



Cisco Nexus 3232C and 3264Q NX-OS Verified Scalability Guide, Release 7.0(3)I4(1)

[Verified Scalability Limits](#) 2

[Introduction](#) 2

[Verified Scalability Limits \(Unidimensional\)](#) 2

Revised: May 23, 2016,

Verified Scalability Limits

This document describes the Cisco NX-OS configuration limits for the Cisco Nexus 3232C and 3264Q switches.

Introduction

The values provided in this guide should not be interpreted as theoretical system limits for Cisco Nexus 3232C or 3264Q 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 tables in this section list the unidimensional verified scalability limits for Cisco NX-OS Release 7.0(3)I4(1) on the Cisco Nexus 3232C and 3264Q switches. 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: Interfaces Verified Scalability Limits (Unidimensional)

Feature	3232C Verified Limit	3264Q Verified Limit
DHCP clients per switch	10 (IPv4) + 10 (IPv6)	10 (IPv4) + 10 (IPv6)
BFD sessions	256	256
Port channel links	32	32
SVIs	400	400
vPCs	100	60
Static network address translation (NAT)	1023	1023
Dynamic network address translation (NAT)	1023	1023
Static twice network address translation (NAT)	768	768
Dynamic twice network address translation (NAT)	1023	1023

Table 2: Label Switching Verified Scalability Limits (Unidimensional)

Feature	3232C Verified Limit	3264Q Verified Limit
Forwarding Equivalence Classes (FECs)	128	128
Equal-cost multipaths (ECMPs)	16	16
FECs * ECMPs	1000	1000
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 3: Layer 2 Switching Verified Scalability Limits (Unidimensional)

Feature	3232C Verified Limit	3264Q Verified Limit
MST instances	64	64
MST virtual ports	48,000	48,000
RPVST virtual ports	12,000	12,000
VLANs	3900	3900
VLANs in RPVST mode	500	500
Total number of VLANs × ports with switchport isolated (3967 VLANs x 48 ports)	190,000	190,000



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.

Table 4: Multicast Routing Verified Scalability Limits (Unidimensional)

Feature	3232C Verified Limit	3264Q Verified Limit
IPv4 multicast routes	8000 (Layer 2 + Layer 3)	8000 (Layer 2 + Layer 3)
Outgoing interfaces (OIFs)	40 (SVI + physical Layer 3)	40 (SVI + physical Layer 3)
IGMP snooping groups	8000	8000
PIM neighbors	250	250



Note These multicast scalability numbers are for Layer 2 and Layer 3.



Note Graceful restart is not supported when unicast or multicast aggressive timers are configured at any scale.

Table 5: Programmability Verified Scalability Limits (Unidimensional)

Feature	3232C Verified Limit	3264Q Verified Limit
OpenFlow		
OpenFlow Layer 3 flows	760	760
OpenFlow Lite Layer 3 flows	2001	2001

Table 6: Security Verified Scalability Limits (Unidimensional)

Feature	3232C Verified Limit	3264Q Verified Limit
DHCP snooping bindings	2048	2048
IPv4 ingress ACLs (lite TCAM region)	2304 (per network forwarding engine)	2304 (per network forwarding engine)
IPv4 ingress ACLs (non-lite TCAM region)	768 (per network forwarding engine)	768 (per network forwarding engine)
IPv4 egress ACLs (non-lite TCAM region)	768 (per network forwarding engine)	768 (per network forwarding engine)
IPv6 ingress ACLs (non-lite TCAM region)	768 (per network forwarding engine)	768 (per network forwarding engine)
IPv6 egress ACLs (non-lite TCAM region)	256 (per network forwarding engine)	256 (per network forwarding engine)



Note These ACL scalability numbers were verified with all TCAM regions freed up except the CoPP system and ingress system TCAM regions.

Table 7: System Management Verified Scalability Limits (Unidimensional)

Feature	3232C Verified Limit	3264Q Verified Limit
PTP		
10G physical ports enabled for PTP	44	44
sFlow		
sFlow ports	64	64
SPAN and ERSPAN		
Configurable SPAN and ERSPAN sessions	32	32
Active SPAN and ERSPAN sessions ^{1 2}	4	4
Active localized SPAN or ERSPAN sessions ³	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

- ¹ A single forwarding engine instance supports four SPAN or ERSPAN sessions. If the first three sessions have bidirectional sources, the fourth session might have hardware resources only for Rx sources, depending on the SPAN or ERSPAN source's forwarding engine instance mappings.
- ² If the ERSPAN ACL contains access control entries (ACEs) with the **set-erspan-gre-proto** or **set-erspan-dscp** action, then only one ERSPAN session can be up.
- ³ The number of SPAN or ERSPAN sessions reduces to two if the same interface is configured as the bidirectional source in more than one session.

Table 8: Unicast Routing Verified Scalability Limits (Unidimensional)

Feature	3232C Verified Limit	3264Q Verified Limit
Unicast Routing		
BGP neighbors	512 (IPv4 only) 512 (IPv6 only) 256 (IPv4) + 256 (IPv6)	512 (IPv4 only) 512 (IPv6 only) 256 (IPv4) + 256 (IPv6)

Feature	3232C Verified Limit	3264Q Verified Limit
EIGRP neighbors	256 (IPv4) 256 (IPv6) 128 (IPv4) + 128 (IPv6)	256 (IPv4) 256 (IPv6) 128 (IPv4) + 128 (IPv6)
EIGRP routes	20,000	20,000
HSRP groups	400	400
IPv4 ARP	Default system routing mode: 32,000 ALPM routing mode: 8000	Default system routing mode: 32,000 ALPM routing mode: 8000
IPv4 host routes	Default system routing mode: 104,000 (hash table and there will be more collisions after 80%) ALPM routing mode: 128,000 (the host routes programmed in the ALPM table)	Default system routing mode: 104,000 (hash table and there will be more collisions after 80%) ALPM routing mode: 128,000 (the host routes programmed in the ALPM table)
IPv6 host routes	Default system routing mode: 52,000 ALPM routing mode: 4000	Default system routing mode: 52,000 ALPM routing mode: 4000
IPv6 ND	Default system routing mode: 32,000 ALPM routing mode: 4000	Default system routing mode: 32,000 ALPM routing mode: 4000
IPv4 unicast routes (LPM)	Default system routing mode: 12,000 ALPM routing mode: 128,000	Default system routing mode: 12,000 ALPM routing mode: 128,000
IPv6 unicast routes (LPM)	Default system routing mode: 6000 (≤64) + 1000 (65-127) ALPM routing mode: 10,000 (15,000 with TCAM ALPM carving)	Default system routing mode: 6000 (≤64) + 1000 (65-127) ALPM routing mode: 10,000 (15,000 with TCAM ALPM carving)
MAC addresses	30,000	30,000
OSPFv2 neighbors	256	256
OSPFv3 neighbors	256	256
VRFs	2000	2000
VRRP groups	250	250



Note Graceful restart 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 multidimensional scenarios.

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. (<http://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: <http://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)

© 2016 Cisco Systems, Inc. All rights reserved.



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

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV
Amsterdam, The Netherlands

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.