Verified Scalability for Cisco Nexus 6000 Series NX-OS Release 7.1(1)N1(1)

First Published: 2014-01-15

Last Modified: 2017-12-05

Americas Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
http://www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 527-0883

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 NONINFRINGEMENT 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.

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)

© 2014-2017 Cisco Systems, Inc. All rights reserved.

CONTENTS

Preface

Preface v

Preface v

Audience v

Document Conventions v

Related Documentation vii

Documentation Feedback vii

Obtaining Documentation and Submitting a Service Request vii

CHAPTER 1

Verified Scalability for Cisco Nexus 6000 Series NX-OS Release 7.1(1)N1(1) 1

Overview of Verified Scalability 1

Verified Scalability for a Layer 2 Switching Deployment 2

Verified Scalability for a Layer 2 Switching and Layer 3 Routing Deployment 4

Verified Scalability for a Layer 3 Aggregation Routing Deployment 6

Verified Scalability for Fiber Channel (FC) and Fiber Channel over Ethernet (FCoE) 8

Verified Scalability for Multicast Routing 8

Verified Scalability for Unicast Routing 9

Preface

The preface contains the following sections:

- Preface, page v

Preface

This preface describes the audience, organization, and conventions of the Book Title. It also provides information on how to obtain related documentation.

This chapter includes the following topics:

Audience

This publication is for experienced network administrators who configure and maintain Cisco NX-OS on Cisco Nexus 5000 Series Platform switches.

Document Conventions

Note

- As part of our constant endeavor to remodel our documents to meet our customers' requirements, we have modified the manner in which we document configuration tasks. As a result of this, you may find a deviation in the style used to describe these tasks, with the newly included sections of the document following the new format.
- The Guidelines and Limitations section contains general guidelines and limitations that are applicable to all the features, and the feature-specific guidelines and limitations that are applicable only to the corresponding feature.

Command descriptions use the following conventions:

Convention	Description
bold	Bold text indicates the commands and keywords that you enter literally as shown.
Italic	Italic text indicates arguments for which the user supplies the values.
[x]	Square brackets enclose an optional element (keyword or argument).
[x | y]	Square brackets enclosing keywords or arguments separated by a vertical bar indicate an optional choice.
{x | y}	Braces enclosing keywords or arguments separated by a vertical bar indicate a required choice.
[x {y | z}]	Nested set of square brackets or braces indicate optional or required choices within optional or required elements. Braces and a vertical bar within square brackets indicate a required choice within an optional element.
`variable`	Indicates a variable for which you supply values, in context where italics cannot be used.
string	A nonquoted set of characters. Do not use quotation marks around the string or the string will include the quotation marks.

Examples use the following conventions:

Convention	Description
`screen font`	Terminal sessions and information the switch displays are in screen font.
`boldface screen font`	Information you must enter is in boldface screen font.
italic screen font	Arguments for which you supply values are in italic screen font.
< >	Nonprinting characters, such as passwords, are in angle brackets.
[ ]	Default responses to system prompts are in square brackets.
!, #	An exclamation point (!) or a pound sign (#) at the beginning of a line of code indicates a comment line.

This document uses the following conventions:

Note Means *reader take note*. Notes contain helpful suggestions or references to material not covered in the manual.

Caution Means *reader be careful*. In this situation, you might do something that could result in equipment damage or loss of data.

Related Documentation

Documentation for Cisco Nexus 6000 Series Switches is available at:

- Configuration Guides

  http://www.cisco.com/c/en/us/support/switches/nexus-6000-series-switches/products-installation-and-configuration-guides-list.html

- Command Reference Guides

  http://www.cisco.com/c/en/us/support/switches/nexus-6000-series-switches/products-command-reference-list.html

- Release Notes

  http://www.cisco.com/c/en/us/support/switches/nexus-6000-series-switches/products-release-notes-list.html

- Install and Upgrade Guides

  http://www.cisco.com/c/en/us/support/switches/nexus-6000-series-switches/products-installation-guides-list.html

- Licensing Guide

  http://www.cisco.com/c/en/us/support/switches/nexus-6000-series-switches/products-licensing-information-listing.html

Documentation for Cisco Nexus 6000 Series Switches and Cisco Nexus 2000 Series Fabric Extenders is available at:

http://www.cisco.com/c/en/us/support/switches/nexus-2000-series-fabric-extenders/products-installation-and-configuration-guides-list.html

Documentation Feedback

To provide technical feedback on this document, or to report an error or omission, please send your comments to nexus6k-docfeedback@cisco.com. We appreciate your feedback.

