform Verified Limit	9300-FX Platform Verified Limit	9300-FX2 Platform Verified Limit	9500 with X97xx-EX/FX Verified Limit
IPv6 host routes	7000	7000	7000	24,000	24,000	24,000	34,000
Overlay IPv4 LPM routes	12,000	12,000	8000	530,500	530,500	530,500	656,000
Overlay IPv6 LPM routes	7000	7000	2000	266,000 <sup>33</sup>	266,000 <sup>34</sup>	266,000 <sup>35</sup>	174,000 <sup>36</sup>
VXLAN VLAN logical port VP count	7000	7000	Not applicable	Not applicable	Not applicable	Not applicable	Not applicable
VXLAN VLANs per FEX port (host interface)	75	75	Not applicable	Not applicable	Not applicable	Not applicable	Not applicable
IGMP groups	8192	8192	8192	8192	8192	8192	8192

<sup>19</sup> In case of IR, each VNI can have a max of 64 peers.

<sup>20</sup> In case of IR, each VNI can have a max number of 64 peers

<sup>21</sup> This is the limit for the Cisco Nexus 93180YC-EX and other fiber based switches. All copper based 9300-EX switches are not applicable.

<sup>22</sup> Only 1900 SVI are supported if dual stack is used/IPv6 is used.

<sup>23</sup> Only 1900 SVI are supported if dual stack is used/IPv6 is used.

<sup>24</sup> ECMP objects are not shared across multiple VRFs.

<sup>25</sup> All /64 routes + 4000 for non /64 routes.

<sup>26</sup> All /64 routes + 4000 for non /64 routes.

<sup>27</sup> All /64 routes + 4000 for non /64 routes.

<sup>28</sup> All /64 routes + 4000 for non /64 routes.

<sup>29</sup> This particular combination has not been validated but the feature is supported.

<sup>30</sup> Only 1900 SVI are supported if dual stack is used/IPv6 is used.

<sup>31</sup> Only 1900 SVI are supported if dual stack is used/IPv6 is used.

<sup>32</sup> ECMP objects are not shared across multiple VRFs.

<sup>33</sup> All /64 routes + 4000 for non /64 routes.

<sup>34</sup> All /64 routes + 4000 for non /64 routes.

<sup>35</sup> All /64 routes + 4000 for non /64 routes.

<sup>36</sup> All /64 routes + 4000 for non /64 routes.

**Table 14: Tetration Verified Scalability Limits (Unidimensional)**

	<b>92160YC-X Verified Limit</b>	<b>9300-EX Platform Verified Limit</b>	<b>9300-FX Platform Verified Limit</b>
TCAM size	1024 entries	1024 entries	1024 entries
	IPv4 – 4 entries per rule (TCP, UDP, ICMP, and IP)	IPv4 – 2 entries per rule (ICMP and IP)	IPv4 – 2 entries per rule (ICMP and IP)
	IPv6 – 16 entries per rule (4 entries per TCP, UDP, ICMPv6, and IPv6 for a total of 16 entries)	IPv6 – 8 entries per rule (4 entries per ICMP and IPv6 for a total of 8 entries)	IPv6 – 8 entries per rule (4 entries per ICMP and IPv6 for a total of 8 entries)
	(24 entries out of 1000 is consumed for default)	(24 entries out of 1000 is consumed for default)	(24 entries out of 1000 is consumed for default)
TCAM scale	250 (IPv4) or 62 (IPv6)	500 (IPv4) or 125 (IPv6)	500 (IPv4) or 125 (IPv6)
VRF match	Not applicable	Not applicable	Not applicable

The entire Cisco Tetration Analytics documentation set is available at the following URL:

<https://www.cisco.com/c/en/us/support/data-center-analytics/tetration-analytics/tsd-products-support-series-home.html>

## Deployment Case Studies

This section provides sample topologies for some common deployments. For each topology, the scalability numbers are the limits with all of the listed features enabled at the same time.




---

**Attention** These numbers are not the maximum verified values if each feature is viewed in isolation. For these numbers, see the "Verified Scalability Limits" section.

