



Cisco Nexus 9000 Series NX-OS Verified Scalability Guide, Release 7.0(3)I2(1)

[Verified Scalability Limits](#) 2

Verified Scalability Limits

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

Introduction

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)I2(1). 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) Verified Scalability Limits (Unidimensional)

Feature	9500 Series Verified Limit ¹	9300 Series Verified Limit ²
Fabric Extenders and Fabric Extender server interfaces	32 and 1536	16 and 768
VLANs per Fabric Extender	2000 (across all Fabric Extenders)	2000 (across all Fabric Extenders)
VLANs per Fabric Extender server interface ³	75	75
Port channels	426	378
Unique Fabric Extenders per Cisco Nexus 9500 Series supported line card	12	Not applicable

¹ 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.

² 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.

³ 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: Interfaces Verified Scalability Limits (Unidimensional)

Feature	9500 Series Verified Limit	9300 Series Verified Limit
DHCP clients	5	5
Generic routing encapsulation (GRE) tunnels	8	8
Port channel links	32	32
SVIs	490 (with HSRP), 1500 (without HSRP)	450 (with HSRP)
vPCs	300	48
Static network address translation (NAT)	Not applicable	1023
Dynamic network address translation (NAT)	Not applicable	1023
Static twice network address translation (NAT)	Not applicable	768
Dynamic twice network address translation (NAT)	Not applicable	1023

Table 3: Label Switching Verified Scalability Limits (Unidimensional)

Feature	9500 Series Verified Limit	9300 Series Verified Limit
Forwarding Equivalence Classes (FECs)	128	128
Equal-cost multipaths (ECMPs)	32	16
Flex counters for static MPLS in egress direction	4000	4000
Flex counters per adjacency	2	2
Adjacencies	1024	1024



Note The maximum number of FECs and ECMPs cannot be configured at the same time. For example, if you have 128 FECs and all of those FECs have 8 ECMPs, you will have $128 * 8 = 1024$ adjacencies, so egress statistics will be supported for all. In contrast, if you have 100 FECs and all of those FECs have 16 ECMPs, you will have $100 * 16 = 1600$ adjacencies. Because a maximum of 1024 adjacencies are supported, the statistics might not work as expected.

Table 4: Layer 2 Switching Verified Scalability Limits (Unidimensional)

Feature	9500 Series Verified Limit	9300 Series Verified Limit
MST instances	64	64

Feature	9500 Series Verified Limit	9300 Series Verified Limit
MST virtual ports	85,000	48,000
RPVST virtual ports	22,000	12,000
VLANs	3967 (the remaining 127 VLANs are reserved)	3967 (the remaining 127 VLANs are reserved)
VLANs in RPVST mode	500	500
Total number of VLANs × ports with switchport isolated	96,000	50,000
Private VLANs (PVLANS)		
Primary VLANs	16	16
Secondary VLANs	20	20
Ports in Community host mode	40	40
Ports in isolated host mode	20	40
Ports in isolated trunk host mode	22	40
Ports in promiscuous mode	48	5
Ports in promiscuous trunk mode	80	5
PVLANS allowed on a PVLAN port	16	16



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 5: Multicast Routing Verified Scalability Limits (Unidimensional)

Feature	9500 Series Verified Limit	9300 Series Verified Limit
IPv4 multicast routes	32,000 (Layer 2 + Layer 3)	8000 (Layer 2 + Layer 3)
Outgoing interfaces (OIFs)	40 (SVI + physical Layer 3) or 256 (physical Layer 3)	40 (SVI + physical Layer 3)



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 6: Programmability Verified Scalability Limits (Unidimensional)

Feature	9500 Series Verified Limit	9300 Series Verified Limit
OpenFlow		
OpenFlow ports	Not applicable	96
OpenFlow Layer 2 flows	Not applicable	32,000
OpenFlow Layer 3 flows	Not applicable	3000

Table 7: Security Verified Scalability Limits (Unidimensional)

Feature	9500 Series Verified Limit	9300 Series Verified Limit
DHCP snooping bindings	2048	2048
IPv4 ingress access control entries (ACEs)	3072 (per network forwarding engine)	3072 (per network forwarding engine)
IPv4 egress access control entries (ACEs)	768 (per network forwarding engine)	768 (per network forwarding engine)
IPv6 ingress access control entries (ACEs)	1536 (per network forwarding engine)	1536 (per network forwarding engine)
IPv6 egress access control entries (ACEs)	256 (per network forwarding engine)	256 (per network forwarding engine)



Note The ACE scalability limits also apply to policy-based ACLs (PBACLs).