Obtaining Documentation and Submitting a Service Request

For information on obtaining documentation, using the Cisco Bug Search Tool (BST), submitting a service request, and gathering additional information, see What's New in Cisco Product Documentation.

To receive new and revised Cisco technical content directly to your desktop, you can subscribe to the What's New in Cisco Product Documentation RSS feed. RSS feeds are a free service.

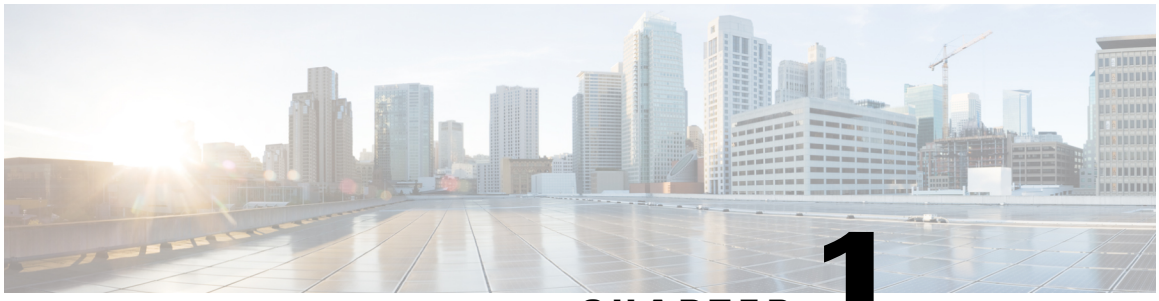

CHAPTER 1

Verified Scalability for Cisco Nexus 6000 Series NX-OS Release 7.1(1)N1(1)

This chapter contains the following sections:

- Overview of Verified Scalability, page 1
- Verified Scalability for a Layer 2 Switching Deployment, page 2
- Verified Scalability for a Layer 2 Switching and Layer 3 Routing Deployment, page 4
- Verified Scalability for a Layer 3 Aggregation Routing Deployment, page 6
- Verified Scalability for Fiber Channel (FC) and Fiber Channel over Ethernet (FCoE), page 8
- Verified Scalability for Multicast Routing, page 8
- Verified Scalability for Unicast Routing, page 9

Overview of Verified Scalability

This document lists the Cisco verified scalability limits.

Note The following scaling capabilities apply to Cisco Nexus 6000 Series switches. For example, 6001 and 6004.

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

The Verified Maximum column lists the maximum scale capability tested for the corresponding feature individually. This number is the absolute maximum currently supported by the Cisco NX-OS Release software for the corresponding feature. If the hardware is capable of a higher scale, future software releases may increase this verified maximum limit.

Verified Scalability for a Layer 2 Switching Deployment

This table lists the verified scalability for a Layer 2 switching deployment.

Table 1: Verified Scalability for a Layer 2 Switching Deployment

Feature	Verified Topology	Verified Maximum
Active VLANs/VSANs per switch	4000	4013 (31 are reserved for VSANs, and the remaining are for VLANs.)
VLAN/VSAN ID Space	4013	4013 Unreserved Space
Logical Interfaces[1]	96,000 (MST) [2]	64,000 (96,000 MST) [3]
PVLAN's (Primay+Secondary)	16	16
Number of PVLAN ports	960	960
Port Security enabled interfaces	960	960
QoS enabled interfaces	960	960
VLAN ACLs (VACLs)	128 (10 Unique VACLs)	1024 (512 unique VACLs with up to 1024 ACE entries across all VACLs)
Maximum interfaces per EtherChannel	16	16
IGMP Snooping Groups	4000 (in FEX deployments) 8000 (in non-FEX deployments)	4000 (in FEX deployments) 16,000 (in non-FEX deployments)
Maximum FEXs per Switch	48 [4]	48 [5]
Maximum FEXs dual-homed to a vPC Switch Pair	48	48
MAC Table Size (Entries)	64,000	115,000

Feature	Verified Topology	Verified Maximum
Number of Switchport Etherchannels	48	96 (Single member port-channel for 40G ports) 384 (Single member port-channel for 10G ports) 64 (Multi member port-channel)
Number of HIF FEX port channels/vPCs (across the maximum number of FEXs)	576	576
SPAN Sessions	4 active sessions 32 source VLANs as a RX source	16 active sessions 32 source VLANs as a RX source
SVIs	2	256
FabricPath VLANs	4000[6]	4000
FabricPath Switch IDs	500	500
FabricPath Multicast Trees	1	1
Number of FabricPath Topologies	2	2
Number of FabricPath Core Port-Channels	4 core links with 4 ports each	16
FEX Host Interface Storm Control	1936 [7]	1936
Segmentation ID	3000 (1000 global segments, 2000 local segments)	3000 (1000 global segments, 2000 local segments)

1 Logical interfaces are a product of the number of VLANs times the number of ports. This parameter reflects the load of handling port programming, and is not dependent on the spanning-tree mode or configuration.