---

### Layer 2/Layer 3 Aggregation Topology (Max-Host Routing Mode)

This Layer 2/Layer 3 aggregation topology consists of Cisco Nexus 9508 switches as virtual port channel (vPC) aggregation pairs. These aggregation nodes are fully loaded with N9K-X9564TX, N9K-X9564PX, and N9K-X9636PQ line cards. The N9K-X9636PQ line cards are used in normal mode and breakout mode. Cisco Nexus 9396PX and 93128TX switches are used as top-of-rack units with Cisco Nexus 3000 Series switches to achieve the desired vPC scale.

The Cisco Nexus 9508 switch is also used as a core Layer 3 node that connects to a pair of vPC aggregation nodes. The focus of the topology is to test IPv4 ARP, IPv6 neighbor discovery (ND), and Layer 2 scalability and other routing, switching, and Layer 4 through Layer 7 features for management and operations. All Layer 3 interfaces are configured for dual stack, and the traffic is dual stack for all VLANs.

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

**Table 15: Layer 2/Layer 3 Aggregation Topology (Max-Host Routing Mode)**

<b>Feature</b>	<b>9508 Verified Limit (Max-Host Routing Mode)</b>
Fully loaded chassis	1 N9K-X9636PQ, 1 N9K-X9564TX, 2 N9K-X9564PX, 1 N9K-X9432PQ, 1 N9K-X9536PQ
Physical interfaces enabled	276
Multicast S,G routes	653
Multicast *,G routes	500
IPv4 unicast routes (LPM)	5000
IPv6 unicast routes (LPM)	850
IPv4 ARP	65,000
IPv6 ND	40,000
MAC addresses	90,000
VLANs	490
vPCs*	200
OSPFv2 neighbors	20
OSPFv3 neighbors	4
BGP (IPv4) neighbors	65
BGP (IPv6) neighbors	65
SVIs	490
STP logical ports	2800 (RPVST)
HSRP VLANs (IPv4/IPv6)	490
Virtual ports	700
Port channel links	8

\* The number of VLANs per vPC supported should be within the MST or RPVST virtual port count specified in this table, depending on the topology.

### **Layer 2/Layer 3 Aggregation Topology (Default Routing Mode)**

This Layer 2/Layer 3 aggregation topology consists of Cisco Nexus 9516 switches as virtual port channel (vPC) aggregation pairs. These aggregation nodes are fully loaded with N9K-X9564TX, N9K-X9564PX, and N9K-X9536PQ line cards. The chassis is fully loaded with five line cards configured for breakout mode. The Cisco Nexus 9396PX and 93128TX switches are used as top-of-rack units with Cisco Nexus 3000 Series switches to achieve the desired vPC scale. The Cisco Nexus 9516 nodes are running in default routing mode. The Cisco Nexus 3164Q switch is also used as a core Layer 3 node that connects to a pair of vPC aggregation nodes.



The focus of the topology is to test IPv4 ARP, IPv6 neighbor discovery (ND), Layer 2 scalability, IPv4 and IPv6 LPM routing, Layer 2 and Layer 3 multicast routing for IPv4, and Layer 4 through Layer 7 features for management and operations. All Layer 3 interfaces are configured for dual stack, and the traffic is dual stack for all VLANs.

In the following table, the Verified Limit column lists the verified scaling capabilities with all listed features enabled at the same time. These numbers are not the maximum verified values if each feature is viewed in isolation.

**Table 16: Layer 2/Layer 3 Aggregation Topology (Default Routing Mode)**

<b>Feature</b>	<b>9516 Switch Verified Limit (Default Routing Mode)</b>	<b>9300 Platform Verified Limit (Default Routing Mode)</b>
Chassis configuration	5 N9K-X9432PQ line cards 4 N9K-X9464PX line cards 3 N9K-X9536PQ line cards 3 N9K-X9464TX line cards 1 N9K-X9564TX line card	9372
Physical ports	1335	50
vPCs	303	24
SVIs	450	450
VRFs	100	100
IPv4 ARP	40,000	40,000
IPv6 ND	10,000	10,000
STP logical ports	10,000	6000
BGP neighbors (IPv4 + IPv6)	502 + 502	502 + 502
IPv4 LPM routes	50,000	6000
IPv6 LPM routes	10,000	1000
BFD (IPv4 + IPv6)	300	102
IGP OSPFv2 neighbors	502	502
IGP OSPFv3 neighbors	502	502
HSRP (IPv4 + IPv6)	450 + 450	450 + 450
IGMP groups	2000	2000
Multicast *,G routes	2000	2000
Multicast S,G routes	8000	6000
Tracking objects	450	450

