



Cisco Nexus 7000 Series NX-OS Unicast Routing Configuration Guide

First Published: 2009-10-21

Last Modified: 2020-03-02

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 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.

All printed copies and duplicate soft copies of this document are considered uncontrolled. See the current online version for the latest version.

Cisco has more than 200 offices worldwide. Addresses and phone numbers are listed on the Cisco website at www.cisco.com/go/offices.

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/c/en/us/about/legal/trademarks.html>. 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. (1721R)

© 2009–2020 Cisco Systems, Inc. All rights reserved.



CONTENTS

PREFACE

Preface	xxix
Audience	xxix
Document Conventions	xxix
Related Documentation for Cisco Nexus 7000 Series NX-OS Software	xxx
Documentation Feedback	xxxii
Communications, Services, and Additional Information	xxxiii

CHAPTER 1

New and Changed Information	1
New and Changed Information	1

CHAPTER 2

Overview	3
Overview	3
Licensing Requirements	3
Information About Layer 3 Unicast Routing	3
Routing Fundamentals	3
Packet Switching	4
Routing Metrics	5
Router IDs	6
Autonomous Systems	7
Convergence	8
Load Balancing and Equal Cost Multipath	8
Route Redistribution	8
Administrative Distance	9
Stub Routing	9
Routing Algorithms	10
Static Routes and Dynamic Routing Protocols	10

- Interior and Exterior Gateway Protocols 10
 - Distance Vector Protocols 10
 - Link-State Protocols 11
- Layer 3 Virtualization 11
- Cisco NX-OS Forwarding Architecture 12
 - Unicast RIB 12
 - Adjacency Manager 13
 - Unicast Forwarding Distribution Module 13
 - FIB 13
 - Hardware Forwarding 13
 - Software Forwarding 14
- Layer 3 Interoperation with the N7K-F132-15 Module 14
- Summary of Layer 3 Routing Features 15
 - IPv4 and IPv6 15
 - IP Services 15
 - OSPF 15
 - EIGRP 15
 - IS-IS 15
 - BGP 15
 - RIP 16
 - Static Routing 16
 - Layer 3 Virtualization 16
 - Route Policy Manager 16
 - Policy-Based Routing 16
 - First Hop Redundancy Protocols 16
 - Object Tracking 17
- Related Documents for Layer 3 Unicast Routing 17

PART I IP 19

CHAPTER 3 Configuring IPv4 21

- Finding Feature Information 21
- Information About IPv4 21
- Multiple IPv4 Addresses 22

Address Resolution Protocol	22
ARP Caching	23
Static and Dynamic Entries in the ARP Cache	23
Devices That Do Not Use ARP	23
Reverse ARP	24
Proxy ARP	24
Local Proxy ARP	25
Gratuitous ARP	25
ARP Refresh on MAC Delete	25
Glean Throttling	25
Path MTU Discovery	26
ICMP	26
Virtualization Support for IPv4	26
Prerequisites for IPv4	26
Guidelines and Limitations for IPv4	27
Default Settings for IPv4 Parameters	27
Configuring IPv4	27
Configuring IPv4 Addressing	27
Configuring Multiple IPv4 Addresses	28
Configuring a Static ARP Entry	29
Configuring Proxy ARP	30
Configuring Local Proxy ARP	30
Configuring Gratuitous ARP	31
Configuring the IP ARP Cache Limit	32
Configuring Glean Optimization	32
Configuring Bloom Filter Support for Glean Adjacencies	33
Configuring Path MTU Discovery	34
Configuring IP Packet Verification	34
Configuring IP Glean Throttling	35
Configuring the Hardware IP Glean Throttle Maximum	36
Configuring the Hardware IP Glean Throttle Timeout	37
Configuring the Hardware IP Glean Throttle Syslog	38
Verifying the IPv4 Configuration	38
Configuration Examples for IPv4	39

Example: Reserving All Ports on a Module for Proxy Routing	39
Example: Reserving Ports for Proxy Routing	41
Example: Excluding Ports From Proxy Routing	41
Related Documents for IPv4	42
Standards for IPv4	42
Feature History for IPv4	43

CHAPTER 4
Configuring IPv6 45

Finding Feature Information	45
Information About IPv6	45
IPv6 Address Formats	46
IPv6 Unicast Addresses	47
Aggregatable Global Addresses	47
Link-Local Addresses	48
IPv4-Compatible IPv6 Addresses	49
Unique Local Addresses	49
Site Local Addresses	50
IPv6 Anycast Addresses	50
IPv6 Multicast Addresses	50
IPv4 Packet Header	52
Simplified IPv6 Packet Header	52
DNS for IPv6	55
Path MTU Discovery for IPv6	55
CDP IPv6 Address Support	56
ICMP for IPv6	56
IPv6 Neighbor Discovery	57
IPv6 Neighbor Solicitation Message	57
IPv6 Router Advertisement Message	59
IPv6 Router Advertisement Options for DNS Configuration	60
IPv6 Neighbor Redirect Message	60
Virtualization Support for IPv6	62
Prerequisites for IPv6	62
Guidelines and Limitations for Configuring IPv6	62
Default Settings for IPv6	63

Configuring IPv6	63
Configuring IPv6 Addressing	63
Configuring IPv6 Neighbor Discovery	64
Configuring Optional IPv6 Neighbor Discovery	67
Configuring Recursive DNS Server (RDNSS)	69
Configuring DNS Search List (DNSSL)	70
Configuring IPv6 Packet Verification	71
Verifying the IPv6 Configuration	72
Configuration Example for IPv6	72
Related Documents for IPv6	72
Standards for IPv6	72
Feature History for IPv6	73

CHAPTER 5
Configuring DNS 75

Finding Feature Information	75
Information About DNS Clients	75
DNS Client Overview	75
DNS Name Servers	76
DNS Operation	76
High Availability for DNS Clients	76
Virtualization Support for DNS Clients	76
Prerequisites for DNS Clients	76
Guidelines and Limitations for DNS Clients	77
Default Settings for DNS Client Parameters	77
Configuring DNS Clients	77
Configuring the DNS Client	77
Verifying the DNS Client Configuration	79
Configuration Examples for DNS Clients	79
Related Documents for DNS Clients	79
Standards for DNS Clients	79
Feature History for DNS	79

CHAPTER 6
Configuring WCCPv2 81

Finding Feature Information	81
-----------------------------	----

- Information About WCCPv2 81
 - WCCPv2 Overview 81
 - WCCPv2 Service Types 82
 - WCCPv2 Service Groups 82
 - WCCPv2 Service Group Lists 83
 - WCCPv2 Designated Cache Engine 84
 - WCCPv2 Redirection 84
 - WCCPv2 Authentication 86
 - WCCPv2 Redirection Method 87
 - WCCPv2 Packet Return Method 87
 - High Availability for WCCPv2 87
 - Virtualization Support for WCCPv2 87
 - WCCPv2 Error Handling for SPM Operations 88
- Prerequisites for WCCPv2 88
- Guidelines and Limitations for WCCPv2 88
- WCCPv2 Default Settings 90
- Configuring WCCPv2 90
 - Enabling and Disabling WCCPv2 90
 - Configuring a WCCPv2 Service Group 91
 - Applying WCCPv2 Redirection to an Interface 92
 - Configuring WCCPv2 in a VRF 94
- Verifying the WCCPv2 Configuration 96
- Configuration Examples for WCCPv2 96
- Related Documents for WCCPv2 97
- Standards for the WCCPv2 97
- Feature History for WCCPv2 97

PART II

Routing 99

CHAPTER 7

Configuring OSPFv2 101

- Finding Feature Information 101
- Information About OSPFv2 101
 - Hello Packet 102
 - Neighbors 102

Adjacency	103
Designated Routers	103
Areas	104
Link-State Advertisements	105
Link-State Advertisements Types	105
OSPFv2 and the Unicast RIB	107
Authentication	107
Simple Password Authentication	107
MD5 Authentication	107
Advanced Features for OSPFv2	108
Stub Area	108
Not-So-Stubby Area	108
Virtual Links	109
Route Redistribution	109
Route Summarization	110
High Availability and Graceful Restart	110
OSPFv2 Stub Router Advertisements	111
Multiple OSPFv2 Instances	111
SPF Optimization	111
BFD	112
Virtualization Support for OSPFv2	112
Prerequisites for OSPFv2	112
Guidelines and Limitations for OSPFv2	112
Default Settings for OSPFv2	114
Configuring Basic OSPFv2	115
Enabling OSPFv2	115
Creating an OSPFv2 Instance	115
Configuring OSPF Packet Size	116
Configuring Optional Parameters on an OSPFv2 Instance	118
Configuring Networks in OSPFv2	119
Configuring Authentication for an Area	121
Configuring Authentication for an Interface	122
Configuring Advanced OSPFv2	124
Configuring Filter Lists for Border Routers	124

Configuring Stub Areas	125
Configuring a Totally Stubby Area	125
Configuring NSSA	126
Configuring Virtual Links	127
Configuring Redistribution	129
Limiting the Number of Redistributed Routes	131
Configuring Route Summarization	132
Configuring Stub Route Advertisements	133
Configuring the Administrative Distance of Routes	134
Modifying the Default Timers	136
Configuring Graceful Restart	138
Restarting an OSPFv2 Instance	139
Configuring OSPFv2 with Virtualization	140
Verifying the OSPFv2 Configuration	141
Monitoring OSPFv2	142
Configuration Examples for OSPFv2	142
Related Documents for OSPFv2	143
Feature History for OSPFv2	143

CHAPTER 8

Configuring OSPFv3	145
Finding Feature Information	145
Information About OSPFv3	145
Comparison of OSPFv3 and OSPFv2	146
Hello Packet	146
Neighbors	147
Adjacency	147
Designated Routers	147
Areas	148
Link-State Advertisement Types	149
Link Cost	151
Flooding and LSA Group Pacing	151
Link-State Database	151
Multi-Area Adjacency	151
OSPFv3 and the IPv6 Unicast RIB	152

Address Family Support	152
Authentication	152
Encryption	153
Guidelines and Limitations for configuring ESP on OSPFv3	153
Advanced Features	154
Stub Area	154
Not-So-Stubby Area	154
Virtual Links	155
Route Redistribution	155
Route Summarization	156
High Availability and Graceful Restart	156
Multiple OSPFv3 Instances	157
SPF Optimization	157
Virtualization Support	157
Prerequisites for OSPFv3	158
Guidelines and Limitations for OSPFv3	158
Default Settings for OSPFv3	160
Configuring Basic OSPFv3	160
Enabling OSPFv3	160
Creating an OSPFv3 Instance	161
Configuring OSPFv3 Packet Size	163
Configuring Networks in OSPFv3	164
Configuring Advanced OSPFv3	166
Configuring Filter Lists for Border Routers	166
Configuring Stub Areas	167
Configuring a Totally Stubby Area	168
Configuring NSSA	168
Configuring Multi-Area Adjacency	170
Configuring Virtual Links	171
Configuring Redistribution	172
Limiting the Number of Redistributed Routes	174
Configuring Route Summarization	175
Configuring the Administrative Distance of Routes	176
Modifying the Default Timers	179

Configuring the OSPFv3 Max-Metric Router LSA	181
Configuring Graceful Restart	181
Restarting an OSPFv3 Instance	182
Configuring OSPFv3 with Virtualization	183
Configuring OSPFv3 Authentication at Router Level	184
Configuring OSPFv3 Authentication at Area Level	185
Configuring OSPFv3 Authentication at Interface Level	185
Configuring OSPFv3 Encryption at Router Level	187
Configuring OSPFv3 Encryption at Area Level	188
Configuring OSPFv3 Encryption at Interface Level	188
Configuring OSPFv3 Encryption for Virtual Links	190
Configuration Examples for OSPFv3	191
Related Documents for OSPFv3	191
Feature History for OSPFv3	192

CHAPTER 9
Configuring EIGRP 193

Finding Feature Information	193
Information About EIGRP	193
EIGRP Components	194
Reliable Transport Protocol	194
Neighbor Discovery and Recovery	194
Diffusing Update Algorithm	194
EIGRP Route Updates	195
Internal Route Metrics	195
Wide Metrics	196
External Route Metrics	196
EIGRP and the Unicast RIB	197
Advanced EIGRP	197
Address Families	197
Authentication	197
Stub Routers	198
Route Summarization	198
Route Redistribution	198
Load Balancing	199

Split Horizon	199
BFD	199
Virtualization Support for EIGRP	199
Graceful Restart and High Availability	200
Multiple EIGRP Instances	200
Prerequisites for EIGRP	201
Guidelines and Limitations for EIGRP	201
Default Settings for EIGRP Parameters	202
Configuring Basic EIGRP	203
Enabling or Disabling the EIGRP Feature	203
Creating an EIGRP Instance	203
Restarting an EIGRP Instance	205
Shutting Down an EIGRP Instance	205
Configuring a Passive Interface for EIGRP	206
Shutting Down EIGRP on an Interface	207
Configuring Advanced EIGRP	207
Configuring Authentication in EIGRP	207
Configuring EIGRP Stub Routing	209
Configuring a Summary Address for EIGRP	210
Redistributing Routes into EIGRP	211
Limiting the Number of Redistributed Routes	212
Configuring the Administrative Distance of Routes	213
Configuring Route-Map Filtering	214
Configuring Load Balancing in EIGRP	215
Configuring Graceful Restart for EIGRP	216
Adjusting the Interval Between Hello Packets and the Hold Time	218
Disabling Split Horizon	219
Enabling Wide Metrics	219
Tuning EIGRP	220
Configuring Virtualization for EIGRP	222
Verifying the EIGRP Configuration	224
Displaying EIGRP Statistics	224
Configuration Example for EIGRP	224
Related Documents for EIGRP	225

MIBs	226
Feature History for EIGRP	226

CHAPTER 10
Configuring IS-IS 227

Finding Feature Information	227
Information About IS-IS	227
IS-IS Overview	228
IS-IS Areas	228
NET and System ID	229
Designated Intermediate System	229
IS-IS Authentication	229
Mesh Groups	229
Overload Bit	230
Route Summarization	230
Route Redistribution	230
Administrative Distance	230
Load Balancing	231
BFD	231
Virtualization Support	231
High Availability and Graceful Restart	231
Multiple IS-IS Instances	232
Prerequisites for IS-IS	232
Guidelines and Limitations for IS-IS	232
Default Settings for IS-IS	232
Configuring IS-IS	233
IS-IS Configuration Modes	233
Router Configuration Mode Example	233
Router Address Family Configuration Mode Example	233
Enabling the IS-IS Feature	234
Creating an IS-IS Instance	234
Restarting an IS-IS Instance	235
Shutting Down IS-IS	236
Configuring IS-IS on an Interface	236
Configuring IS-IS Authentication in an Area	237

Configuring IS-IS Authentication on an Interface	238
Configuring a Mesh Group	238
Configuring a Designated Intermediate System	239
Configuring Dynamic Host Exchange	239
Setting the Overload Bit	239
Configuring the Attached Bit	240
Configuring the Transient Mode for Hello Padding	240
Configuring a Summary Address	240
Configuring Redistribution	241
Limiting the Number of Redistributed Routes	242
Configuring the Administrative Distance of Routes	244
Disabling Strict Adjacency Mode	244
Configuring a Graceful Restart	245
Configuring Virtualization	246
Tuning IS-IS	247
Monitoring IS-IS	249
Configuration Examples for IS-IS	250
Related Documents for IS-IS	250
Standards for IS-IS	250
Feature History for IS-IS	250
<hr/>	
CHAPTER 11	Configuring Basic BGP 253
Finding Feature Information	253
Information About Basic BGP	253
BGP Autonomous Systems	254
4-Byte AS Number Support	254
Administrative Distance	254
BGP Peers	255
BGP Sessions	255
Dynamic AS Numbers for Prefix Peers	255
BGP Router Identifier	255
BGP Path Selection	256
BGP Path Selection - Comparing Pairs of Paths	256
BGP Path Selection - Determining the Order of Comparisons	258