2 32,000 STP logical interfaces are verified in the unified fabric topology.

3 16,000 Port-VLAN scaling number applies to Rapid PVST+ and non-STP modes. For MST mode only, the maximum verified limit for Port-VLAN scaling is 96,000. 64000 limit is verified when a switch is running in MST mode and performing Rapid PVST+ stimulation. 48,000 is for MST and 16,000 is for Rapid PVST+.

4 48 maximum FEXs per switch applies to the Cisco Nexus 6004. The limit is 24 maximum FEXs per switch for the Cisco Nexus 6001.

5 48 maximum FEXs per switch applies to the Cisco Nexus 6004. The limit is 24 maximum FEXs per switch for the Cisco Nexus 6001.

6 FabricPath VLANs are verified in the unified fabric topology

7 This is the target maximum number that HIF-SC can support. Beyond this number, NIF-SC is recommended for deployment.

Verified Scalability for a Layer 2 Switching and Layer 3 Routing Deployment

This table contains the verified scalability for a Layer 2 switching and Layer 3 routing deployment.

Note The currently tested values do not provide an indication for the maximum scalability of the control plane. These numbers vary based on the load of the system in terms of routing protocols, timers settings, and other values. Proof of concept testing should be used to determine the scalability of a given feature for your environment.

Table 2: Verified Scalability for a Layer 2 Switching and Layer 3 Routing Deployment

Feature	Verified Topology	Verified Maximum
Active VLANs/VSANs per Switch	1000 8	4013 (31 are set reserved for VSANs and the remaing are for VLANs)
VLAN/VSAN ID Space	4013 Unreserved space	4013 Unreserved space
STP Instances	16,000	16,000
Maximum Interfaces per EtherChannel	16	16
IGMP Snooping Groups	4000 (in FEX deployments) 8000 (in non-FEX deployments)	4000 (in FEX deployments) 16,000 (in non-FEX deployments)
Maximum FEXs per Switch	24 32 (for N6004)	24
Maximum FEXs Dual-homed to a vPC Switch Pair 9	24 32 (for N6004)	24
MAC Table Size (Entries)	32,000 10	64,000 11
Number of FEX Port Channels/vPCs (across the maximum number of FEXs)	691	768

Feature	Verified Topology	Verified Maximum
SPAN Sessions	4 active sessions 32 source VLANs as an RX source	16 active sessions 32 source VLANs as an RX source
Number of SVIs	564	256
Dynamic IPv4 Routes [12]	8000	24,000
Dynamic IPv6 Routes [13]	4000 [14]	6000
Multicast IPv4 Routes [15]	4000	8000
ARPs (IPv4 Hosts) [16]	32,000	64,000
IPv6 Hosts	16,000	32,000
VRFs	25	1000
RACLs	64 Ingress RACLs with up to 1152 ACE entries across all the RACLs	896 Ingress RACLs with up to 1152 ACE entries across all the RACLs
HSRP Groups [17]	256	500
VRRP Groups [18]	256	500
BFD sessions over L3-intf for CE mode **Note** BFD scaling tested with 5 physical links.	8 sessions (250ms intvl, 750ms dead-intvl)	30 (250ms intvl, 750ms dead-intvl)
BFD sessions over SVI for FabricPath mode **Note** BFD scaling tested with 5 physical links.	64 (250ms intvl, 750ms dead-intvl)	64 (250ms intvl, 750ms dead-intvl)
PBR IPv4	95	95
PBR IPv6	95	95

8 4,013 VLANs are verified in Layer 2 switching, Fibre Channel;, and FCoE topologies.
9 FEXs are verified in the Layer 2 topology
10 128,000 entries are reserved for unicast MAC entries and 128,000 entries are reserved for IP host routes.

11 128,000 entries are reserved for unicast MAC entries and 128,000 entries are reserved for IP host routes.

12 The maximum number of entries that can be supported is 24,000. This table is shared between IPv4 and IPv6. An IPv4 route takes up one entry in the table and an IPv6 route takes up four entries.

13 The maximum number of entries that can be supported is 24,000. This table is shared between IPv4 and IPv6. An IPv4 route takes up one entry in the table and an IPv6 route takes up four entries.

14 With no IPv4, the number increases to 6000.

15 This includes (*,G) entries, (S,G) entries, and the entries required for vPC with bind-vrf configured. When bind-vrf is configured, each (*,G) and (S,G) entry is replicated.