Feature	9516 Switch Verified Limit (Default Routing Mode)	9300 Platform Verified Limit (Default Routing Mode)
VLANs	500	500
PIM neighbors	502	502
MAC addresses	60,000	60,000
Network address translation (NAT)	Not applicable	756
sFlow	256	32

## FEX System Topology

The FEX 9500 multi-dimensional scale topology consists of Cisco Nexus 9508 switches as virtual port channel (vPC) pairs. Each switch has multiple X9564PX line cards. Each switch has 32 FEX uplinks connected to them. The FEX 9300 multi-dimensional scale topology consists of two Cisco Nexus 9396PX switches used in vPC mode along with 16 FEX uplinks connected to each switch. Multiple FEXs of type Nexus 2248TP-E, 2232PP, 2248PQ, and 2348UPQ are used.

The switches are used at the Layer 2 and Layer 3 boundary and are also configured as VXLAN VTEPs. The FEX host ports are operating as Layer 2 ports. The switches are configured as gateways with the use of SVI interfaces.

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

**Table 17: FEX System Topology**

Feature	9500 Platform Verified Limit	9300 Platform Verified Limit
Fabric Extenders	32	16
Up interfaces	1100	560
Port channels	426	256
vPC members	390	360
VLANs	744	416
PVLAN VLANs	56	56
Secondary VLANs per primary VLAN	25	25
MAC addresses	45,000	25,000
HSRP	365	365
ARP	12,000	10,000
Neighbor discovery (ND)	5000	5000
Multicast (*,G)	4000	4000

Feature	9500 Platform Verified Limit	9300 Platform Verified Limit
Multicast (S,G)	4000	4000

## Multicast System Topology

Two Cisco Nexus 9508 switches are configured as vPC peers in one domain, and two Cisco Nexus 9372PX switches are configured as vPC peers in the other domain. The chassis are fully loaded with N9K-X9432PQ, N9K-X9464PX, N9K-X9536PQ, N9K-X9564PX, N9K-X9564TX, and N9K-X9636PQ line cards. eBGP routing is used to connect these two PIM domains. OSPF is used as IGP in one domain, and EIGRP is configured in the other domain. This setup is configured with multiple rendezvous points (RPs) to serve different multicast group ranges. BSR is used to advertise RP information in both of these PIM domains. PIM anycast is used in one domain, and MSDP anycast is used in the other domain for redundancy and load balancing. Static RP configuration is also used for a range of multicast groups.

The Cisco Nexus 9516 and Cisco Nexus 7000 Series switches are used as Layer 3 core routers in one domain. The Cisco Nexus 3164Q switches are used as Layer 3 core routers in the other domain. This topology also includes the Cisco Nexus 9396PX, Cisco Nexus 9372PX, and Cisco Nexus 3016/3064T switches in the access layer.

In addition to including Layer 2/Layer 3 IPv4 multicast routing, this topology also covers IPv4 and IPv6 host and LPM routing and Layer 2 unicast forwarding. All interfaces are configured for dual stack.

In the following table, the Verified Limit column lists the verified scaling capabilities with all listed features enabled at the same time. These numbers are not the maximum verified values if each feature is viewed in isolation.