BGP Path Selection - Determining the Best-Path Change Suppression	258
BGP and the Unicast RIB	259
BGP Prefix Independent Convergence	259
BGP PIC Feature Support Matrix	259
BGP PIC Core	260
BGP PIC Edge	260
BGP PIC Edge Unipath	261
BGP PIC Edge with Multipaths	262
BGP Virtualization	263
Prerequisites for BGP	263
Guidelines and Limitations for BGP	263
Default Settings	264
CLI Configuration Modes	265
Global Configuration Mode	265
Address Family Configuration Mode	265
Neighbor Configuration Mode	265
Neighbor Address Family Configuration Mode	266
Configuring Basic BGP	266
Enabling BGP	266
Creating a BGP Instance	267
Restarting a BGP Instance	268
Shutting Down BGP	268
Configuring BGP Peers	269
Configuring AS-4 Dot Notation	270
Configuring Dynamic AS Numbers for Prefix Peers	271
Configuring BGP PIC Edge	272
Clearing BGP Information	274
Verifying the Basic BGP Configuration	277
Monitoring BGP Statistics	279
Configuration Examples for Basic BGP	280
Related Documents for Basic BGP	280
MIBs	280
Feature History for BGP	280

CHAPTER 12

Configuring Advanced BGP	283
Finding Feature Information	283
Information About Advanced BGP	283
Peer Templates	284
Authentication	284
Route Policies and Resetting BGP Sessions	284
eBGP	285
BGP Next Hop Unchanged	285
iBGP	285
AS Confederations	286
Route Reflector	286
Capabilities Negotiation	287
Route Dampening	287
Load Sharing and Multipath	288
BGP Additional Paths	288
Route Aggregation	289
BGP Conditional Advertisement	289
BGP Next-Hop Address Tracking	290
Route Redistribution	291
BGP Support for Importing Routes from Default VRF	291
BGP Support for Exporting Routes to Default VRF	292
BFD	292
Tuning BGP	292
BGP Timers	292
Tuning the Best-Path Algorithm	292
Multiprotocol BGP	292
Graceful Restart and High Availability	293
Low Memory Handling	294
ISSU	294
Virtualization Support	294
Prerequisites for Advanced BGP	295
Guidelines and Limitations for Advanced BGP	295
Default Settings	296

Configuring Advanced BGP	297
Configuring BGP Session Templates	297
Configuring BGP Peer-Policy Templates	298
Configuring BGP Peer Templates	300
Configuring Prefix Peering	302
Configuring BGP Authentication	303
Resetting a BGP Session	303
Modifying the Next-Hop Address	304
Configuring BGP Next-Hop Address Tracking	304
Configuring Next-Hop Filtering	305
Disabling Capabilities Negotiation	305
Configuring BGP Additional Paths	305
Advertising the Capability of Sending and Receiving Additional Paths	305
Configuring the Sending and Receiving of Additional Paths	306
Configuring Advertised Paths	307
Configuring Additional Path Selection	308
Configuring eBGP	309
Disabling eBGP Single-Hop Checking	309
Configuring eBGP Multihop	309
Disabling a Fast External Fallover	309
Limiting the AS-path Attribute	310
Configuring Local AS Support	310
Configuring AS Confederations	311
Configuring Route Reflector	311
Configuring Next-Hops on Reflected Routes Using an Outbound Route-Map	313
Configuring Route Dampening	314
Configuring Load Sharing and ECMP	315
Configuring Maximum Prefixes	315
Configuring Dynamic Capability	316
Configuring Aggregate Addresses	316
Unsuppressing the Advertisement of Aggregated Routes	317
Configuring BGP Conditional Route Injection	317
Configuring BGP Conditional Advertisement	318
Configuring Route Redistribution	320

Advertising the Default Route	321
Configuring Route Import from Default VRF to any other VRF	322
Configuring Route Export from BGP VRF to Default VRF	322
Configuring Multiprotocol BGP	324
Configuring Policy-Based Administrative Distance	325
Tuning BGP	326
Configuring a Graceful Restart	329
Configuring Virtualization	330
Verifying the Advanced BGP Configuration	331
Displaying Advanced BGP Statistics	333
Related Documents	333
RFCs	334
MIBs	334
Feature History for Advanced BGP	334

CHAPTER 13**Configuring RIP 337**

Finding Feature Information	337
Information About RIP	337
RIPv2 Authentication	338
Split Horizon	338
Route Filtering	339
Route Summarization	339
Route Redistribution	339
Load Balancing	339
High Availability for RIP	339
Virtualization Support	340
Prerequisites for RIP	340
Guidelines and Limitations for RIP	340
Default Settings for RIP Parameters	340
Configuring RIP	340
Enabling RIP	340
Creating a RIP Instance	341
Restarting a RIP Instance	342
Configuring RIP on an Interface	343

Configuring RIP Authentication	343
Configuring a Passive Interface	344
Configuring Split Horizon with Poison Reverse	345
Configuring Route Summarization	346
Configuring Route Redistribution	346
Configuring Cisco NX-OS RIP for Compatibility with Cisco IOS RIP	347
Configuring Virtualization	348
Tuning RIP	350
Verifying the RIP Configuration	351
Displaying RIP Statistics	352
Configuration Examples for RIP	352
Related Documents for RIP	353
Standards for RIP	353
Feature History for RIP	353

CHAPTER 14

Configuring Static Routing	355
Finding Feature Information	355
Information About Static Routing	355
Administrative Distance	356
Directly Connected Static Routes	356
Fully Specified Static Routes	356
Floating Static Routes	356
Remote Next-Hops for Static Routes	356
Reliable Static Routing Backup Using Object Tracking Deployment	357
IP Service Level Agreements	357
BFD	357
Virtualization Support	357
Prerequisites for Static Routing	358
Guidelines and Limitations for Static Routing	358
Default Settings for Static Routing Parameters	358
Configuring Static Routing	358
Configuring a Static Route for IPv4	358
Configuring a Static Route for IPv6	359
Configuring a Static Route over a VLAN	360

Configuring Reliable Static Routing Backup Using Object Tracking	361
Configuring Virtualization for IPv4	362
Configuring Virtualization for IPv6	363
Verifying the Static Routing Configuration	364
Related Documents for Static Routing	364
Feature History for Static Routing	364

CHAPTER 15**Configuring the Interoperability of Modules for Unicast Routing 367**

Finding Feature Information	367
Configuring the Interoperability of Modules for Unicast Routing	367
Information About the Interoperability of Modules for Unicast Routing	368
Guidelines and Limitations for the Interoperability of Modules for Unicast Routing	368
Configuring the Interoperability of Modules for Unicast Routing	368
Verifying the Configuration for the Interoperability of Modules for Unicast Routing	369
Configuration Examples for the Interoperability of Modules for Unicast Routing	369
Related Documents for the Interoperability of Modules for Unicast Routing	370
Feature History for the Interoperability of Modules for Unicast Routing	370

CHAPTER 16**Configuring Layer 3 Virtualization 371**

Finding Feature Information	371
Information About Layer 3 Virtualization	371
VRF and Routing	373
VRF-Aware Services	373
Reachability	374
Filtering	374
Combining Reachability and Filtering	375
Guidelines and Limitations for VRF	375
Default Settings for VRF	376
Configuring VRFs	376
Creating a VRF	376
Assigning VRF Membership to an Interface	377
Configuring VRF Parameters for a Routing Protocol	378
Configuring VRF Aware Service	379
Setting the VRF Scope	379

Verifying the VRF Configuration	380
Configuration Examples for VRF	380
Related Documents for VRF	381
Standards for VRF	382
Feature History for VRF	382

CHAPTER 17

Managing the Unicast RIB and FIB	383
Finding Feature Information	383
Information About the Unicast RIB and FIB	383
Layer 3 Consistency Checker	384
Dynamic TCAM Allocation	384
Maximum TCAM Entries and FIB Scale Limits	385
Default Settings for the Unicast RIB and FIB	386
Managing the Unicast RIB and FIB	386
Displaying Module FIB Information	386
Configuring Load Sharing in the Unicast FIB	387
Configuring Per-Packet Load Sharing	389
Displaying Routing and Adjacency Information	390
Triggering the Layer 3 Consistency Checker	391
Clearing Forwarding Information in the FIB	392
Configuring Maximum Routes for the Unicast RIB	392
Estimating Memory Requirements for Routes	393
Clearing Routes in the Unicast RIB	394
Verifying the Unicast RIB and FIB	395
Related Documents for the Unicast RIB and FIB	395
Feature History for the Unicast RIB and FIB	395

CHAPTER 18

Configuring Route Policy Manager	397
Finding Feature Information	397
Information About Route Policy Manager	397
Prefix Lists	398
MAC Lists	398
Route Maps	398
Match Criteria	399

Set Changes	399
Access Lists	399
AS Numbers for BGP	400
AS-path Lists for BGP	400
Community Lists for BGP	400
Extended Community Lists for BGP	400
Route Redistribution and Route Maps	401
Route Map Support Matrix for Routing Protocols	401
Policy-Based Routing	404
Prerequisites for Route Policy Manager	405
Guidelines and Limitations	405
Default Settings for Route Policy Manager Parameters	405
Configuring Route Policy Manager	406
Configuring IP Prefix Lists	406
Configuring MAC Lists	407
Configuring AS-path Lists	408
Configuring Community Lists	408
Configuring Extended Community Lists	409
Optional Match Parameters for Route Maps	410
Optional Set Parameters for Route Maps	412
Verifying the Route Policy Manager Configuration	415
Configuration Examples for Route Policy Manager	416
Related Documents for Route Policy Manager	416
Standards for Route Policy Manager	416
Feature History for Route Policy Manager	416

CHAPTER 19

Configuring Policy-Based Routing	419
Finding Feature Information	419
Information About Policy Based Routing	419
Policy Route Maps	420
Set Criteria for Policy-Based Routing	420
Local Policy Routing	421
Route Map Support Matrix for Policy-Based Routing	421
Prerequisites for Policy-Based Routing	422

Guidelines and Limitations for Policy-Based Routing	422
Default Settings for Policy-Based Routing	424
Configuring Policy-Based Routing	424
Enabling the Policy-Based Routing	424
Configuring a Route Policy	424
Configuring Local Policy Routing	428
Configuring a Deny ACE	429
Verifying the Policy-Based Routing Configuration	430
Configuration Examples for Policy Based-Routing	430
Configuration Example for Local Policy Routing	430
Related Documents for Policy-Based Routing	431
Standards for Policy-Based Routing	431
Feature History for Policy-Based Routing	431

PART III
First-Hop Redundancy Protocols 433

CHAPTER 20
Configuring GLBP 435

Finding Feature Information	435
Information About GLBP	435
GLBP Active Virtual Gateway	436
GLBP Virtual MAC Address Assignment	436
GLBP Virtual Gateway Redundancy	436
GLBP Virtual Forwarder Redundancy	437
GLBP Authentication	438
GLBP Load Balancing and Tracking	438
High Availability and Extended Nonstop Forwarding	439
Virtualization Support	439
Prerequisites for GLBP	440
Guidelines and Limitations for GLBP	440
Default Settings for GLBP	440
Configuring GLBP	441
Enabling GLBP	441
Configuring GLBP Authentication	442
Configuring GLBP Load Balancing	443

Configuring GLBP Weighting and Tracking	443
Customizing GLBP	445
Configuring Extended Hold Timers for GLBP	447
Enabling a GLBP Group	447
Verifying the GLBP Configuration	448
Configuration Examples for GLBP	449
Related Documents for GLBP	449
Standards for GLBP	449
Feature History for GLBP	450

CHAPTER 21
Configuring HSRP 451

Finding Feature Information	451
Information About HSRP	451
HSRP for IPv4	453
HSRP for IPv6	453
HSRP for IPv6 Addresses	454
Multiple Group Optimization for HSRP	454
HSRP Versions	455
HSRP Authentication	455
HSRP Messages	455
HSRP Load Sharing	455
Object Tracking and HSRP	456
vPC and HSRP	456
vPC Peer Gateway and HSRP	457
FabricPath Anycast HSRP	457
BFD	457
High Availability and Extended Nonstop Forwarding	457
Virtualization Support	458
HSRP VIP	458
Prerequisites for HSRP	459
Guidelines and Limitations for HSRP	459
Default Settings for HSRP Parameters	461
Configuring HSRP	461
Enabling HSRP	461

Configuring the HSRP Version	462
Configuring an HSRP Group for IPv4	462
Configuring an HSRP Group for IPv6	464
Configuring an HSRP Master Group Task	465
Configuring an HSRP Slave Group	467
Configuring the HSRP Virtual MAC Address Manually	469
Configuring the HSRP Virtual MAC Address Using Burned-in MAC Address	469
Configuring MAC Address Table Reservation for HSRP	470
Clearing MAC Address Table Reservation for HSRP	472
Authenticating HSRP	473
Configuring HSRP Object Tracking	474
Configuring the HSRP Priority	475
Customizing HSRP in HSRP Configuration Mode	476
Customizing HSRP in Interface Configuration Mode	478
Configuring Extended Hold Timers for HSRP	478
Verifying the HSRP Configuration	479
Configuration Examples for HSRP	480
Related Documents for HSRP	480
MIBs	481
Feature History for HSRP	481

CHAPTER 22**Configuring VRRP 483**

Finding Feature Information	483
Information About VRRP	483
VRRP Operation	483
VRRP Benefits	485
Multiple VRRP Groups	485
VRRP Router Priority and Preemption	486
vPC and VRRP	487
VRRP Advertisements	487
VRRP Authentication	487
VRRP Tracking	487
VRRPv3 and VRRS	488
BFD for VRRP	488

High Availability	488
Virtualization Support	489
Guidelines and Limitations for VRRP	489
Default Settings for VRRP Parameters	490
Configuring VRRP	490
Enabling VRRP	490
Configuring VRRP Groups	491
Configuring VRRP Priority	492
Configuring VRRP Authentication	493
Configuring Time Intervals for Advertisement Packets	494
Disabling Preemption	495
Configuring VRRP Interface State Tracking	496
Enabling the VRRPv3 Feature	498
Creating VRRPv3 Groups	498
Configuring the Delay Period for FHRP Client Initialization	500
Configuring VRRPv3 Control Groups	501
Configuring VRRS Pathways	502
Verifying the VRRP Configuration	503
Monitoring VRRP Statistics	504
Configuration Example for VRRP	504
Related Documents for VRRP	506
Feature History for VRRP	506

CHAPTER 23

Configuring Object Tracking	507
Finding Feature Information	507
Information About Object Tracking	507
Object Track List	508
High Availability	509
Virtualization Support	509
Prerequisites for Object Tracking	509
Guidelines and Limitations for Object Tracking	509
Default Settings for Object Tracking Parameters	509
Configuring Object Tracking	510
Configuring Object Tracking for an Interface	510

Deleting a Tracking Object	511
Configuring Object Tracking for Route Reachability	511
Configuring an Object Track List with a Boolean Expression	512
Configuring an Object Track List with a Percentage Threshold	513
Configuring an Object Track List with a Weight Threshold	514
Configuring an Object Tracking Delay	515
Configuring Object Tracking for a Nondefault VRF	517
Verifying the Object Tracking Configuration	518
Configuration Example for Object Tracking	519
Related Documents for Object Tracking	519
Standards for Object Tracking	519
Feature History for Object Tracking	519

APPENDIX A	IETF RFCs Supported by Cisco NX-OS Unicast Features Release 6.x	521
	BGP RFCs	521
	First-Hop Redundancy Protocols RFCs	522
	IP Services RFCs	522
	IPv6 RFCs	523
	IS-IS RFCs	523
	OSPF RFCs	524
	RIP RFCs	524

APPENDIX B	Configuration Limits for Cisco NX-OS Layer 3 Unicast Features	527
-------------------	--	------------



Preface

This preface describes the audience, organization and conventions of the *Cisco Nexus 7000 Series NX-OS Unicast Routing Configuration Guide*. It also provides information on how to obtain related documentation.