Table 8: System Management Verified Scalability Limits (Unidimensional)

Feature	9500 Series Verified Limit	9300 Series Verified Limit
MPLS Stripping		
Labels	12,000	12,000
PTP		

Feature	9500 Series Verified Limit	9300 Series Verified Limit
10G physical ports enabled for PTP	44	44
sFlow		
sFlow ports	256	64
SPAN and ERSPAN		
Configurable SPAN or ERSPAN sessions	32	4
Active SPAN or ERSPAN sessions ⁴	4 to 32, based on the number of line cards and the session configuration	4
Active localized SPAN or ERSPAN sessions per line card ⁵	4	4
Source interfaces per SPAN or ERSPAN session (Rx and Tx, Rx, or Tx)	48	48
Destination interfaces per SPAN session	1 (physical interface)	1 (physical interface)
Source VLANs per SPAN or ERSPAN session	32	32
TAP aggregation		
Redirect interfaces in the redirect port list	12	12
Redirect port lists (or fan outs) per system	100	100

⁴ 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.

⁵ 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 9: Unicast Routing Verified Scalability Limits (Unidimensional)

Feature	9500 Series Verified Limit	9300 Series Verified Limit
Unicast Routing		
BFD sessions (echo mode)	512	256
BGP neighbors	2000	512

Feature	9500 Series Verified Limit	9300 Series Verified Limit
EIGRP routes	20,000	20,000
EIGRP neighbors	512	256
HSRP groups	490	490
IPv4 ARP	48,000	48,000
IPv4 host routes	208,000 (hash table and there will be more collisions after 80%)	208,000 (hash table and there will be more collisions after 80%)
IPv6 host routes	104,000 (hash table and there will be more collisions after 80%)	104,000 (hash table and there will be more collisions after 80%)
IPv6 ND	48,000	48,000
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)
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)
IPv4 and IPv6 unicast routes (LPM) in 64-bit ALPM routing mode	128,000 (IPv4) 80,000 (IPv6)	128,000 (IPv4) 80,000 (IPv6)
IS-ISv4 adjacencies (either L1, L2, or sum of L1 and L2 with default timers)	255	255
IS-ISv4 BFD sessions (with default timers)	255	255
IS-ISv4 routes	10,000	10,000
IS-ISv4 network type	Point to point, broadcast	Point to point, broadcast
MAC addresses	90,000	90,000
OSPFv2 neighbors	1000	256
OSPFv3 neighbors	1000	256
VRFs	1000	1000
VRRP groups per interface or I/O module	250	250
Policy-based routing (PBR)		
Configured sequences per policy	256	256

Feature	9500 Series Verified Limit	9300 Series Verified Limit
Next-hop addresses per policy	32	32
IPv4 ACEs (unidimensional)	3072 (per network forwarding engine)	3072 (per network forwarding engine)
IPv6 ACEs (unidimensional)	1536 (per network forwarding engine)	1536 (per network forwarding engine)
IPv4 and IPv6s ACEs	2048 IPv4 + 256 IPv6	2048 IPv4 + 256 IPv6
Interfaces with PBR policy	512	512
VRRPv3		
VRRPv3 groups per interface	255	255
VRRPv3 groups with default timers (1 s)	490	490
VRRPv3 groups with aggressive timers (100 ms)	290	290
VRRPv3 groups with relaxed timers (3 s)	490	490
Pathways with one VRRPv3 group with default timer (1 s)	489	489
VRRPv3 groups and pathways combined	490	490



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.



Note 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.