**Table 18: Multicast System Topology**

Feature	9500 Platform Verified Limit	9300 Platform Verified Limit
Chassis configuration	N9K-X9636PQ, N9K-X9536PQ, N9K-X9564PX, N9K-X9564TX, N9K-X9432PQ, N9K-X9464PX, N9K-X9432PQ, C3164PQ	C9372PX, C9396PX, C3164PQ
Multicast S,G routes	17,500	5000
Multicast *,G routes	2500 (IGMP) 12500 (snooping)	500 (IGMP) 2500 (snooping)
Sources	2000, 200, 100, 40, 10, 3, 2, 1	2000, 200, 100, 40, 10, 3, 2, 1
Replications	40	20
ECMPs	16	8
SVIs	200	200
HSRP/VRRP	200 HSRP	100 VRRP
MAC addresses	40,000	10,000
ARP	20,000	4000
Unicast LPM IPv4 routes	20,000	4000
Unicast LPM IPv6 routes	10,000	1000

Feature	9500 Platform Verified Limit	9300 Platform Verified Limit
IPv4 ARP	18,000	4000
IPv6 ND	4000	2000
MSDP peers (fully mesh)	4	4
Anycast RPs (MSDP and PIM anycast) <sup>37</sup>	2 MSDP	2 PIM anycast
IPv4 multicast routes with PIM bidirectional groups	8000	8000

<sup>37</sup> This multicast system topology consists of two multicast PIM domains. The Multicast Source Discovery Protocol (MSDP) is used to exchange multicast source information between these two domains.

### VXLAN BGP/eVPN iBGP Centric Topology

This VXLAN BGP/eVPN iBGP centric topology consists of Cisco Nexus 9300 and 9500 Platform switches acting as VXLAN vPC tunnel endpoints (VTEPs) and VXLAN non-vPC VTEPs. VXLAN VTEPs establish iBGP sessions to a Cisco Nexus 9508 switch (route reflector) acting as a spine node. VXLAN-distributed anycast gateway SVIs are configured for dual stack, and the traffic is dual stack.

The focus of this topology is to test VXLAN overlay network scale and underlay Layer 2 switching and other routing, multicast, and Layer 4 through Layer 7 features for management and operations. Underlay PIM neighbors and IS-IS adjacency were tested with the default timer and Bidirectional Forwarding Detection (BFD) enabled on all links.

In the following table, the Verified Limit column lists the verified scaling capabilities with all listed features enabled at the same time. These numbers are not the maximum verified values if each feature is viewed in isolation.

**Table 19: VXLAN BGP/eVPN iBGP Centric Topology**

Feature	9500 Platform Verified Limit	9300 Platform Verified Limit	9200 Platform Verified Limit	9300-EX Platform Verified Limit	9300-FX Platform Verified Limit	9364C Platform Verified Limit	9500 with X9700-EX Verified Limit
System Routing Template	default	default	default	default <sup>38</sup>	default <sup>39</sup>	Not applicable	default <sup>40</sup>
VXLAN VTEPs	128	128	128	128	128	128	128
VXLAN Layer 2 VNIs	1000	1000	1000	1000	1000	1000	1000
VXLAN Layer 3 VNIs/VRFs	500	500	500	500	500	500	500
VXLAN multicast groups	128	128	128	128	128	128	128

Feature	9500 Platform Verified Limit	9300 Platform Verified Limit	9200 Platform Verified Limit	9300-EX Platform Verified Limit	9300-FX Platform Verified Limit	9364C Platform Verified Limit	9500 with X9700-EX Verified Limit
VXLAN overlay MAC addresses	60,000	60,000	60,000	60,000	60,000	60,000	60,000
VXLAN overlay IPv4 host routes	60,000	60,000	60,000	60,000	60,000	60,000	60,000
VXLAN overlay IPv6 host routes	4000	4000	4000	4000	4000	Not applicable	4000
VXLAN overlay IGMP Snooping groups	2000 <sup>41</sup>	2000	2000	2000	2000	2000	2000
VXLAN IPv4 LPM routes	10000	10000	5120	5120	5120	5120	5120
VXLAN IPv6 LPM routes	2000	2000	1500	1500	1500	Not applicable	1500
VXLAN VLAN logical port VP count	5200	5200	5200	5200	5200	Not applicable	5200
VLANs on VTEP node	1700 (total VLANs) 1500 (VXLAN VLANs) 200 (non-VXLAN VLANs)	1700 (total VLANs) 1500 (VXLAN VLANs) 200 (non-VXLAN VLANs)	1700 (total VLANs) 1500 (VXLAN VLANs) 200 (non-VXLAN VLANs)	1700 (total VLANs) 1500 (VXLAN VLANs) 200 (non-VXLAN VLANs)	1700 (total VLANs) 1500 (VXLAN VLANs) 200 (non-VXLAN VLANs)	1700 (total VLANs) 1500 (VXLAN VLANs) 200 (non-VXLAN VLANs)	
MST instances	40	40	40	40	40	Not applicable	40
STP logical ports	3500	3500	3500	3500	3500	Not applicable	3500
vPC port channels	50	20	20	20	20	Not applicable	20
Underlay IS-IS neighbors	64	32	32	32	32	Not applicable	32
Underlay PIM neighbors	200	200	200	200	200	Not applicable	200