-
- [Audience, on page xxix](#)
- [Document Conventions, on page xxix](#)
- [Related Documentation for Cisco Nexus 7000 Series NX-OS Software, on page xxx](#)
- [Documentation Feedback, on page xxxii](#)
- [Communications, Services, and Additional Information, on page xxxiii](#)

Audience

This publication is for network administrators who configure and maintain Cisco Nexus devices.

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.

Command descriptions use the following conventions:

Convention	Description
bold	Bold text indicates the commands and keywords that you enter literally as shown.
<i>Italic</i>	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.

Convention	Description
{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.
<i>variable</i>	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
<code>screen font</code>	Terminal sessions and information the switch displays are in screen font.
boldface screen font	Information you must enter is in boldface screen font.
<i>italic screen font</i>	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 for Cisco Nexus 7000 Series NX-OS Software

The entire Cisco Nexus 7000 Series NX-OS documentation set is available at the following URL:

<https://www.cisco.com/c/en/us/support/switches/nexus-7000-series-switches/series.html#~tab-documents>

Release Notes

The release notes are available at the following URL:

http://www.cisco.com/en/US/products/ps9402/prod_release_notes_list.html

Configuration Guides

These guides are available at the following URL:

http://www.cisco.com/en/US/products/ps9402/products_installation_and_configuration_guides_list.html

The documents in this category include:

- *Cisco Nexus 7000 Series NX-OS Configuration Examples*
- *Cisco Nexus 7000 Series NX-OS FabricPath Configuration Guide*
- *Cisco Nexus 7000 Series NX-OS Fundamentals Configuration Guide*
- *Cisco Nexus 7000 Series NX-OS Interfaces Configuration Guide*
- *Cisco Nexus 7000 Series NX-OS IP SLAs Configuration Guide*
- *Cisco Nexus 7000 Series NX-OS Layer 2 Switching Configuration Guide*
- *Cisco Nexus 7000 Series NX-OS LISP Configuration Guide*
- *Cisco Nexus 7000 Series NX-OS MPLS Configuration Guide*
- *Cisco Nexus 7000 Series NX-OS Multicast Routing Configuration Guide*
- *Cisco Nexus 7000 Series NX-OS OTV Configuration Guide*
- *Cisco Nexus 7000 Series NX-OS Quality of Service Configuration Guide*
- *Cisco Nexus 7000 Series NX-OS SAN Switching Guide*
- *Cisco Nexus 7000 Series NX-OS Security Configuration Guide*
- *Cisco Nexus 7000 Series NX-OS System Management Configuration Guide*
- *Cisco Nexus 7000 Series NX-OS Unicast Routing Configuration Guide*
- *Cisco Nexus 7000 Series NX-OS Verified Scalability Guide*
- *Cisco Nexus 7000 Series NX-OS Virtual Device Context Configuration Guide*
- *Cisco Nexus 7000 Series NX-OS Virtual Device Context Quick Start*
- *Cisco Nexus 7000 Series NX-OS OTV Quick Start Guide*
- *Cisco NX-OS FCoE Configuration Guide for Cisco Nexus 7000 and Cisco MDS 9500*
- *Cisco Nexus 2000 Series Fabric Extender Software Configuration Guide*

Command References

These guides are available at the following URL:

http://www.cisco.com/en/US/products/ps9402/prod_command_reference_list.html

The documents in this category include:

- *Cisco Nexus 7000 Series NX-OS Command Reference Master Index*
- *Cisco Nexus 7000 Series NX-OS FabricPath Command Reference*
- *Cisco Nexus 7000 Series NX-OS Fundamentals Command Reference*
- *Cisco Nexus 7000 Series NX-OS High Availability Command Reference*
- *Cisco Nexus 7000 Series NX-OS Interfaces Command Reference*
- *Cisco Nexus 7000 Series NX-OS Layer 2 Switching Command Reference*
- *Cisco Nexus 7000 Series NX-OS LISP Command Reference*
- *Cisco Nexus 7000 Series NX-OS MPLS Configuration Guide*
- *Cisco Nexus 7000 Series NX-OS Multicast Routing Command Reference*
- *Cisco Nexus 7000 Series NX-OS OTV Command Reference*
- *Cisco Nexus 7000 Series NX-OS Quality of Service Command Reference*
- *Cisco Nexus 7000 Series NX-OS SAN Switching Command Reference*
- *Cisco Nexus 7000 Series NX-OS Security Command Reference*
- *Cisco Nexus 7000 Series NX-OS System Management Command Reference*
- *Cisco Nexus 7000 Series NX-OS Unicast Routing Command Reference*
- *Cisco Nexus 7000 Series NX-OS Virtual Device Context Command Reference*
- *Cisco NX-OS FCoE Command Reference for Cisco Nexus 7000 and Cisco MDS 9500*

Other Software Documents

You can locate these documents starting at the following landing page:

<https://www.cisco.com/c/en/us/support/switches/nexus-7000-series-switches/series.html#~tab-documents>

- *Cisco Nexus 7000 Series NX-OS MIB Quick Reference*
- *Cisco Nexus 7000 Series NX-OS Software Upgrade and Downgrade Guide*
- *Cisco Nexus 7000 Series NX-OS Troubleshooting Guide*
- *Cisco NX-OS Licensing Guide*
- *Cisco NX-OS System Messages Reference*
- *Cisco NX-OS Interface User Guide*

Documentation Feedback

To provide technical feedback on this document, or to report an error or omission, please send your comments to: .

We appreciate your feedback.

Communications, Services, and Additional Information

- To receive timely, relevant information from Cisco, sign up at [Cisco Profile Manager](#).
- To get the business impact you're looking for with the technologies that matter, visit [Cisco Services](#).
- To submit a service request, visit [Cisco Support](#).
- To discover and browse secure, validated enterprise-class apps, products, solutions and services, visit [Cisco Marketplace](#).
- To obtain general networking, training, and certification titles, visit [Cisco Press](#).
- To find warranty information for a specific product or product family, access [Cisco Warranty Finder](#).

Cisco Bug Search Tool

[Cisco Bug Search Tool](#) (BST) is a web-based tool that acts as a gateway to the Cisco bug tracking system that maintains a comprehensive list of defects and vulnerabilities in Cisco products and software. BST provides you with detailed defect information about your products and software.



CHAPTER 1

New and Changed Information

- [New and Changed Information](#), on page 1

New and Changed Information

Table 1: New and Changed Unicast Routing Features

Feature	Description	Changed in Release
Support for M3 module	Added support for M3 module	7.3(0)DX(1)
BGP PIC Edge	Added support for IPv6 address family	7.3(0)D1(1)
OSPFv3 IPsec Authentication	Added IPsec authentication of OSPFv3 packets	7.3(0)D1(1)
Advanced BGP	Added support for exporting routes to default VRF	7.3(0)D1(1)
vIP HSRP Enhancement	Added support for HSRP vIP configuration.	7.2(0)D1(1)



CHAPTER 2

Overview

This chapter contains the following sections:

- [Overview, on page 3](#)

Overview

This chapter introduces the underlying concepts for the Layer 3 unicast routing protocols in Cisco NX-OS.

Licensing Requirements

For a complete explanation of Cisco NX-OS licensing recommendations and how to obtain and apply licenses, see the [Cisco NX-OS Licensing Guide](#).

Information About Layer 3 Unicast Routing

Layer 3 unicast routing involves two basic activities: determining optimal routing paths and packet switching. You can use routing algorithms to calculate the optimal path from the router to a destination. This calculation depends on the algorithm selected, route metrics, and other considerations such as load balancing and alternate path discovery.

Routing Fundamentals

Routing protocols use a metric to evaluate the best path to the destination. A metric is a standard of measurement, such as a path bandwidth, that routing algorithms use to determine the optimal path to a destination. To aid path determination, routing algorithms initialize and maintain routing tables that contain route information such as the IP destination address, the address of the next router, or the next hop. Destination and next-hop associations tell a router that an IP destination can be reached optimally by sending the packet to a particular router that represents the next hop on the way to the final destination. When a router receives an incoming packet, it checks the destination address and attempts to associate this address with the next hop.

Routing tables can contain other information, such as the data about the desirability of a path. Routers compare metrics to determine optimal routes, and these metrics differ depending on the design of the routing algorithm used.

Routers communicate with one another and maintain their routing tables by transmitting a variety of messages. The routing update message is one such message that consists of all or a portion of a routing table. By analyzing routing updates from all other routers, a router can build a detailed picture of the network topology. A link-state

advertisement, which is another example of a message sent between routers, informs other routers of the link state of the sending router. You can also use link information to enable routers to determine optimal routes to network destinations.


Packet Switching

In packet switching, a host determines that it must send a packet to another host. Having acquired a router address by some means, the source host sends a packet that is addressed specifically to the router physical (Media Access Control [MAC]-layer) address but with the IP (network layer) address of the destination host.

The router examines the destination IP address and tries to find the IP address in the routing table. If the router does not know how to forward the packet, it typically drops the packet. If the router knows how to forward the packet, it changes the destination MAC address to the MAC address of the next-hop router and transmits the packet.

The next hop might be the ultimate destination host or another router that executes the same switching decision process. As the packet moves through the internetwork, its physical address changes, but its protocol address remains constant (see the following figure).

Figure 1: Packet Header Updates Through a Network



18/29/78

Routing Metrics

Routing algorithms use many different metrics to determine the best route. Sophisticated routing algorithms can base route selection on multiple metrics.

Path Length

The path length is the most common routing metric. Some routing protocols allow you to assign arbitrary costs to each network link. In this case, the path length is the sum of the costs associated with each link traversed. Other routing protocols define the hop count, which is a metric that specifies the number of passes through internetworking products, such as routers, that a packet must take from a source to a destination.

Reliability

The reliability, in the context of routing algorithms, is the dependability (in terms of the bit-error rate) of each network link. Some network links might go down more often than others. After a network fails, certain network links might be repaired more easily or more quickly than other links. The reliability factors that you can take into account when assigning the reliability rating are arbitrary numeric values that you usually assign to network links.

Routing Delay

The routing delay is the length of time required to move a packet from a source to a destination through the internetwork. The delay depends on many factors, including the bandwidth of intermediate network links, the port queues at each router along the way, the network congestion on all intermediate network links, and the physical distance that the packet must travel. Because the routing delay is a combination of several important variables, it is a common and useful metric.

Bandwidth

The bandwidth is the available traffic capacity of a link. For example, a 10-Gigabit Ethernet link is preferable to a 1-Gigabit Ethernet link. Although the bandwidth is the maximum attainable throughput on a link, routes through links with greater bandwidth do not necessarily provide better routes than routes through slower links. For example, if a faster link is busier, the actual time required to send a packet to the destination could be greater.

Load

The load is the degree to which a network resource, such as a router, is busy. You can calculate the load in a variety of ways, including CPU usage and packets processed per second. Monitoring these parameters on a continual basis can be resource intensive.

Communication Cost

The communication cost is a measure of the operating cost to route over a link. The communication cost is another important metric, especially if you do not care about performance as much as operating expenditures. For example, the line delay for a private line might be longer than a public line, but you can send packets over your private line rather than through the public lines that cost money for usage time.

Router IDs

Each routing process has an associated router ID. You can configure the router ID to any interface in the system. If you do not configure the router ID, Cisco NX-OS selects the router ID based on the following criteria:

- Cisco NX-OS prefers loopback0 over any other interface. If loopback0 does not exist, then Cisco NX-OS prefers the first loopback interface over any other interface type.
- If you have not configured a loopback interface, Cisco NX-OS uses the first interface in the configuration file as the router ID. If you configure any loopback interface after Cisco NX-OS selects the router ID, the loopback interface becomes the router ID. If the loopback interface is not loopback0 and you configure loopback0 with an IP address, the router ID changes to the IP address of loopback0.
- If the interface that the router ID is based on changes, that new IP address becomes the router ID. If any other interface changes its IP address, there is no router ID change.

Autonomous Systems

An autonomous system (AS) is a network controlled by a single technical administration entity. Autonomous systems divide global external networks into individual routing domains, where local routing policies are applied. This organization simplifies routing domain administration and simplifies consistent policy configuration.

Each autonomous system can support multiple interior routing protocols that dynamically exchange routing information through route redistribution. The Regional Internet Registries assign a unique number to each public autonomous system that directly connects to the Internet. This autonomous system number (AS number) identifies both the routing process and the autonomous system.

The Border Gateway Protocol (BGP) supports 4-byte AS numbers that can be represented in asplain and asdot notations:

- asplain—A decimal value notation where both 2-byte and 4-byte AS numbers are represented by their decimal value. For example, 65526 is a 2-byte AS number, and 234567 is a 4-byte AS number.
- asdot—An AS dot notation where 2-byte AS numbers are represented by their decimal value and 4-byte AS numbers are represented by a dot notation. For example, 2-byte AS number 65526 is represented as 65526, and 4-byte AS number 65546 is represented as 1.10.

The BGP 4-byte AS number capability is used to propagate 4-byte-based AS path information across BGP speakers that do not support 4-byte AS numbers. Beginning with Cisco NX-OS Release 6.2(2), you can configure 4-byte AS numbers in asdot notation. The default value is asplain.

The following table lists the AS number ranges.

Table 2: AS Numbers

2-Byte Numbers	4-Byte Numbers in AS.dot Notation	4-Byte Numbers in plaintext Notation	Purpose
1 to 64511	N/A	1 to 64511	Public AS (assigned by RIR) 1
64512 to 65534	N/A	64512 to 65534	Private AS (assigned by local administrator)
65535	N/A	65535	Reserved
N/A	1.0 to 65535.65535	65536 to 4294967295	Public AS (assigned by RIR)

¹ RIR=Regional Internet Registries



Note RFC 5396 is partially supported. The asplain and asdot notations are supported, but the asdot+ notation is not.

Private autonomous system numbers are used for internal routing domains but must be translated by the router for traffic that is routed out to the Internet. You should not configure routing protocols to advertise private

autonomous system numbers to external networks. By default, Cisco NX-OS does not remove private autonomous system numbers from routing updates.



Note The autonomous system number assignment for public and private networks is governed by the Internet Assigned Number Authority (IANA). For information about autonomous system numbers, including the reserved number assignment, or to apply to register an autonomous system number, refer to the following URL: <http://www.iana.org/>

Convergence

A key aspect to measure for any routing algorithm is how much time a router takes to react to network topology changes. When a part of the network changes for any reason, such as a link failure, the routing information in different routers might not match. Some routers will have updated information about the changed topology, while other routers will still have the old information. The convergence is the amount of time before all routers in the network have updated, matching routing information. The convergence time varies depending on the routing algorithm. Fast convergence minimizes the chance of lost packets caused by inaccurate routing information.