Note 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 10: VXLAN Verified Scalability Limits (Unidimensional)

Feature	9500 Series Verified Limit	9300 Series Verified Limit
VXLAN Flood and Learn		
Virtual network identifiers (VNIs) or VXLAN-mapped VLANs	1000	1000
Underlay multicast groups	128	128
Overlay MAC addresses	64,000	64,000
Remote VXLAN tunnel endpoints (VTEPs)	256	256
Ingress replication peers	256	256
Ingress replication Layer 2 VNIs	1000	1000
MAC addresses for ingress replication	64,000	64,000
Port VLAN translations under an interface	100	100
Port VLAN translations in a switch	2000	2000
Static MAC addresses pointing to a remote VTEP	1000	1000
VXLAN VLAN logical port VP count	7000	7000
VXLAN VLANs per FEX port (host interface)	75	75
Layer 2 routed VNIs for vPC-centralized gateway	450	450
VXLAN BGP eVPN		
Layer 2 VNIs	1000	1000
Layer 3 VNIs / VRFs ⁶	750	900
Underlay multicast groups	128	128
VTEPs	256	256
MAC addresses	64,000	64,000
IPv4 host routes	32,000	32,000
IPv6 host routes	7,000	7,000
Overlay IPv4 LPM routes	12,000	12,000

Feature	9500 Series Verified Limit	9300 Series Verified Limit
Overlay IPv6 LPM routes	7000	7000
VXLAN VLAN logical port VP count	7000	7000
VXLAN VLANs per FEX port (host interface)	75	75
VXLAN BGP eVPN Ingress Replication		
Layer 2 VNIs	1000	1000
Layer 3 VNIs / VRFs ⁷	750	900
VTEPs	128	128
MAC addresses	64,000	64,000
IPv4 host routes	32,000	32,000
IPv6 host routes	7000	7000
Overlay IPv4 LPM routes	12,000	12,000
Overlay IPv6 LPM routes	7000	7000
VXLAN VLAN logical port VP count	7000	7000
VXLAN VLANs per FEX port (host interface)	75	75

⁶ ECMP objects are not shared across multiple VRFs.

⁷ ECMP objects are not shared across multiple VRFs.

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

Feature	9508 Verified Limit (Max-Host Routing Mode)
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 12: Layer 2/Layer 3 Aggregation Topology (Default Routing Mode)

Feature	9516 Series Verified Limit (Default Routing Mode)	9300 Series Verified Limit (Default Routing Mode)
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

Feature	9516 Series Verified Limit (Default Routing Mode)	9300 Series Verified Limit (Default Routing Mode)
IGMP groups	2000	2000
Multicast *,G routes	2000	2000
Multicast S,G routes	8000	6000
Tracking objects	450	450
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 13: 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

Feature	9500 Platform Verified Limit	9300 Platform Verified Limit
HSRP	365	365
ARP	12,000	10,000
Neighbor discovery (ND)	5000	5000
Multicast (*,G)	4000	4000
Multicast (S,G)	4000	4000

Multicast System Topology

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.

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 14: Multicast System Topology

Feature	9500 Series Verified Limit	9300 Series 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

Feature	9500 Series Verified Limit	9300 Series Verified Limit
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
IPv4 ARP	18,000	4000
IPv6 ND	4000	2000
MSDP peers (fully mesh)	4	4
Anycast RPs (MSDP and PIM anycast)	2 MSDP	2 PIM anycast
IPv4 multicast routes with PIM bidirectional groups	2500	2500

VXLAN BGP/eVPN iBGP Centric Topology Using Multicast

This VXLAN BGP/eVPN iBGP centric topology consists of Cisco Nexus 9300 and 9500 Series 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 15: VXLAN BGP/eVPN iBGP Centric Topology Using Multicast

Feature	9500 Series Verified Limit	9300 Series Verified Limit
VXLAN VTEPs	128	128
VXLAN Layer 2 VNIs	1000	1000
VXLAN Layer 3 VNIs/VRFs	500	500
VXLAN multicast groups	25	25
VXLAN overlay MAC addresses	60,000	60,000
VXLAN overlay IPv4 host routes	60,000	60,000

Feature	9500 Series Verified Limit	9300 Series Verified Limit
VXLAN overlay IPv6 host routes	4000	4000
VXLAN IPv4 LPM routes	10000	10000
VXLAN IPv6 LPM routes	2000	2000
VXLAN VLAN logical port VP count	5200	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)
MST instances	40	40
STP logical ports	3500	3500
vPC port channels	50	20
Underlay IS-IS neighbors	64	32
Underlay PIM neighbors	200	200
Underlay HSRP groups for regular VLANs	200	200
Underlay vPC SVIs	200	200
Underlay multicast S,G routes	4000 (ASM) 1000 (SSM)	4000 (ASM) 1000 (SSM)

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)

© 2015–2018 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.