16 The maximum number of hosts supported is listed under ARPs. This includes IPv4 and IPv6 hosts. IPv4 hosts take up one entry and IPv6 hosts take up 2 entries in hardware. So, for the verified limit, the switch supports one of the following: 32,000 IPv4 hosts and 0 IPv6 hosts, 16,000 IPv6 hosts and 0 IPv4 hosts, or a combination of IPv4 and IPv6 hosts.

17 The limit of the table that holds the Router MAC and Virtual MAC entries that determines whether the packet needs to be bridged or routed is 500 entries. The Virtual MAC entries can be shared across Layer 3 interfaces. So, we recommend that you configure the same group ID across all or multiple Layer 3 interfaces/SVIs. If multiple group IDs are configured on an Layer 3 interface, then we recommend that you configure the same set of group IDs across all or multiple Layer 3 interfaces. This way, HSRP/VRRP can be supported on more interfaces. Please refer to the unicast configuration guide for more information.

18 The limit of the table that holds the Router MAC and Virtual MAC entries for determining packet routing or switching is 500 entries. The Virtual MAC entries can be shared across Layer 3 interfaces. So, we recommend that you configure the same group ID across all or multiple Layer 3 interfaces/SVIs. If multiple group IDs are configured on an Layer 3 interface, then we recommend that you configure the same set of group IDs across all or multiple Layer 3 interfaces. This way, HSRP/VRRP can be supported on more interfaces. Please refer to the unicast configuration guide for more information.

Verified Scalability for a Layer 3 Aggregation Routing Deployment

This table lists the verified scalability for a Layer 3 aggregation routing deployment.

Table 3: Verified Scalability for a Layer 3 Aggregation Routing Deployment

Feature	Verified Topology	Verified Maximum
Active VLANs/VSANs per Switch	4000 19	4000
VLAN/VSAN ID Space	4013 unreserved space	4013 unreserved space
STP Instances	64,000	64,000
Maximum Interfaces per EtherChannel	16	16
IGMP Snooping Groups	8000 (in non-FEX deployments)	16,000 (in non-FEX deployments)
MAC Table Size	64,000 20	115,000 21
SPAN Sessions	4 active sessions 32 source VLANs as an RX source	16 active sessions 32 source VLANs as an RX source
SVIs	1000	4000
Dynamic IPv4 Routes 22	8000	24,000

Feature	Verified Topology	Verified Maximum
Dynamic IPv6 Routes [23]	4000 [24]	6000
Multicast IPv4 Routes [25]	8000	16,000
RACLs	64 ingress RACLs with up to 1,152 ACE entries across all of the RACLs.	512 ingress RACLs with up to 1,152 ACE entries across all of the RACLs.
VRFs	25	1,000
ARPs (IPv4 Hosts) [26]	64,000	64,000
IPv6 Hosts [27]	10,000 [28]	32,000
IGP Peers	64	100
HSRP Groups [29]	1721 (911 IPv4, 810 IPv6)	1721 (911 IPv4, 810 IPv6)
VRRP Groups [30]	500	500
FabricPath Switch IDs	500	500
FabricPath Multicast Trees	2	2
Number of FabricPath Topologies	2	2
Number of FabricPath Core Links	32	32
PBR IPv4	110	110
PBR IPv6	110	110
BFD sessions over L3-intf for CE Mode **Note** BFD scaling tested under a single physical link with 100 sub-interfaces.	100 sessions (150ms intvl, 750ms dead-intvl)	100 sessions (250ms intvl, 750ms dead-intvl)

Feature	Verified Topology	Verified Maximum
BFD sessions over SVI for FabricPath mode **Note** BFD scaling tested under a single physical link with 100 sub-interfaces.	64 sessions (150ms intvl, 750ms dead-intvl)	64 sessions (250ms intvl, 750ms dead-intvl)

19 4,013 VLANs are verified in Layer 2 switching, Fibre Channel, and FCoE topologies.

20 128,000 entries are reserved for Unicast MAC entries, and 128,000 entries are reserved for IP host routes.

21 128,000 entries are reserved for Unicast MAC entries, and 128,000 entries are reserved for IP host routes.

22 The maximum number of entries that can be supported is 24,000. This table is shared between IPv4 and IPv6. An IPv4 route takes up one entry in the table and an IPv6 route takes up four entries.

23 The maximum number of entries that can be supported is 24,000. This table is shared between IPv4 and IPv6. An IPv4 route takes up one entry in the table and an IPv6 route takes up four entries.