Load Balancing and Equal Cost Multipath

Routing protocols can use load balancing or equal cost multipath (ECMP) to share traffic across multiple paths. When a router learns multiple routes to a specific network, it installs the route with the lowest administrative distance in the routing table. If the router receives and installs multiple paths with the same administrative distance and cost to a destination, load balancing can occur. Load balancing distributes the traffic across all the paths, sharing the load. The number of paths used is limited by the number of entries that the routing protocol puts in the routing table. Cisco NX-OS supports up to 16 paths to a destination. Starting from Cisco NX-OS Release 8.4(1), the BGP feature supports up to 64 paths to a destination on M3- and F3-Series I/O modules. Starting from Cisco NX-OS Release 8.4(2), the BGP feature supports up to 64 paths to a destination on F4-Series I/O modules.

The Enhanced Interior Gateway Routing Protocol (EIGRP) also supports unequal cost load balancing.

Route Redistribution

If you have multiple routing protocols configured in your network, you can configure these protocols to share routing information by configuring route redistribution in each protocol. For example, you can configure the Open Shortest Path First (OSPF) protocol to advertise routes learned from the Border Gateway Protocol (BGP). You can also redistribute static routes into any dynamic routing protocol. The router that is redistributing routes from another protocol sets a fixed route metric for those redistributed routes, which prevents incompatible route metrics between the different routing protocols. For example, routes redistributed from EIGRP into OSPF are assigned a fixed link cost metric that OSPF understands.



Note You are required to use route maps when you configure redistribution of routing information.

Route redistribution also uses an administrative distance to distinguish between routes learned from two different routing protocols. The preferred routing protocol is given a lower administrative distance so that its routes are picked over routes from another protocol with a higher administrative distance assigned.

Administrative Distance

An administrative distance is a rating of the trustworthiness of a routing information source. A higher value indicates a lower trust rating. Typically, a route can be learned through more than one protocol. Administrative distance is used to discriminate between routes learned from more than one protocol. The route with the lowest administrative distance is installed in the IP routing table.


Stub Routing

You can use stub routing in a hub-and-spoke network topology, where one or more end (stub) networks are connected to a remote router (the spoke) that is connected to one or more distribution routers (the hub). The remote router is adjacent only to one or more distribution routers. The only route for IP traffic to follow into the remote router is through a distribution router. This type of configuration is commonly used in WAN topologies in which the distribution router is directly connected to a WAN. The distribution router can be connected to many more remote routers. Often, the distribution router is connected to 100 or more remote routers. In a hub-and-spoke topology, the remote router must forward all nonlocal traffic to a distribution router, so it becomes unnecessary for the remote router to hold a complete routing table. Generally, the distribution router sends only a default route to the remote router.

Only specified routes are propagated from the remote (stub) router. The stub router responds to all queries for summaries, connected routes, redistributed static routes, external routes, and internal routes with the message "inaccessible." A router that is configured as a stub sends a special peer information packet to all neighboring routers to report its status as a stub router.

Any neighbor that receives a packet that informs it of the stub status does not query the stub router for any routes, and a router that has a stub peer does not query that peer. The stub router depends on the distribution router to send the proper updates to all peers. The following figure shows a simple hub-and-spoke network.

Figure 2: Simple Hub-and-Spoke Network



Stub routing does not prevent routes from being advertised to the remote router. This figure shows that the remote router can access the corporate network and the Internet through the distribution router only. A full route table on the remote router, in this example, serves no functional purpose because the path to the corporate network and the Internet is always through the distribution router. A larger route table only increases the amount of memory consumed by the remote router. The bandwidth and memory used can be lessened by summarizing and filtering routes in the distribution router. In this network topology, the remote router does not need to receive routes that have been learned from other networks because the remote router must send all non-local traffic, regardless of its destination, to the distribution router. To configure a true stub network, you should configure the distribution router to send only a default route to the remote router.

OSPF supports stub areas and EIGRP supports stub routers.

Routing Algorithms

Routing algorithms determine how a router gathers and reports reachability information, how it deals with topology changes, and how it determines the optimal route to a destination. Various types of routing algorithms exist, and each algorithm has a different impact on network and router resources. Routing algorithms use a variety of metrics that affect calculation of optimal routes. You can classify routing algorithms by type, such as static or dynamic, and interior or exterior.

Static Routes and Dynamic Routing Protocols

Static routes are route table entries that you manually configure. These static routes do not change unless you reconfigure them. Static routes are simple to design and work well in environments where network traffic is relatively predictable and where network design is relatively simple.

Because static routing systems cannot react to network changes, you should not use them for large, constantly changing networks. Most routing protocols today use dynamic routing algorithms that adjust to changing network circumstances by analyzing incoming routing update messages. If the message indicates that a network change has occurred, the routing software recalculates routes and sends out new routing update messages. These messages permeate the network, triggering routers to rerun their algorithms and change their routing tables accordingly.

You can supplement dynamic routing algorithms with static routes where appropriate. For example, you should configure each subnetwork with a static route to the IP default gateway or router of last resort (a router to which all unroutable packets are sent).

Interior and Exterior Gateway Protocols

You can separate networks into unique routing domains or autonomous systems. An autonomous system is a portion of an internetwork under common administrative authority that is regulated by a particular set of administrative guidelines. Routing protocols that route between autonomous systems are called exterior gateway protocols or interdomain protocols. The Border Gateway Protocol (BGP) is an example of an exterior gateway protocol. Routing protocols used within an autonomous system are called interior gateway protocols or intradomain protocols. EIGRP and OSPF are examples of interior gateway protocols.

Distance Vector Protocols

Distance vector protocols use distance vector algorithms (also known as Bellman-Ford algorithms) that call for each router to send all or some portion of its routing table to its neighbors. Distance vector algorithms define routes by distance (for example, the number of hops to the destination) and direction (for example, the next-hop router). These routes are then broadcast to the directly connected neighbor routers. Each router uses these updates to verify and update the routing tables.

To prevent routing loops, most distance vector algorithms use split horizon with poison reverse which means that the routes learned from an interface are set as unreachable and advertised back along the interface that they were learned on during the next periodic update. This process prevents the router from seeing its own route updates coming back.

Distance vector algorithms send updates at fixed intervals but can also send updates in response to changes in route metric values. These triggered updates can speed up the route convergence time. The Routing Information Protocol (RIP) is a distance vector protocol.

Link-State Protocols

The link-state protocols, also known as shortest path first (SPF), share information with neighboring routers. Each router builds a link-state advertisement (LSA) that contains information about each link and directly connected neighbor router.

Each LSA has a sequence number. When a router receives an LSA and updates its link-state database, the LSA is flooded to all adjacent neighbors. If a router receives two LSAs with the same sequence number (from the same router), the router does not flood the last LSA that it received to its neighbors because it wants to prevent an LSA update loop. Because the router floods the LSAs immediately after it receives them, the convergence time for link-state protocols is minimized.

Discovering neighbors and establishing adjacency is an important part of a link state protocol. Neighbors are discovered using special Hello packets that also serve as keepalive notifications to each neighbor router. Adjacency is the establishment of a common set of operating parameters for the link-state protocol between neighbor routers.

The LSAs received by a router are added to the router's link-state database. Each entry consists of the following parameters:

- Router ID (for the router that originated the LSA)
- Neighbor ID
- Link cost
- Sequence number of the LSA
- Age of the LSA entry

The router runs the SPF algorithm on the link-state database, building the shortest path tree for that router. This SPF tree is used to populate the routing table.

In link-state algorithms, each router builds a picture of the entire network in its routing tables. The link-state algorithms send small updates everywhere, while distance vector algorithms send larger updates only to neighboring routers.

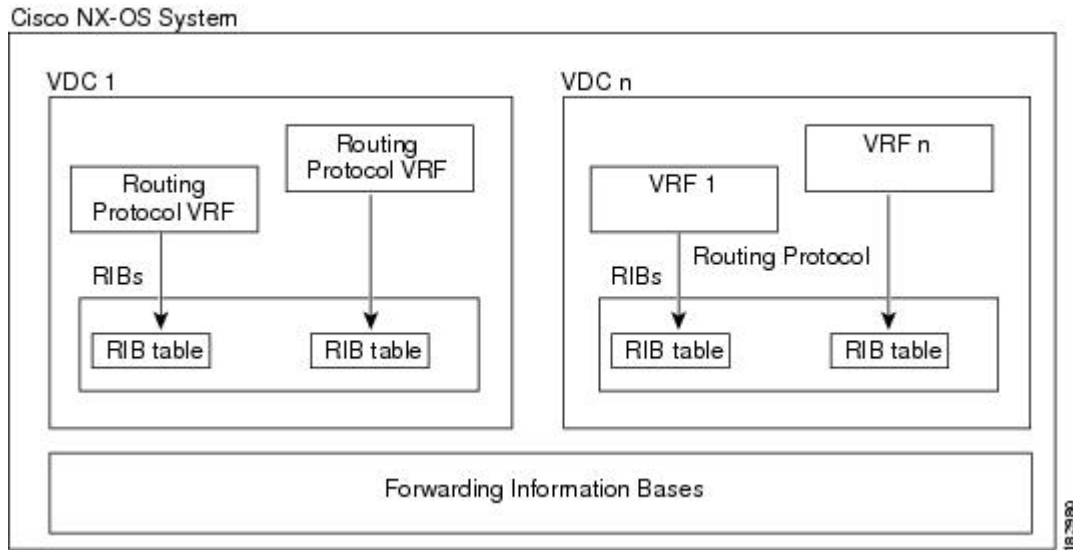
Because they converge more quickly, link-state algorithms are less likely to cause routing loops than distance vector algorithms. However, link-state algorithms require more CPU power and memory than distance vector algorithms and they can be more expensive to implement and support. Link-state protocols are generally more scalable than distance vector protocols.

OSPF is an example of a link-state protocol.

Layer 3 Virtualization

Cisco NX-OS uses a virtual device context (VDC) to provide separate management domains per VDC and software fault isolation. Each VDC supports multiple virtual routing and forwarding instances and multiple routing information bases (RIBs) to support multiple address domains. Each VRF is associated with a RIB and this information is collected by the Forwarding Information Base (FIB). The following figure shows the relationship between a VDC, a VRF, and a Cisco NX-OS device.

Figure 3: Layer 3 Virtualization Example



A VRF represents a Layer 3 addressing domain. Each Layer 3 interface (logical or physical) belongs to one VRF. A VRF belongs to one VDC. Each VDC can support multiple VRFs.

See the *Cisco Nexus 7000 Series NX-OS Virtual Device Context Configuration Guide* for information about VDCs.


Cisco NX-OS Forwarding Architecture

The Cisco NX-OS forwarding architecture is responsible for processing all routing updates and populating the forwarding information to all modules in the chassis.

Unicast RIB

The Cisco NX-OS forwarding architecture consists of multiple components, as shown in the following figure.

Figure 4: Cisco NX-OS Forwarding Architecture



The unicast RIB exists on the active supervisor. It maintains the routing table with directly connected routes, static routes, and routes learned from dynamic unicast routing protocols. The unicast RIB also collects adjacency information from sources such as the Address Resolution Protocol (ARP). The unicast RIB determines the

best next hop for a given route and populates the unicast forwarding information base (FIB) by using the services of the unicast FIB distribution module (FDM).

Each dynamic routing protocol must update the unicast RIB for any route that has timed out. The unicast RIB then deletes that route and recalculates the best next hop for that route (if an alternate path is available).

Adjacency Manager

The adjacency manager exists on the active supervisor and maintains adjacency information for different protocols including ARP, Neighbor Discovery Protocol (NDP), and static configuration. The most basic adjacency information is the Layer 3 to Layer 2 address mapping discovered by these protocols. Outgoing Layer 2 packets use the adjacency information to complete the Layer 2 header.

The adjacency manager can trigger ARP requests to find a particular Layer 3 to Layer 2 mapping. The new mapping becomes available when the corresponding ARP reply is received and processed. For IPv6, the adjacency manager finds the Layer 3 to Layer 2 mapping information from NDP.

Unicast Forwarding Distribution Module

The unicast Forwarding Distribution Module (FDM) exists on the active supervisor and distributes the forwarding path information from the unicast RIB and other sources. The unicast RIB generates forwarding information that the unicast FIB programs into the hardware forwarding tables on the standby supervisor and the modules. The unicast FDM also downloads the FIB information to newly inserted modules.

The unicast FDM gathers adjacency information, rewrite information, and other platform-dependent information when updating routes in the unicast FIB. The adjacency and rewrite information consists of interface, next hop, and Layer 3 to Layer 2 mapping information. The interface and next-hop information is received in route updates from the unicast RIB. The Layer 3 to Layer 2 mapping is received from the adjacency manager.

FIB

The unicast FIB exists on supervisors and switching modules and builds the information used for the hardware forwarding engine. The unicast FIB receives route updates from the unicast FDM and sends the information to be programmed in the hardware forwarding engine. The unicast FIB controls the addition, deletion, and modification of routes, paths, and adjacencies.

The unicast FIBs are maintained on a per-VRF and per-address-family basis, that is, one for IPv4 and one for IPv6 for each configured VRF. Based on route update messages, the unicast FIB maintains a per-VRF prefix and next-hop adjacency information database. The next-hop adjacency data structure contains the next-hop IP address and the Layer 2 rewrite information. Multiple prefixes could share a next-hop adjacency information structure.

Hardware Forwarding

Cisco NX-OS supports distributed packet forwarding. The ingress port takes relevant information from the packet header and passes the information to the local switching engine. The local switching engine does the Layer 3 lookup and uses this information to rewrite the packet header. The ingress module forwards the packet to the egress port. If the egress port is on a different module, the packet is forwarded using the switch fabric to the egress module. The egress module does not participate in the Layer 3 forwarding decision.

The forwarding tables are identical on the supervisor and all the modules.

You also use the **show platform fib** or **show platform forwarding** commands to display details on hardware forwarding.

Software Forwarding

The software forwarding path in Cisco NX-OS is used mainly to handle features that are not supported in the hardware or to handle errors encountered during the hardware processing. Typically, packets with IP options or packets that need fragmentation are passed to the CPU on the active supervisor. All packets that should be switched in the software or terminated go to the supervisor. The supervisor uses the information provided by the unicast RIB and the adjacency manager to make the forwarding decisions. The module is not involved in the software forwarding path.

Software forwarding is controlled by control plane policies and rate limiters. For more information, see the *Cisco Nexus 7000 Series NX-OS Security Configuration Guide*.

Layer 3 Interoperation with the N7K-F132-15 Module



Note You must install one of the N7K-M Series modules in the Cisco Nexus 7000 Series chassis to run Layer 3 routing with the N7K-F132-15 module. You must have interfaces from both the M Series and the N7K-F132-15 modules in the same VDC. (See the *Cisco Nexus 7000 Series NX-OS Virtual Device Context Configuration Guide* for more information about VDCs.)



Note You cannot use F2 Series modules in the Cisco Nexus 7000 Series chassis to run Layer 3 routing with the N7K-F132-15 module

Layer 3 routing functionality comes up automatically when you have one of the N7K-M Series modules installed in the chassis with the N7K-F132-15 module. You would usually position a chassis with both the N7K-F132-15 and M Series modules, or a mixed chassis, at the boundary between the Layer 2 and Layer 3 networks.

You must configure a VLAN interface for each VLAN on the N7K-F132-15 module that you want to use the proxy-routing functionality in a mixed chassis. (See the *Cisco Nexus 7000 Series NX-OS Interfaces Configuration Guide* for information about configuring VLAN interfaces.)