<b>Feature</b>	<b>9500 Platform Verified Limit</b>	<b>9300 Platform Verified Limit</b>	<b>9200 Platform Verified Limit</b>	<b>9300-EX Platform Verified Limit</b>	<b>9300-FX Platform Verified Limit</b>	<b>9364C Platform Verified Limit</b>	<b>9500 with X9700-EX Verified Limit</b>
Underlay HSRP groups for regular VLANs	200	200	200	200	200	Not applicable	200
Underlay vPC SVIs	200	200	200	200	200	Not applicable	200

<sup>38</sup> The vxlan-routing-template needs to be configured on 7.0(3)I5(1).

<sup>39</sup> The vxlan-routing-template needs to be configured on 7.0(3)I5(1).

<sup>40</sup> The vxlan-routing-template needs to be configured on 7.0(3)I5(1).

<sup>41</sup> IGMP Snooping on vxlan vlan on 9500 series switch supported from 7.0(3)I5(2) release onwards.

THE SPECIFICATIONS AND INFORMATION REGARDING THE PRODUCTS IN THIS MANUAL ARE SUBJECT TO CHANGE WITHOUT NOTICE. ALL STATEMENTS, INFORMATION, AND RECOMMENDATIONS IN THIS MANUAL ARE BELIEVED TO BE ACCURATE BUT ARE PRESENTED WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. USERS MUST TAKE FULL RESPONSIBILITY FOR THEIR APPLICATION OF ANY PRODUCTS.

THE SOFTWARE LICENSE AND LIMITED WARRANTY FOR THE ACCOMPANYING PRODUCT ARE SET FORTH IN THE INFORMATION PACKET THAT SHIPPED WITH THE PRODUCT AND ARE INCORPORATED HEREIN BY THIS REFERENCE. IF YOU ARE UNABLE TO LOCATE THE SOFTWARE LICENSE OR LIMITED WARRANTY, CONTACT YOUR CISCO REPRESENTATIVE FOR A COPY.

The Cisco implementation of TCP header compression is an adaptation of a program developed by the University of California, Berkeley (UCB) as part of UCB's public domain version of the UNIX operating system. All rights reserved. Copyright © 1981, Regents of the University of California.

NOTWITHSTANDING ANY OTHER WARRANTY HEREIN, ALL DOCUMENT FILES AND SOFTWARE OF THESE SUPPLIERS ARE PROVIDED "AS IS" WITH ALL FAULTS. CISCO AND THE ABOVE-NAMED SUPPLIERS DISCLAIM ALL WARRANTIES, EXPRESSED OR IMPLIED, INCLUDING, WITHOUT LIMITATION, THOSE OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NON-INFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE.

IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THIS MANUAL, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

Any Internet Protocol (IP) addresses and phone numbers used in this document are not intended to be actual addresses and phone numbers. Any examples, command display output, network topology diagrams, and other figures included in the document are shown for illustrative purposes only. Any use of actual IP addresses or phone numbers in illustrative content is unintentional and coincidental.

This product includes cryptographic software written by Eric Young (eay@cryptsoft.com).

This product includes software developed by the OpenSSL Project for use in the OpenSSL Toolkit. (<https://www.openssl.org>)

This product includes software written by Tim Hudson (tjh@cryptsoft.com).

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: <https://www.cisco.com/go/trademarks>. Third-party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1110R)

© 2018–2021 Cisco Systems, Inc. All rights reserved.



**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](http://www.cisco.com/go/offices).