24 Entries shared between IPv4, IPv6 network routes.

25 This includes (*,G) entries, (S,G) entries, and the entries required for vPC with bind-vrf configured. When bind-vrf is configured, each (*,G) and (S,G) entry is replicated. This includes (*,G) entries, (S,G) entries, and IGMP-snooping entries combined.

26 The maximum number of hosts supported is listed under ARPs. This includes IPv4 and IPv6 hosts. IPv4 hosts take up one entry and IPv6 hosts take up 2 entries in hardware. So, for the verified limit, the switch supports one of the following: 64,000 IPv4 hosts and 0 IPv6 hosts, 16,000 IPv6 hosts and 0 IPv4 hosts, or a combination of IPv4 and IPv6 hosts.

27 The maximum number of hosts supported is listed under ARPs. This includes IPv4 and IPv6 hosts. IPv4 hosts take up one entry and IPv6 hosts take up 2 entries in hardware. So, for the verified limit, the switch supports one of the following: 64,000 IPv4 hosts and 0 IPv6 hosts, 16,000 IPv6 hosts and 0 IPv4 hosts, or a combination of IPv4 and IPv6 hosts.

28 Entries shared between IPv4 multicast, IPv4, IPv6 host routes .

29 The limit of the table that holds the Router MAC and Virtual MAC entries that determine whether the packet needs to be bridged or routed is 500 entries. The Virtual MAC entries can be shared across Layer 3 interfaces. So we recommend that you configure the same group ID across all or multiple Layer 3 interfaces/SVIs. If multiple group IDs are configured on a Layer 3 interface, then we recommend that you configure the same set of group IDs across all or multiple Layer 3 interfaces. This way, HSRP/VRRP can be supported on more interfaces. Please refer to the Unicast Routing Configuration Guide for more information.

30 The limit of the table that holds the Router MAC and Virtual MAC entries for determining packet routing or switching is 500 entries. The Virtual MAC entries can be shared across Layer 3 interfaces. So we recommend that you configure the same group ID across all or multiple Layer 3 interfaces/SVIs. If multiple group IDs are configured on a Layer 3 interface, then we recommend that you configure the same set of group IDs across all or multiple Layer 3 interfaces. This way, HSRP/VRRP can be supported on more interfaces. Please refer to the Unicast Routing Configuration Guide for more information.

Verified Scalability for Fiber Channel (FC) and Fiber Channel over Ethernet (FCoE)

Refer to Cisco Nexus 5000 and 6000 Series FC and FCoE Configuration Limits for NX-OS Release 7.3(0)N1(1)

Verified Scalability for Multicast Routing

This table lists the verified scalability for multicast routing.

Table 4: Verified Scalability for Multicast Routing

Feature	Parameter	Verified Maximum
Protocol Independent Multicast (PIM)	Number of neighbors	500
	Number of neighbors/total routes per system with aggressive hello timers (5 seconds)	16/4,000
Multicast Source Discovery Protocol (MSDP)	Number of MSDP Source-Active (SA) cache entries	6,000

Note

- In vPC setup, TCAM exhaustion failure will lead to some routes not getting programmed in the hardware. Hence, there might exist a condition where mrib will show the route exists but mfib may not have it programmed.
- In bind-vrf configuration, for every mroute, additional mroute is programmed in the hardware and this could lead to TCAM exhaustion. Hence, ensure that the mroute count does not exceed ((max-limit/2) - 4) default routes.

  For example: If the hardware profile multicast max-limit is 8000, then mroute count should not exceed ((8000/2)-4) default routes.

Verified Scalability for Unicast Routing

Guidelines and Limitations for Unicast Routing

- You can have up to four instances of OSPFv2.
- You can have up to four instances of OSPFv3.

This table lists the verified scalability for unicast routing.

Table 5: Verified Scalability for Unicast Routing

Feature	Parameter	Verified Maximum
OSPFv2	Number of active interfaces	256
	Number of passive interfaces	256
	Number of neighbors/total routes with aggressive timers (1 sec/ 3 sec)	16/6,000

Feature	Parameter	Verified Maximum
OSPFv3	Number of active interfaces	256
	Number of passive interfaces	256
EIGRP	Number of active interfaces	50
BGP	Number of peers (iBGP and eBGP, active)	256
	Number of AS path entries	512
	Number of prefix-list entries in a single prefix-list	10,000
HSRP	Number of groups with aggressive timers (1 sec/3 sec)	500
L3 ISIS	Number of adjacencies	100
Unicast Adjacencies	Number of regular adjacencies	16374
	Number of ECMP adjacencies. The regular adjacencies are also part of the maximum ECMP adjacencies	28470