By default, all of the physical interfaces on the N7K-M series modules in the VDC become proxy routing ports for the VLANs that are configured with VLAN interfaces on the Layer 2-only N7K-F132-15 module in the same VDC. The physical interfaces on the M Series module can be administratively down and still pass traffic as proxy forwarding.

Packets that enter an interface on the N7K-F132-15 module are automatically forwarded to one of the interfaces on the M Series modules in the same VDC to be routed. The interface on the M Series module also performs egress replication for Layer 3 multicast packets that enter an interface on the N7K-F132-15 module in the same VDC.

Because the Layer 3 (proxy routing) traffic from the N7K-F132-15 modules adds to the traffic that the M Series modules are already processing, the device automatically provides load balancing for the total traffic load among the front panel ports of the available M Series modules in the VDC. If you add or remove interfaces to the M Series modules in the VDC, the device automatically rebalances the traffic. Note that proxy routing is sharing the forwarding capacity of the M Series modules. Removing interfaces reduces the amount of capacity available.

Instead of using the automatically configured proxy-routing interfaces on the M Series modules, you can optionally configure which interfaces on the M Series modules in the VDC performs proxy routing.

Summary of Layer 3 Routing Features

This section provides a brief introduction to the Layer 3 unicast features and protocols supported in Cisco NX-OS.

IPv4 and IPv6

Layer 3 uses either the IPv4 or IPv6 protocol. IPv6 is a new IP protocol designed to replace IPv4, the Internet protocol that is predominantly deployed and used throughout the world. IPv6 increases the number of network address bits from 32 bits (in IPv4) to 128 bits.

IP Services

IP Services includes Dynamic Host Configuration Protocol (DHCP) and Domain Name System (DNS Client) clients.

OSPF

The Open Shortest Path First (OSPF) protocol is a link-state routing protocol used to exchange network reachability information within an autonomous system. Each OSPF router advertises information about its active links to its neighbor routers. Link information consists of the link type, the link metric, and the neighbor router that is connected to the link. The advertisements that contain this link information are called link-state advertisements.

EIGRP

The Enhanced Interior Gateway Routing Protocol (EIGRP) is a unicast routing protocol that has the characteristics of both distance vector and link-state routing protocols. It is an improved version of IGRP, which is a Cisco proprietary routing protocol. EIGRP relies on its neighbors to provide the routes. It constructs the network topology from the routes advertised by its neighbors, similar to a link-state protocol, and uses this information to select loop-free paths to destinations.

IS-IS

The Intermediate System-to-Intermediate System (IS-IS) protocol is an intradomain Open System Interconnection (OSI) dynamic routing protocol specified in the International Organization for Standardization (ISO) 10589. The IS-IS routing protocol is a link-state protocol. IS-IS features are as follows:

- Hierarchical routing
- Classless behavior
- Rapid flooding of new information
- Fast Convergence
- Very scalable

BGP

The Border Gateway Protocol (BGP) is an inter-autonomous system routing protocol. A BGP router advertises network reachability information to other BGP routers using Transmission Control Protocol (TCP) as its reliable transport mechanism. The network reachability information includes the destination network prefix,

a list of autonomous systems that needs to be traversed to reach the destination, and the next-hop router. Reachability information contains additional path attributes such as preference to a route, origin of the route, community and others..

RIP

The Routing Information Protocol (RIP) is a distance-vector protocol that uses a hop count as its metric. RIP is widely used for routing traffic in the global Internet and is an Interior Gateway Protocol (IGP), which means that it performs routing within a single autonomous system.

Static Routing

Static routing allows you to enter a fixed route to a destination. This feature is useful for small networks where the topology is simple. Static routing is also used with other routing protocols to control default routes and route distribution.

Layer 3 Virtualization

Virtualization allows you to share physical resources across separate management domains. Cisco NX-OS supports Virtual Device Contexts (VDCs) that allow you to create separate virtual systems within a Cisco NX-OS system. Each VDC is isolated from the others, which means that a problem in one VDC does not affect any other VDCs. VDCs are also secure from each other. You can assign separate network operators to each VDC and these network operators cannot control or view the configuration of a different VDC.

Cisco NX-OS also supports Layer 3 virtualization with virtual routing and forwarding (VRF). VRF provides a separate address domain for configuring Layer 3 routing protocols.

Route Policy Manager

The Route Policy Manager provides a route filtering capability in Cisco NX-OS. It uses route maps to filter routes distributed across various routing protocols and between different entities within a given routing protocol. Filtering is based on specific match criteria, which is similar to packet filtering by access control lists.

Policy-Based Routing

Policy-based routing uses the Route Policy Manager to create policy route filters. These policy route filters can forward a packet to a specified next hop based on the source of the packet or other fields in the packet header. Policy routes can be linked to extended IP access lists so that routing might be based on protocol types and port numbers.

First Hop Redundancy Protocols

First hop redundancy protocols (FHRP), such as Gateway Load Balancing Protocol (GLBP), Hot Standby Router Protocol (HSRP), and Virtual Router Redundancy Protocol (VRRP), allow you to provide redundant connections to your hosts. If an active first-hop router fails, the FHRP automatically selects a standby router to take over. You do not need to update the hosts with new IP addresses since the address is virtual and shared between each router in the FHRP group.

Object Tracking

Object tracking allows you to track specific objects on the network, such as the interface line protocol state, IP routing, and route reachability, and take action when the tracked object's state changes. This feature allows you to increase the availability of the network and shorten the recovery time if an object state goes down.

Related Documents for Layer 3 Unicast Routing

Feature Name	Feature Information
Layer 3 features	<i>Cisco NEXUS 7000 Series NX-OS Multicast Routing Configuration Guide</i> <i>Cisco NEXUS 7000 Series NX-OS High Availability and Redundancy Guide</i> <i>Cisco NEXUS 7000 Series NX-OS Virtual Device Context</i> Exploring Autonomous System Numbers: http://www.cisco.com/web/about/ac123/ac147/archived_issues/ipj_9-1/autonomous_system_numbers.html



PART I

IP

- [Configuring IPv4, on page 21](#)
- [Configuring IPv6, on page 45](#)
- [Configuring DNS, on page 75](#)
- [Configuring WCCPv2, on page 81](#)



CHAPTER 3

Configuring IPv4

This chapter contains the following sections:

- [Finding Feature Information, on page 21](#)
- [Information About IPv4, on page 21](#)
- [Virtualization Support for IPv4, on page 26](#)
- [Prerequisites for IPv4, on page 26](#)
- [Guidelines and Limitations for IPv4, on page 27](#)
- [Default Settings for IPv4 Parameters, on page 27](#)
- [Configuring IPv4, on page 27](#)
- [Verifying the IPv4 Configuration, on page 38](#)
- [Configuration Examples for IPv4, on page 39](#)
- [Related Documents for IPv4, on page 42](#)
- [Standards for IPv4, on page 42](#)
- [Feature History for IPv4, on page 43](#)

Finding Feature Information

Your software release might not support all the features documented in this module. For the latest caveats and feature information, see the Bug Search Tool at <https://tools.cisco.com/bugsearch/> and the release notes for your software release. To find information about the features documented in this module, and to see a list of the releases in which each feature is supported, see the "New and Changed Information" chapter or the Feature History table in this chapter.

Information About IPv4

You can configure IP on the device to assign IP addresses to network interfaces. When you assign IP addresses, you enable the interfaces and allow communication with the hosts on those interfaces.

You can configure an IP address as primary or secondary on a device. An interface can have one primary IP address and multiple secondary addresses. All networking device on an interface should share the same primary IP address because the packets that are generated by the device always use the primary IPv4 address. Each IPv4 packet is based on the information from a source or destination IP address.

You can use a subnet to mask the IP addresses. A mask is used to determine what subnet an IP address belongs to. An IP address contains the network address and the host address. A mask identifies the bits that denote

the network number in an IP address. When you use the mask to subnet a network, the mask is then referred to as a subnet mask. Subnet masks are 32-bit values that allow the recipient of IP packets to distinguish the network ID portion of the IP address from the host ID portion of the IP address.

The IP feature is responsible for handling IPv4 packets that terminate in the supervisor module, as well as forwarding of IPv4 packets, which includes IPv4 unicast/multicast route lookup, reverse path forwarding (RPF) checks, and software access control list/policy-based routing (ACL/PBR) forwarding. The IP feature also manages the network interface IP address configuration, duplicate address checks, static routes, and packet send/receive interface for IP clients.

Multiple IPv4 Addresses

Cisco NX-OS supports multiple IP addresses per interface. You can specify an unlimited number of secondary addresses for a variety of situations.

The most common situations are as follows:

- When there are not enough host IP addresses for a particular network interface. For example, if your subnetting allows up to 254 hosts per logical subnet, but on one physical subnet you must have 300 host addresses, then you can use secondary IP addresses on the routers or access servers to allow you to have two logical subnets that use one physical subnet.
- Two subnets of a single network might otherwise be separated by another network. You can create a single network from subnets that are physically separated by another network by using a secondary address. In these instances, the first network is extended, or layered on top of the second network. A subnet cannot appear on more than one active interface of the router at a time.



Note If any device on a network segment uses a secondary IPv4 address, other devices on that same network segment that require a secondary address must use a secondary address from the same network or subnet. The inconsistent use of secondary addresses on a network segment can quickly cause routing loops.

Address Resolution Protocol

Networking devices and Layer 3 switches use Address Resolution Protocol (ARP) to map IP (network layer) addresses to (Media Access Control [MAC]-layer) addresses to enable IP packets to be sent across networks. Before a device sends a packet to another device, it looks in its own ARP cache to see if there is a MAC address and corresponding IP address for the destination device. If there is no entry, the source device sends a broadcast message to every device on the network.

Each device compares the IP address to its own. Only the device with the matching IP address replies to the device that sends the data with a packet that contains the MAC address for the device. The source device adds the destination device MAC address to its ARP table for future reference, creates a data-link header and trailer that encapsulates the packet, and proceeds to transfer the data.

Figure 5: ARP Process



When the destination device lies on a remote network that is beyond another device, the process is the same except that the device that sends the data sends an ARP request for the MAC address of the default gateway. After the address is resolved and the default gateway receives the packet, the default gateway broadcasts the destination IP address over the networks connected to it. The device on the destination device network uses ARP to obtain the MAC address of the destination device and delivers the packet. ARP is enabled by default.

The default system-defined CoPP policy rate limits ARP broadcast packets bound for the supervisor module. The default system-defined CoPP policy prevents an ARP broadcast storm from affecting the control plane traffic but does not affect bridged packets.

ARP Caching

ARP caching minimizes broadcasts and limits wasteful use of network resources. The mapping of IP addresses to MAC addresses occurs at each hop (device) on the network for every packet sent over an internetwork, which may affect network performance.

ARP caching stores network addresses and the associated data-link addresses in the memory for a period of time, which minimizes the use of valuable network resources to broadcast for the same address each time that a packet is sent. You must maintain the cache entries that are set to expire periodically because the information might become outdated. Every device on a network updates its tables as addresses are broadcast.

To maintain the ARP entry, active MAC address-table entries and host routing adjacencies, Cisco NX-OS sends up to 3 unicast ARP request messages to devices that are present in the ARP cache. The first message is sent at 75% of the configured ARP timeout value, followed by two retries 30 and 60 seconds later if the cached entry has not already been refreshed.

Static and Dynamic Entries in the ARP Cache

Static routing requires that you manually configure the IP addresses, subnet masks, gateways, and corresponding MAC addresses for each interface of each device. Static routing requires more work to maintain the route table. You must update the table each time you add or change routes.

Dynamic routing uses protocols that enable the devices in a network to exchange routing table information with each other. Dynamic routing is more efficient than static routing because the route table is automatically updated unless you add a time limit to the cache. The default time limit is 25 minutes but you can modify the time limit if the network has many routes that are added and deleted from the cache.

Devices That Do Not Use ARP

When a network is divided into two segments, a bridge joins the segments and filters traffic to each segment based on MAC addresses. The bridge builds its own address table, which uses MAC addresses only. A device has an ARP cache that contains both IP addresses and the corresponding MAC addresses.

Passive hubs are central-connection devices that physically connect other devices in a network. They send messages out on all their ports to the devices and operate at Layer 1 but do not maintain an address table.


Layer 2 switches determine which port is connected to a device to which the message is addressed and sent only to that port. However, Layer 3 switches are devices that build an ARP cache (table).

Reverse ARP

Reverse ARP (RARP) as defined by RFC 903 works the same way as ARP, except that the RARP request packet requests an IP address instead of a MAC address. RARP often is used by diskless workstations because this type of device has no way to store IP addresses to use when they boot. The only address that is known is the MAC address because it is burned into the hardware.

Use of RARP requires an RARP server on the same network segment as the router interface.

Figure 6: Reverse ARP



RARP has several limitations. Because of these limitations, most businesses use DHCP to assign IP addresses dynamically. DHCP is cost effective and requires less maintenance than RARP. The following are the most important limitations:

- Since RARP uses hardware addresses, if the internetwork is large with many physical networks, a RARP server must be on every segment with an additional server for redundancy. Maintaining two servers for every segment is costly.
- Each server must be configured with a table of static mappings between the hardware addresses and IP addresses. Maintenance of the IP addresses is difficult.
- RARP only provides IP addresses of the hosts and not subnet masks or default gateways.

Proxy ARP

Proxy ARP enables a device that is physically located on one network appear to be logically part of a different physical network connected to the same device or firewall. Proxy ARP allows you to hide a device with a public IP address on a private network behind a router and still have the device appear to be on the public network in front of the router. By hiding its identity, the router accepts responsibility for routing packets to the real destination. Proxy ARP can help devices on a subnet reach remote subnets without configuring routing or a default gateway.

When devices are not in the same data link layer network but in the same IP network, they try to transmit data to each other as if they are on the local network. However, the router that separates the devices does not send a broadcast message because routers do not pass hardware-layer broadcasts and the addresses cannot be resolved.

When you enable Proxy ARP on the device and it receives an ARP request, it identifies the request as a request for a system that is not on the local LAN. The device responds as if it is the remote destination for which the broadcast is addressed, with an ARP response that associates the device's MAC address with the remote destination's IP address. The local device believes that it is directly connected to the destination, while in reality its packets are being forwarded from the local subnetwork toward the destination subnetwork by their local device. By default, Proxy ARP is disabled.

Local Proxy ARP

You can use local Proxy ARP to enable a device to respond to ARP requests for IP addresses within a subnet where normally no routing is required. When you enable local Proxy ARP, ARP responds to all ARP requests for IP addresses within the subnet and forwards all traffic between hosts in the subnet. Use this feature only on subnets where hosts are intentionally prevented from communicating directly by the configuration on the device to which they are connected.

Gratuitous ARP

Gratuitous ARP sends a request with an identical source IP address and a destination IP address to detect duplicate IP addresses. Cisco NX-OS supports enabling or disabling gratuitous ARP requests or ARP cache updates.

ARP Refresh on MAC Delete

On MAC delete, by default the ARP entry corresponding to the deleted MAC is refreshed.” Below mentioned are the ARP refresh behavior:

1. The physical interface corresponding to the adjacency is updated as NULL to FIB.
2. ARP entry in moved to INCOMPLETE state and ARP refresh is sent to the host.
3. ARP entry moved to RESOLVED state, after receiving ARP response.
4. If an ARP response is not received, the ARP process retries for ARP refresh at intervals of 2, 4, 8, and 16.
 - a. If the ARP response is not received at these intervals, the ARP entry is deleted after 30 seconds which is the summation of these interval times.

