Cisco Nexus 3164Q NX-OS Verified Scalability Guide, Release 7.0(3)I1(2)

First Published: May 15, 2015

Verified Scalability Limits

This chapter describes the Cisco NX-OS configuration limits for the Cisco Nexus 3164Q switch.

Introduction

The values provided in this guide should not be interpreted as theoretical system limits for Cisco Nexus 3164Q 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)I1(2) on the Cisco Nexus 3164Q switch. 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	3164Q Verified Limit
BFD sessions	128 (IPv4) + 128 (IPv6)
Generic routing encapsulation (GRE) tunnels	8
Port channel links	32
SVIs	250
vPCs	60

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

Feature	3164Q Verified Limit
MST instances	64

Feature	3164Q Verified Limit
MST virtual ports	48,000
RPVST virtual ports	12,000
VLANs	3900
VLANs in RPVST mode	500

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 3: Multicast Routing Verified Scalability Limits (Unidimensional)

Feature	3164Q Verified Limit
IPv4 multicast routes	32,000
Outgoing interfaces (OIFs)	40 (see CSCum58876)
IGMP snooping groups	32,000
PIM neighbors	250

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 4: Security Verified Scalability Limits (Unidimensional)

Feature	3164Q Verified Limit
DHCP snooping bindings	2000
IPv4 ingress ACLs	3070 (per network forwarding engine)
IPv4 egress ACLs	765 (per network forwarding engine)

Feature	3164Q Verified Limit
IPv6 ingress ACLs	1530 (per network forwarding engine)
IPv6 egress ACLs	250 (per network forwarding engine)

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

Table 5: System Management Verified Scalability Limits (Unidimensional)

Feature	3164Q Verified Limit
PTP	
10G physical ports enabled for PTP	44
SPAN and ERSPAN	
Configurable SPAN or ERSPAN sessions	4
Active SPAN or ERSPAN sessions[1]	4
Active localized SPAN or ERSPAN sessions per line card[2]	4
Source interfaces per SPAN or ERSPAN session (Rx and Tx, Rx, or Tx)	48
Destination interfaces per SPAN session	1 (physical interface)
Source VLANs per SPAN or ERSPAN session	32
TAP Aggregation	
Redirect interfaces in the redirect port list	12
Redirect port lists (or fan outs) per system	100

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 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.

Table 6: Unicast Routing Verified Scalability Limits (Unidimensional)

Feature	3164Q Verified Limit
Unicast Routing	
BGP neighbors	512 (IPv4 only) 512 (IPv6 only) 256 (IPv4) + 256 (IPv6)
eBGP	1000
EIGRP neighbors	128 (IPv4), 128 (IPv6)
EIGRP routes	20,000
HSRP groups	250
IPv4 ARP	48,000
IPv4 host routes	208,000
IPv6 host routes	40,000
IPv6 ND	40,000
IPv4 unicast routes (LPM)	128,000 (default system routing mode) 128,000 with no IPv6 routes (64-bit ALPM routing mode)
IPv6 unicast routes (LPM)	20,000 (default system routing mode) 80,000 with no IPv4 routes (64-bit ALPM routing mode)
IPv4 and IPv6 unicast routes (LPM) in 64-bit ALPM routing mode	*x* IPv6 routes and *y* IPv4 routes, where $2x + y <= 128,000$
MAC addresses	80,000
OSPFv2 neighbors	256
OSPFv3 neighbors	256
VRFs	1000
VRRP groups per interface or I/O module	250
Policy-Based Routing (PBR)	

Feature	3164Q Verified Limit
Configured sequences per policy	256
Next-hop addresses per policy	32
IPv4 ACEs (unidimensional)	3072 (per network forwarding engine)
IPv6 ACEs (unidimensional)	1536 (per network forwarding engine)
IPv4 and IPv6s ACEs	2048 IPv4 + 256 IPv6
Interfaces with PBR policy	512
VRRPv3	
VRRPv3 groups per interface	255
VRRPv3 groups with default timers (1 s)	490
VRRPv3 groups with aggressive timers (100 ms)	200
VRRPv3 groups with relaxed timers (3 s)	490
Pathways with one VRRPv3 group with default timer (1 s)	489
VRRPv3 groups and pathways combined	490

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 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 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 multidimensional scenarios.

Table 7: VXLAN Verified Scalability Limits (Unidimensional)

Feature	3164Q Verified Limit
VXLAN Flood and Learn	
Virtual network identifiers (VNIs) or VXLAN-mapped VLANs	1000
Underlay multicast groups	128
Overlay MAC addresses	64,000
Remote VXLAN tunnel endpoints (VTEPs)	256
Ingress replication peers	256
Ingress replication Layer 2 VNIs	1000
MAC addresses for ingress replication	64,000
VXLAN VLAN logical port VP count	6000
VXLAN BGP eVPN	
Layer 2 VNI	1000
Underlay multicast groups	128
VTEPs	256
BGP VTEP peers	256
MAC addresses	64,000
VXLAN VLAN logical port VP count	6000