Glean Throttling

When forwarding an incoming IP packet in a line card, if the Address Resolution Protocol (ARP) request for the next hop is not resolved, the line card forwards the packets to the supervisor (glean throttling). The supervisor resolves the MAC address for the next hop and programs the hardware.

The Cisco Nexus 7000 Series device hardware has glean rate limiters to protect the supervisor from the glean traffic. If the maximum number of entries is exceeded, the packets for which the ARP request is not resolved continues to be processed in the software instead of getting dropped in the hardware.

When an ARP request is sent, the software adds a /32 drop adjacency in the hardware to prevent the packets to the same next-hop IP address to be forwarded to the supervisor. When the ARP is resolved, the hardware

entry is updated with the correct MAC address. If the ARP entry is not resolved before a timeout period, the entry is removed from the hardware.

Path MTU Discovery

Path maximum transmission unit (MTU) discovery is a method for maximizing the use of available bandwidth in the network between the endpoints of a TCP connection. It is described in RFC 1191. Existing connections are not affected when this feature is turned on or off.



Note Please ensure you enable **ip unreachable** command between TCP endpoints for the Path MTU discovery feature to work correctly.

ICMP

You can use the Internet Control Message Protocol (ICMP) to provide message packets that report errors and other information that is relevant to IP processing. ICMP generates error messages, such as ICMP destination unreachable messages, ICMP Echo Requests (which send a packet on a round trip between two hosts) and Echo Reply messages. ICMP also provides many diagnostic functions and can send and redirect error packets to the host. By default, ICMP is enabled.

Some of the ICMP message types are as follows:

- Network error messages
- Network congestion messages
- Troubleshooting information
- Timeout announcements



Note ICMP redirects are disabled on interfaces where the local proxy ARP feature is enabled.

Virtualization Support for IPv4

IPv4 supports virtual routing and forwarding (VRF) instances. VRFs exist within virtual device contexts (VDCs). By default, Cisco NX-OS places you in the default VDC and default VRF unless you specifically configure another VDC and VRF. For more information, see the *Cisco NX-OS Virtual Device Context Configuration Guide*.

Prerequisites for IPv4

IPv4 has the following prerequisites:

- IPv4 can only be configured on Layer 3 interfaces.

Guidelines and Limitations for IPv4

IPv4 has the following configuration guidelines and limitations:

- You can configure a secondary IP address only after you configure the primary IP address.
- F2 Series modules do not support IPv4 tunnels.
- If any device on a network segment uses a secondary IPv4 address, other devices on that same network segment that require a secondary address must use a secondary address from the same network or subnet. The inconsistent use of secondary addresses on a network segment can quickly cause routing loops.
- If you are familiar with the Cisco IOS CLI, be aware that the Cisco NX-OS commands for this feature might differ from the Cisco IOS commands that you would use.

Default Settings for IPv4 Parameters

Table 3: Default IPv4 Parameters

Parameters	Default
ARP timeout	1500 seconds
proxy ARP	Disabled
Maximum number of IPv4 ARP entries in the neighbor adjacency table	131,072

Configuring IPv4

Configuring IPv4 Addressing

You can assign a primary IP address for a network interface.

Before you begin

Ensure that you are in the correct VDC (or use the **switchto vdc** command).

Procedure

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# interface ethernet <i>number</i>	Enters interface configuration mode.
Step 3	switch(config-if)# no switchport	Configures the interface as a Layer 3 routed interface.

	Command or Action	Purpose
Step 4	switch(config-if)# ip address <i>ip-address/length</i> [<i>secondary</i>]	Specifies a primary or secondary IPv4 address for an interface. <ul style="list-style-type: none"> • The network mask can be a four-part dotted decimal address. For example, 255.0.0.0 indicates that each bit equal to 1 means the corresponding address bit belongs to the network address. • The network mask can be indicated as a slash (/) and a number—a prefix length. The prefix length is a decimal value that indicates how many of the high-order contiguous bits of the address comprise the prefix (the network portion of the address). A slash must precede the decimal value and there is no space between the IP address and the slash.
Step 5	(Optional) switch(config-if)# show ip interface	Displays interfaces configured for IPv4.
Step 6	(Optional) switch(config)# copy running-config startup-config	Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

Example

This example shows how to assign an IPv4 address:

```
switch# configure terminal
switch(config)# interface ethernet 2/3
switch(config-if)# no switchport
switch(config-if)# ip address 192.2.1.1.255.0.0.0
switch(config-if)# copy running-config startup-config
switch(config-if)#
```

Configuring Multiple IPv4 Addresses

You can only add secondary IP addresses after you configure primary IP addresses.

Before you begin

Ensure that you are in the correct VDC (or use the **switchto vdc** command).

Procedure

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.

	Command or Action	Purpose
Step 2	switch(config)# interface ethernet <i>number</i>	Enters interface configuration mode.
Step 3	switch(config-if)# no switchport	Configures the interface as a Layer 3 routed interface.
Step 4	switch(config-if)# ip address <i>ip-address/length</i> [<i>secondary</i>]	Specifies a the configured address as a secondary IPv4 address.
Step 5	(Optional) switch(config-if)# show ip interface	Displays interfaces configured for IPv4.
Step 6	(Optional) switch(config)# copy running-config startup-config	Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

Configuring a Static ARP Entry

Configure a static ARP entry on the device to map IP addresses to MAC hardware addresses, including static multicast MAC addresses.

Before you begin

Ensure that you are in the correct VDC (or use the **switchto vdc** command).

Procedure

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# interface ethernet <i>number</i>	Enters interface configuration mode.
Step 3	switch(config)# no switchport	Configures the interface as a Layer 3 routed interface.
Step 4	switch(config-if)# ip arp address <i>ip-address</i> <i>mac-address</i>	Associates an IP address with a MAC address as a static entry.
Step 5	(Optional) switch(config)# copy running-config startup-config	Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

Example

This example shows how to assign a static ARP entry:

```
switch# configure terminal
switch(config)# interface ethernet 2/3
switch(config-if)# no switchport
switch(config-if)# ip arp 192.2.1.1.0019.076c.1a78
switch(config-if)# copy running-config startup-config
switch(config-if)#
```

Configuring Proxy ARP

Configure proxy ARP on the device to determine the media addresses of hosts on other networks or subnets.

Before you begin

Ensure that you are in the correct VDC (or use the **switchto vdc** command).

Procedure

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config-if)# interface ethernet number	Enters interface configuration mode.
Step 3	switch(config-if)# no switchport	Configures the interface as a Layer 3 routed interface.
Step 4	switch(config-if)# ip proxy arp	Enables proxy ARP on the interface.
Step 5	switch(config)# copy running-config startup-config	(Optional) Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

Example

This example shows how to configure proxy ARP:

```
switch# configure terminal
switch(config)# interface ethernet 2/3
switch(config-if)# no switchport
switch(config-if)# ip proxy-arp
switch(config-if)# copy running-config startup-config
switch(config-if)#
```

Configuring Local Proxy ARP

Before you begin

Ensure that you are in the correct VDC (or use the **switchto vdc** command).

Procedure

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# interface ethernet number	Enters interface configuration mode.

	Command or Action	Purpose
Step 3	switch(config)# no switchport	Configures the interface as a Layer 3 routed interface.
Step 4	switch(config-if)# ip local-proxy-arp	Enables local proxy ARP on the interface.
Step 5	(Optional) switch(config)# copy running-config startup-config	Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

Example

This example shows how to configure local proxy ARP:

```
switch# configure terminal
switch(config)# interface ethernet 2/3
switch(config-if)# no switchport
switch(config-if)# ip local-proxy-arp
switch(config-if)# copy running-config startup-config
switch(config-if)#
```

Configuring Gratuitous ARP

Before you begin

Ensure that you are in the correct VDC (or use the **switchto vdc** command).

Procedure

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# interface ethernet number	Enters interface configuration mode.
Step 3	switch(config-if)# no switchport	Configures the interface as a Layer 3 routed interface.
Step 4	switch(config-if)# ip arp gratuitous {request update}	Enables gratuitous ARP on the interface. Gratuitous ARP is enabled by default.
Step 5	(Optional) switch(config)# copy running-config startup-config	Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

Example

This example shows how to configure gratuitous ARP:

```

switch# configure terminal
switch(config)# interface ethernet 2/3
switch(config-if)# no switchport
switch(config-if)# ip arp gratuitous request
switch(config-if)# copy running-config startup-config
switch(config-if)#

```

Configuring the IP ARP Cache Limit

Before you begin

Ensure that you are in the correct VDC (or use the **switchto vdc** command).

Procedure

	Command or Action	Purpose
Step 1	switch(config)# configure terminal	Enters global configuration mode.
Step 2	switch(config)# ip arp cache limit <i>max-arp-entries</i> [syslog <i>syslogs-per-second</i>]	Configures the maximum number of ARP entries in the neighbor adjacency table. The range is from 1 to 409600. The syslog keyword configures the number of syslogs per second. The range is from 1 to 1000. If you do not configure a limit, system logs appear on the console if you try to add an adjacency after reaching the default limit. If you configure a limit for IPv4 ARP entries, system logs appear if you try to add an adjacency after reaching the configured limit.
Step 3	switch(config)# show ip adjacency summary	Displays the global limit of the neighbor adjacency table and a summary of throttle adjacencies.
Step 4	(Optional) switch(config)# copy running-config startup-config	Saves this configuration change.

Configuring Glean Optimization

You can configure glean optimization to improve the performance of glean packets by reducing the processing of the packets in the supervisor. Glean optimization applies to glean packets where the destination IP address is part of the same subnet and does not apply to packets where the destination IP address is in a different subnet. The default is enabled.

Before you begin

Ensure that you are in the correct VDC (or use the **switchto vdc** command).

Procedure

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# interface ethernet <i>number</i>	Enters interface configuration mode.
Step 3	switch(config-if)# [no] ip arp fast-path	Enables glean optimization. Use the no form of the command to disable this feature.
Step 4	(Optional) switch(config)# copy running-config startup-config	Saves this configuration change.

Configuring Bloom Filter Support for Glean Adjacencies

Bloom Filter Support for Glean Adjacencies is introduced in Cisco NX-OS Release 8.4(2).

When a routed frame has an ARP cache miss, the packet hits a glean adjacency (which means the IP DA hits on the FIB table but cannot resolve MAC DA for the routed frame), and it is punted to the supervisor module. Until the ARP cache is updated, all packets belonging to this flow will hit the glean adjacency and are punted to the supervisor module. To avoid this punting of the supervisor module, the L3 engine hashes a flow to set a bit in a leak table to indicate that the packet has been punted to the supervisor module. Subsequent frames are dropped until the software clears the leak table bit. This helps to forward the packets without any further delay.

The Bloom Filter Support for Glean Adjacencies feature is supported on M3 and F4 modules.

Before you perform the configuration, ensure that you are in the correct VDC or use the **switchto vdc** command. This command is a global, system CLI on the default vdc and it is not configurable on a non-default vdc.

Procedure

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# hardware forwarding glean-bloom-filter	Enables the bloom filter forwarding. This command is disabled by default.
Step 3	switch(config)# no hardware forwarding glean-bloom-filter	Disables the bloom filter forwarding.
Step 4	(Optional) switch(config)# copy running-config startup-config	Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.
Step 5	switch(config)# exit	Exits the global configuration mode.
Step 6	(Optional) switch# show system internal forwarding route summary	Displays the glean routes from all supported modules.

Example

This example shows how to enable IP glean throttling:

```

switch# configure terminal
switch(config)# hardware forwarding glean-bloom-filter
switch(config)# copy running-config startup-config
switch(config)# exit
switch# show system internal forwarding route summary

```

Configuring Path MTU Discovery

Before you begin

Ensure that you are in the correct VDC (or use the **switchto vdc** command).

Procedure

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# ip tcp path-mtu-discovery	Enables path MTU discovery.
Step 3	(Optional) switch(config)# copy running-config startup-config	Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

Configuring IP Packet Verification

Cisco NX-OS supports an Intrusion Detection System (IDS) that checks for IP packet verification. You can enable or disable these IDS checks.

Procedure

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# hardware ip verify address { destination zero identical reserved source { broadcast multicast }}	Performs the following IDS checks on the IP address: <ul style="list-style-type: none"> • destination zero—Drops IP packets if the destination IP address is 0.0.0.0. • identical—Drops IP packets if the source IP address is identical to the destination IP address. • reserved—Drops IP packets if the IP address is in the 127.x.x.x range.

	Command or Action	Purpose
		<ul style="list-style-type: none"> source—Drops IP packets if the IP source address is either 255.255.255.255 (broadcast) or in the 224.x.x.x range (multicast).
Step 3	switch(config)# hardware ip verify checksum	Drops IP packets if the packet checksum is invalid.
Step 4	switch(config)# hardware ip verify fragment	Drops IP packets if the packet fragment has a nonzero offset and the DF bit is active.
Step 5	switch(config)# hardware ip verify length { consistent maximum { max-frag max-tcp udp } minimum }	<p>Performs the following IDS checks on the IP address:</p> <ul style="list-style-type: none"> consistent— Drops IP packets where the Ethernet frame size is greater than or equal to the IP packet length plus the Ethernet header. maximum max-frag—Drops IP packets if the maximum fragment offset is greater than 65536. maximum max-tcp—Drops IP packets if the TCP length is greater than the IP payload length. maximum udp—Drops IP packets if the IP payload length is less than the UDP packet length. minimum—Drops IP packets if the Ethernet frame length is less than the IP packet length plus four octets (the CRC length).
Step 6	switch(config)# hardware ip verify tcp tiny-frag	Drops TCP packets if the IP fragment offset is 1, or if the IP fragment offset is 0 and the IP payload length is less than 16.
Step 7	switch(config)# hardware ip verify version	Drops IP packets if the ethertype is not set to 4 (IPv4).

What to do next

Use the **show hardware forwarding ip verify** command to display the IP packet verification configuration.

Configuring IP Glean Throttling

Cisco NX-OS software supports glean throttling rate limiters to protect the supervisor from the glean traffic.



Note We recommend that you configure the IP glean throttle feature by using the **hardware ip glean throttle** command to filter the unnecessary glean packets that are sent to the supervisor for ARP resolution for the next hops that are not reachable or do not exist. IP glean throttling boosts software performance and helps to manage traffic more efficiently.

Before you begin

Ensure that you are in the correct VDC (or use the **switchto vdc** command).

Procedure

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# hardware ip glean throttle	Enables ARP throttling.
Step 3	switch(config)# no hardware ip glean throttle	Disables ARP throttling.
Step 4	(Optional) switch(config)# copy running-config startup-config	Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

Example

This example shows how to enable IP glean throttling:

```
switch# configure terminal
switch(config)# hardware ip glean throttle
switch(config-if)# copy running-config startup-config
```

Configuring the Hardware IP Glean Throttle Maximum

You can limit the maximum number of drop adjacencies that are installed in the Forwarding Information Base (FIB).

Before you begin

Ensure that you are in the correct VDC (or use the **switchto vdc** command).

Procedure

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# hardware ip glean throttle maximum count	Configures the number of drop adjacencies that are installed in the FIB.

	Command or Action	Purpose
Step 3	switch(config)# no hardware ip glean throttle maximum count	Applies the default limits. The default value is 1000. The range is from 0 to 32767 entries.
Step 4	(Optional) switch(config)# copy running-config startup-config	Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

Example

This example shows how to limit the maximum number of drop adjacencies that are installed in the FIB:

```
switch# configure terminal
switch(config)# hardware ip glean throttle maximum 2134
switch(config-if)# copy running-config startup-config
```

Configuring the Hardware IP Glean Throttle Timeout

You can configure a timeout for the installed drop adjacencies to remain in the Forwarding Information Base (FIB).

Before you begin

Ensure that you are in the correct VDC (or use the **switchto vdc** command).

Procedure

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# hardware ip glean throttle maximum timeout timeout-in-seconds	Configures the timeout for the installed drop adjacencies to remain in the FIB.
Step 3	switch(config)# no hardware ip glean throttle maximum timeout timeout-in-seconds	Applies the default limits. The timeout value is in seconds. The range is from 300 seconds (5 minutes) to 1800 seconds (30 minutes). Note After the timeout period is exceeded, the drop adjacencies are removed from the FIB.
Step 4	(Optional) switch(config)# copy running-config startup-config	Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

Configuring the Hardware IP Glean Throttle Syslog

You can a syslog if the number of packets that get dropped for a specific flow exceeds the configured packet count.

Before you begin

Ensure that you are in the correct VDC (or use the **switchto vdc** command).

Procedure

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# hardware ip glean throttle syslog packet-count	Generates a syslog if the number of packets that get dropped for a specific flow exceed the configured packet count.
Step 3	switch(config)# no hardware ip glean throttle syslog packet-count	Applies the default limits. The default is 10000 packets. The range is from 0 to 65535 packets. Note After the timeout period is exceeded, the drop adjacencies are removed from the FIB.
Step 4	(Optional) switch(config)# copy running-config startup-config	Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

Example

This example shows how to generate a syslog if the number of packets that get dropped for a specific flow exceeds the configured packet count:

```
switch# configure terminal
switch(config)# hardware ip glean throttle maximum timeout 300
switch(config-if)# copy running-config startup-config
```

Verifying the IPv4 Configuration

Use one of the following commands to verify the configuration:

Command	Purpose
show hardware forwarding ip verify	Shows the IP packet verification configuration.
show ip adjacency	Displays the adjacency table.
show ip adjacency summary	Displays the summary of number of throttle adjacencies.

Command	Purpose
show ip arp	Displays the ARP table.
show ip arp summary	Displays the summary of the number of throttle adjacencies.
show ip adjacency throttle statistics	Displays only the throttle adjacencies.
show ip interface	Displays IP-related interface information.
show ip arp statistics [vrf vrf-name]	Displays the ARP statistics.

Configuration Examples for IPv4

Example: Reserving All Ports on a Module for Proxy Routing

This example shows how to reserve all ports on a module for proxy routing:

Step 1: Determine which modules are present in the device:

```
switch# show module
Mod Ports Module-Type Model Status
-----
1 32 10 Gbps Ethernet Module N7K-M132XP-12 ok
2 48 10/100/1000 Mbps Ethernet Module N7K-M148GT-11 ok
3 48 1000 Mbps Optical Ethernet Modul N7K-M148GS-11 ok
5 0 Supervisor module-1X N7K-SUP1 active *
6 0 Supervisor module-1X N7K-SUP1 ha-standby
8 32 1/10 Gbps Ethernet Module N7K-F132XP-15 ok
```

The F1 module is in Slot 8, and the M1 modules are in Slots 1 to 3.

Step 2: Determine which ports are available in the VDC:

```
switch# show vdc membership | end "Ethernet3/48"
vdc_id: 0 vdc_name: Unallocated interfaces:
vdc_id: 1 vdc_name: switch interfaces:
Ethernet1/9 Ethernet1/10 Ethernet1/11
Ethernet1/12 Ethernet1/13 Ethernet1/14
Ethernet1/15 Ethernet1/16 Ethernet1/17
Ethernet1/18 Ethernet1/19 Ethernet1/20
Ethernet1/21 Ethernet1/22 Ethernet1/23
Ethernet1/24 Ethernet1/25 Ethernet1/26
Ethernet1/27 Ethernet1/28 Ethernet1/29
Ethernet1/30 Ethernet1/31 Ethernet1/32
Ethernet2/1 Ethernet2/2 Ethernet2/3
Ethernet2/4 Ethernet2/5 Ethernet2/6
Ethernet2/7 Ethernet2/8 Ethernet2/9
Ethernet2/10 Ethernet2/11 Ethernet2/12
Ethernet2/25 Ethernet2/26 Ethernet2/27
Ethernet2/28 Ethernet2/29 Ethernet2/30
Ethernet2/31 Ethernet2/32 Ethernet2/33
Ethernet2/34 Ethernet2/35 Ethernet2/36
Ethernet2/37 Ethernet2/38 Ethernet2/39
Ethernet2/40 Ethernet2/41 Ethernet2/42
Ethernet2/43 Ethernet2/44 Ethernet2/45
Ethernet2/46 Ethernet2/47 Ethernet2/48
```

```

Ethernet3/1 Ethernet3/2 Ethernet3/3
Ethernet3/4 Ethernet3/5 Ethernet3/6
Ethernet3/7 Ethernet3/8 Ethernet3/9
Ethernet3/10 Ethernet3/11 Ethernet3/12
Ethernet3/13 Ethernet3/14 Ethernet3/15
Ethernet3/16 Ethernet3/17 Ethernet3/18
Ethernet3/19 Ethernet3/20 Ethernet3/21
Ethernet3/22 Ethernet3/23 Ethernet3/24
Ethernet3/25 Ethernet3/26 Ethernet3/27
Ethernet3/28 Ethernet3/29 Ethernet3/30
Ethernet3/31 Ethernet3/32 Ethernet3/33
Ethernet3/34 Ethernet3/35 Ethernet3/36
Ethernet3/37 Ethernet3/38 Ethernet3/39
Ethernet3/40 Ethernet3/41 Ethernet3/42
Ethernet3/43 Ethernet3/44 Ethernet3/45
Ethernet3/46 Ethernet3/47 Ethernet3/48

```

Step 3: Determine which ports are available for proxy routing:

```

switch# show hardware proxy layer-3 detail
Global Information:
F1 Modules: Count: 1 Slot: 8
M1 Modules: Count: 3 Slot: 1-3
Replication Rebalance Mode: Manual
Number of proxy layer-3 forwarders: 13
Number of proxy layer-3 replicators: 8
Forwarder Interfaces Status Reason
-----
Eth1/9, Eth1/11, Eth1/13, Eth1/15 up SUCCESS
Eth1/10, Eth1/12, Eth1/14, Eth1/16 up SUCCESS
Eth1/17, Eth1/19, Eth1/21, Eth1/23 up SUCCESS
Eth1/18, Eth1/20, Eth1/22, Eth1/24 up SUCCESS
Eth1/25, Eth1/27, Eth1/29, Eth1/31 up SUCCESS
Eth1/26, Eth1/28, Eth1/30, Eth1/32 up SUCCESS
Eth2/1-12 up SUCCESS
Eth2/25-36 up SUCCESS
Eth2/37-48 up SUCCESS
Eth3/1-12 up SUCCESS
Eth3/13-24 up SUCCESS
Eth3/25-36 up SUCCESS
Eth3/37-48 up SUCCESS
Replicator Interfaces #Interface-Vlan Interface-Vlan
-----
Eth1/1, Eth1/3, Eth1/5, Eth1/7, Eth1/9, 0
Eth1/11, Eth1/13, Eth1/15
Eth1/2, Eth1/4, Eth1/6, Eth1/8, Eth1/10, 0
Eth1/12, Eth1/14, Eth1/16
Eth1/17, Eth1/19, Eth1/21, Eth1/23, 0
Eth1/25, Eth1/27, Eth1/29, Eth1/31
Eth1/18, Eth1/20, Eth1/22, Eth1/24, 0
Eth1/26, Eth1/28, Eth1/30, Eth1/32
Eth2/1-24 0
Eth2/25-48 0
Eth3/1-24 0
Eth3/25-48 0
switch#

```



Note Ports are listed in their respective port groups.

Step 4: Reserve a module for unicast and multicast proxy routing:

```
switch# configure terminal
switch(config)# hardware proxy layer-3 forwarding use module 2
switch(config)# hardware proxy layer-3 replication use module 2
```

Step 5: Verify this configuration:

```
switch(config)# show hardware proxy layer-3 detail
Global Information:
F1 Modules: Count: 1 Slot: 8
M1 Modules: Count: 3 Slot: 1-3
Replication Rebalance Mode: Manual
Number of proxy layer-3 forwarders: 3
Number of proxy layer-3 replicators: 2
Forwarder Interfaces Status Reason
-----
Eth2/1-12 up SUCCESS
Eth2/25-36 up SUCCESS
Eth2/37-48 up SUCCESS

Replicator Interfaces #Interface-Vlan Interface-Vlan
-----
Eth2/1-24 0
Eth2/25-48 0
switch(config)#
```

Example: Reserving Ports for Proxy Routing

This example shows how to reserve some ports on a module for proxy routing:

Step 1: Reserve a subset of ports on a module:

```
switch(config)# hardware proxy layer-3 forwarding use interface ethernet 2/1-6
switch(config)# hardware proxy layer-3 replication use interface ethernet 2/1-6 <----subset
of port group
```

This example reserves a subset of ports from a port group.

Step 2: Verify this configuration:

```
switch(config)# show hardware proxy layer-3 detail
Global Information:
F1 Modules: Count: 1 Slot: 8
M1 Modules: Count: 3 Slot: 1-3
Replication Rebalance Mode: Manual
Number of proxy layer-3 forwarders: 1
Number of proxy layer-3 replicators: 1
Forwarder Interfaces Status Reason
-----
Eth2/1-12 up SUCCESS
Replicator Interfaces #Interface-Vlan Interface-Vlan
-----
Eth2/1-24 0
switch(config)#
```



Note All ports in a port group are reserved for proxy routing.

Example: Excluding Ports From Proxy Routing

The following example excludes some ports on a module for proxy routing:

```

switch(config)# hardware proxy layer-3 forwarding exclude interface ethernet 2/1-12
switch(config)# hardware proxy layer-3 replication exclude interface ethernet 2/1-12
switch(config)# show hardware proxy layer-3 detail
Global Information:
F1 Modules: Count: 1 Slot: 8
M1 Modules: Count: 3 Slot: 1-3
Replication Rebalance Mode: Manual
Number of proxy layer-3 forwarders: 12
Number of proxy layer-3 replicators: 7
Forwarder Interfaces Status Reason
-----
Eth1/9, Eth1/11, Eth1/13, Eth1/15 up SUCCESS
Eth1/10, Eth1/12, Eth1/14, Eth1/16 up SUCCESS
Eth1/17, Eth1/19, Eth1/21, Eth1/23 up SUCCESS
Eth1/18, Eth1/20, Eth1/22, Eth1/24 up SUCCESS
Eth1/25, Eth1/27, Eth1/29, Eth1/31 up SUCCESS
Eth1/26, Eth1/28, Eth1/30, Eth1/32 up SUCCESS
Eth2/25-36 up SUCCESS
Eth2/37-48 up SUCCESS
Eth3/1-12 up SUCCESS
Eth3/13-24 up SUCCESS
Eth3/25-36 up SUCCESS
Eth3/37-48 up SUCCESS

Replicator Interfaces #Interface-Vlan Interface-Vlan
-----
Eth1/1, Eth1/3, Eth1/5, Eth1/7, Eth1/9, 0
Eth1/11, Eth1/13, Eth1/15
Eth1/2, Eth1/4, Eth1/6, Eth1/8, Eth1/10, 0
Eth1/12, Eth1/14, Eth1/16
Eth1/17, Eth1/19, Eth1/21, Eth1/23, 0
Eth1/25, Eth1/27, Eth1/29, Eth1/31
Eth1/18, Eth1/20, Eth1/22, Eth1/24, 0
Eth1/26, Eth1/28, Eth1/30, Eth1/32
Eth2/25-48 0
Eth3/1-24 0
Eth3/25-48 0
switch(config)#

```

Related Documents for IPv4

Related Topic	Document Title
IP CLI commands	https://www.cisco.com/c/en/us/td/docs/switches/datacenter/nw/sw/unicast/command/cisco_nexus7000_unicast_routing_command_ref.html

Standards for IPv4

Standards	Title
No new or modified standards are supported by this feature, and support for existing standards has not been modified by this feature.	—

Feature History for IPv4

This table includes only the updates for those releases that have resulted in additions or changes to the feature.

Table 4: Feature History for IPv4

Feature Name	Release	Feature Information
Glean optimization	6.2(2)	This feature was introduced.
ARP	6.2(2)	Added the ability to configure the maximum number of ARP entries in the neighbor adjacency table.
IP	6.0(1)	Updated for F2 Series modules.
ACL filter for IP directed broadcasts	5.2(1)	Added support to filter IP directed broadcasts through an IP access list.
Glean throttling	5.1(1)	Added support for IPv4 glean throttling.
ARP	4.1(4)	Added support to protect against an ARP broadcast storm.
IP	4.1(3)	Changed the platform ip verify command to the hardware ip verify command.
ARP	4.0(3)	Added support for gratuitous ARP. The ip arp gratuitous {request update} command was added.
IP	4.0(1)	This feature was introduced.



CHAPTER 4

Configuring IPv6

This chapter contains the following sections:

- [Finding Feature Information, on page 45](#)
- [Information About IPv6, on page 45](#)
- [Virtualization Support for IPv6, on page 62](#)
- [Prerequisites for IPv6, on page 62](#)
- [Guidelines and Limitations for Configuring IPv6, on page 62](#)
- [Default Settings for IPv6, on page 63](#)
- [Configuring IPv6, on page 63](#)
- [Verifying the IPv6 Configuration, on page 72](#)
- [Configuration Example for IPv6, on page 72](#)
- [Related Documents for IPv6, on page 72](#)
- [Standards for IPv6, on page 72](#)
- [Feature History for IPv6, on page 73](#)

Finding Feature Information

Your software release might not support all the features documented in this module. For the latest caveats and feature information, see the Bug Search Tool at <https://tools.cisco.com/bugsearch/> and the release notes for your software release. To find information about the features documented in this module, and to see a list of the releases in which each feature is supported, see the "New and Changed Information" chapter or the Feature History table in this chapter.

Information About IPv6

IPv6, which is designed to replace IPv4, increases the number of network address bits from 32 bits (in IPv4) to 128 bits. IPv6 is based on IPv4 but it includes a much larger address space and other improvements such as a simplified main header and extension headers.

The larger IPv6 address space allows networks to scale and provide global reachability. The simplified IPv6 packet header format handles packets more efficiently. The flexibility of the IPv6 address space reduces the need for private addresses and the use of Network Address Translation (NAT), which translates private (not globally unique) addresses into a limited number of public addresses. IPv6 enables new application protocols that do not require special processing by border routers at the edge of networks.

IPv6 functionality, such as prefix aggregation, simplified network renumbering, and IPv6 site multihoming capabilities, enable more efficient routing. IPv6 supports Routing Information Protocol (RIP), Integrated Intermediate System-to-Intermediate System (IS-IS), Open Shortest Path First (OSPF) for IPv6, and multiprotocol Border Gateway Protocol (BGP).

IPv6 Address Formats

An IPv6 address has 128 bits or 16 bytes. The address is divided into eight, 16-bit hexadecimal blocks separated by colons (:) in the format x:x:x:x:x:x:x.

Two examples of IPv6 addresses are as follows:

```
2001:0DB8:7654:3210:FEDC:BA98:7654:3210
```

```
2001:0DB8:0:0:8:800:200C:417A
```

IPv6 addresses contain consecutive zeros within the address. You can use two colons (::) at the beginning, middle, or end of an IPv6 address to replace the consecutive zeros.



Note You can use two colons (::) only once in an IPv6 address to replace the longest string of consecutive zeros within the address.

You can use a double colon as part of the IPv6 address when consecutive 16-bit values are denoted as zero. You can configure multiple IPv6 addresses per interface but only one link-local address.

The hexadecimal letters in IPv6 addresses are not case sensitive.

Table 5: Compressed IPv6 Address Formats

IPv6 Address Type	Preferred Format	Compressed Format
Unicast	2001:0:0:0:DB8:800:200C:417A	2001::0DB8:800:200C:417A
Multicast	FF01:0:0:0:0:0:0:101	FF01::101
Loopback	0:0:0:0:0:0:0:1	::1
Unspecified	0:0:0:0:0:0:0:0	::

A node may use the loopback address listed in the table to send an IPv6 packet to itself. The loopback address in IPv6 is the same as the loopback address in IPv4.



Note You cannot assign the IPv6 loopback address to a physical interface. A packet that contains the IPv6 loopback address as its source or destination address must remain within the node that created the packet. IPv6 routers do not forward packets that have the IPv6 loopback address as their source or destination address.

You cannot assign an IPv6 unspecified address to an interface. You should not use the unspecified IPv6 addresses as destination addresses in IPv6 packets or the IPv6 routing header.

The IPv6-prefix is in the form documented in RFC 2373 where the IPv6 address is specified in hexadecimal using 16-bit values between colons. The prefix length is a decimal value that indicates how many of the

high-order contiguous bits of the address comprise the prefix (the network portion of the address). For example, 2001:0DB8:8086:6502::/32 is a valid IPv6 prefix.

IPv6 Unicast Addresses


An IPv6 unicast address is an identifier for a single interface on a single node. A packet that is sent to a unicast address is delivered to the interface identified by that address.

Aggregatable Global Addresses

An aggregatable global address is an IPv6 address from the aggregatable global unicast prefix. The structure of aggregatable global unicast addresses enables strict aggregation of routing prefixes that limits the number of routing table entries in the global routing table. Aggregatable global addresses are used on links that are aggregated upward through organizations and eventually to the Internet service providers (ISPs).

Aggregatable global IPv6 addresses are defined by a global routing prefix, a subnet ID, and an interface ID. Except for addresses that start with binary 000, all global unicast addresses have a 64-bit interface ID. The IPv6 global unicast address allocation uses the range of addresses that start with binary value 001 (2000::/3). The figure shows the structure of an aggregatable global address.

Figure 7: Aggregatable Global Addresses



Addresses with a prefix of 2000::/3 (001) through E000::/3 (111) are required to have 64-bit interface identifiers in the extended universal identifier (EUI)-64 format. The Internet Assigned Numbers Authority (IANA) allocates the IPv6 address space in the range of 2000::/16 to regional registries.

The aggregatable global address consists of a 48-bit global routing prefix and a 16-bit subnet ID or Site-Level Aggregator (SLA). In the IPv6 aggregatable global unicast address format document (RFC 2374), the global routing prefix included two other hierarchically structured fields called Top-Level Aggregator (TLA) and Next-Level Aggregator (NLA). The IETF decided to remove the TLA and NLA fields from the RFCs because these fields are policy based. Some existing IPv6 networks deployed before the change might still use networks that are on the older architecture.

A subnet ID, which is a 16-bit subnet field, can be used by individual organizations to create a local addressing hierarchy and to identify subnets. A subnet ID is similar to a subnet in IPv4, except that an organization with an IPv6 subnet ID can support up to 65,535 individual subnets.

An interface ID identifies interfaces on a link. The interface ID is unique to the link. In many cases, an interface ID is the same as or based on the link-layer address of an interface. Interface IDs used in aggregatable global unicast and other IPv6 address types have 64 bits and are in the modified EUI-64 format.

Interface IDs are in the modified EUI-64 format in one of the following ways:

- For all IEEE 802 interface types (for example, Ethernet, and Fiber Distributed Data interfaces), the first three octets (24 bits) are the Organizationally Unique Identifier (OUI) of the 48-bit link-layer address (MAC address) of the interface, the fourth and fifth octets (16 bits) are a fixed hexadecimal value of FFFE, and the last three octets (24 bits) are the last three octets of the MAC address. The Universal/Local

(U/L) bit, which is the seventh bit of the first octet, has a value of 0 or 1. Zero indicates a locally administered identifier; 1 indicates a globally unique IPv6 interface identifier.

- For all other interface types (for example, serial, loopback, ATM, Frame Relay, and tunnel interface types—except tunnel interfaces used with IPv6 overlay tunnels), the interface ID is similar to the interface ID for IEEE 802 interface types; however, the first MAC address from the pool of MAC addresses in the router is used as the identifier (because the interface does not have a MAC address).
- For tunnel interface types that are used with IPv6 overlay tunnels, the interface ID is the IPv4 address assigned to the tunnel interface with all zeros in the high-order 32 bits of the identifier



Note For interfaces that use the Point-to-Point Protocol (PPP), where the interfaces at both ends of the connection might have the same MAC address, the interface identifiers at both ends of the connection are negotiated (picked randomly and, if necessary, reconstructed) until both identifiers are unique. The first MAC address in the router is used as the identifier for interfaces using PPP.

If no IEEE 802 interface types are in the router, link-local IPv6 addresses are generated on the interfaces in the router in the following sequence:


1. The router is queried for MAC addresses (from the pool of MAC addresses in the router).
2. If no MAC addresses are available in the router, the serial number of the router is used to form the link-local addresses.
3. If the serial number of the router cannot be used to form the link-local addresses, the router uses a Message Digest 5 (MD5) hash to determine the MAC address of the router from the hostname of the router.

Link-Local Addresses

A link-local address is an IPv6 unicast address that can be automatically configured on any interface using the link-local prefix FE80::/10 (1111 1110 10) and the interface identifier in the modified EUI-64 format. Link-local addresses are used in the Neighbor Discovery Protocol (NDP) and the stateless autoconfiguration process. Nodes on a local link can use link-local addresses to communicate; the nodes do not need globally unique addresses to communicate.

IPv6 routers cannot forward packets that have link-local source or destination addresses to other links.


Figure 8: Link-Local Address Format



IPv4-Compatible IPv6 Addresses

An IPv4-compatible IPv6 address is an IPv6 unicast address that has zeros in the high-order 96 bits of the address and an IPv4 address in the low-order 32 bits of the address. The format of an IPv4-compatible IPv6 address is 0:0:0:0:0:A.B.C.D or ::A.B.C.D. The entire 128-bit IPv4-compatible IPv6 address is used as the IPv6 address of a node and the IPv4 address embedded in the low-order 32 bits is used as the IPv4 address of the node. IPv4-compatible IPv6 addresses are assigned to nodes that support both the IPv4 and IPv6 protocol stacks and are used in automatic tunnels.

Figure 9: IPv4-Compatible IPv6 Address Format



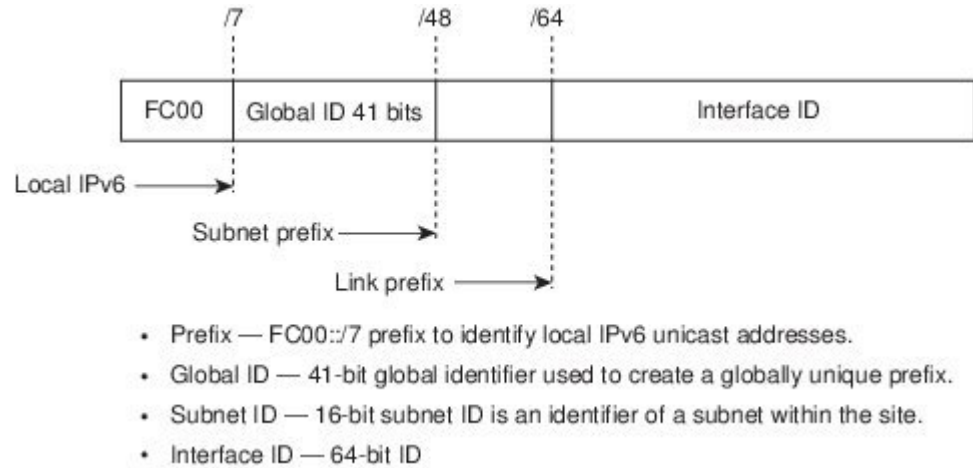
Unique Local Addresses

A unique local address is an IPv6 unicast address that is globally unique and is intended for local communications. It is not expected to be routable on the global Internet and is routable inside of a limited area, such as a site, and it may be routed between a limited set of sites. Applications may treat unique local addresses like global scoped addresses.

A unique local address has the following characteristics:

- It has a globally unique prefix (it has a high probability of uniqueness).
- It has a well-known prefix to allow for easy filtering at site boundaries
- It allows sites to be combined or privately interconnected without creating any address conflicts or requiring renumbering of interfaces that use these prefixes.
- It is ISP-independent and can be used for communications inside of a site without having any permanent or intermittent Internet connectivity.
- If it is accidentally leaked outside of a site through routing or the Domain Name Server (DNS), there is no conflict with any other addresses.

Figure 10: Unique Local Address Structure



23/23/09

Site Local Addresses

Because RFC 3879 deprecates the use of site-local addresses, you should follow the recommendations of unique local addressing (ULA) in RFC 4193 when you configure private IPv6 addresses.

IPv6 Anycast Addresses


An anycast address is an address that is assigned to a set of interfaces that belong to different nodes. A packet sent to an anycast address is delivered to the closest interface—as defined by the routing protocols in use—identified by the anycast address. Anycast addresses are syntactically indistinguishable from unicast addresses because anycast addresses are allocated from the unicast address space. Assigning a unicast address to more than one interface turns a unicast address into an anycast address. You must configure the nodes to which the anycast address to recognize that the address is an anycast address.



Note Anycast addresses can be used only by a router, not a host. Anycast addresses cannot be used as the source address of an IPv6 packet.

The following figure shows the format of the subnet router anycast address; the address has a prefix concatenated by a series of zeros (the interface ID). The subnet router anycast address can be used to reach a router on the link that is identified by the prefix in the subnet router anycast address.

Figure 11: Subnet Router Anycast Address Format




IPv6 Multicast Addresses

An IPv6 multicast address is an IPv6 address that has a prefix of FF00::/8 (1111 1111). An IPv6 multicast address is an identifier for a set of interfaces that belong to different nodes. A packet sent to a multicast address

is delivered to all interfaces identified by the multicast address. The second octet following the prefix defines the lifetime and scope of the multicast address. A permanent multicast address has a lifetime parameter equal to 0; a temporary multicast address has a lifetime parameter equal to 1. A multicast address that has the scope of a node, link, site, or organization, or a global scope, has a scope parameter of 1, 2, 5, 8, or E, respectively. For example, a multicast address with the prefix FF02::/16 is a permanent multicast address with a link scope.

Figure 12: IPv6 Multicast Address Format




IPv6 nodes (hosts and routers) are required to join (where received packets are destined for) the following multicast groups:

- All-nodes multicast group FF02:0:0:0:0:0:0:1 (the scope is link-local)
- Solicited-node multicast group FF02:0:0:0:0:1:FF00:0000/104 for each of its assigned unicast and anycast addresses

IPv6 routers must also join the all-routers multicast group FF02:0:0:0:0:0:0:2 (the scope is link-local).

The solicited-node multicast address is a multicast group that corresponds to an IPv6 unicast or anycast address. IPv6 nodes must join the associated solicited-node multicast group for every unicast and anycast address to which it is assigned. The IPv6 solicited-node multicast address has the prefix FF02:0:0:0:0:1:FF00:0000/104 concatenated with the 24 low-order bits of a corresponding IPv6 unicast or anycast address. For example, the solicited-node multicast address that corresponds to the IPv6 address 2037::01:800:200E:8C6C is FF02::1:FF0E:8C6C. Solicited-node addresses are used in neighbor solicitation messages.

Figure 13: IPv6 Solicited-Node Multicast Address Format



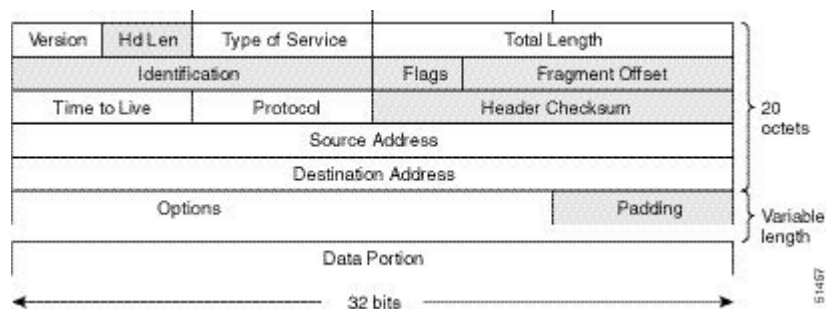


Note IPv6 has no broadcast addresses. IPv6 multicast addresses are used instead of broadcast addresses.

IPv4 Packet Header

The base IPv4 packet header has 12 fields with a total size of 20 octets (160 bits). The 12 fields may be followed by an Options field, which is followed by a data portion that is usually the transport-layer packet. The variable length of the Options field adds to the total size of the IPv4 packet header. The shaded fields of the IPv4 packet header are not included in the IPv6 packet header.

Figure 14: IPv4 Packet Header Format



Simplified IPv6 Packet Header


The base IPv6 packet header has 8 fields with a total size of 40 octets (320 bits). Fragmentation is handled by the source of a packet and checksums at the data link layer and transport layer are used. The User Datagram Protocol (UDP) checksum checks the integrity of the inner packet and the base IPv6 packet header and Options field are aligned to 64 bits, which can facilitate the processing of IPv6 packets.

Table 6: Base IPv6 Packet Header Fields

Field	Description
Version	Similar to the Version field in the IPv4 packet header, except that the field lists number 6 for IPv6 instead of number 4 for IPv4.
Traffic Class	Similar to the Type of Service field in the IPv4 packet header. The Traffic Class field tags packets with a traffic class that is used in differentiated services.
Flow Label	New field in the IPv6 packet header. The Flow Label field tags packets with a specific flow that differentiates the packets at the network layer.
Payload Length	Similar to the Total Length field in the IPv4 packet header. The Payload Length field indicates the total length of the data portion of the packet.

Field	Description
Next Header	Similar to the Protocol field in the IPv4 packet header. The value of the Next Header field determines the type of information that follows the base IPv6 header. The type of information that follows the base IPv6 header can be a transport-layer packet, for example, a TCP or UDP packet, or an Extension Header.
Hop Limit	Similar to the Time to Live field in the IPv4 packet header. The value of the Hop Limit field specifies the maximum number of routers that an IPv6 packet can pass through before the packet is considered invalid. Each router decrements the value by one. Because no checksum is in the IPv6 header, the router can decrement the value without needing to recalculate the checksum, which saves processing resources.
Source Address	Similar to the Source Address field in the IPv4 packet header, except that the field contains a 128-bit source address for IPv6 instead of a 32-bit source address for IPv4.
Destination Address	Similar to the Destination Address field in the IPv4 packet header, except that the field contains a 128-bit destination address for IPv6 instead of a 32-bit destination address for IPv4.

Figure 15: IPv6 Packet Header Format



Optional extension headers and the data portion of the packet are after the eight fields of the base IPv6 packet header. If present, each extension header is aligned to 64 bits. There is no fixed number of extension headers in an IPv6 packet. Each extension header is identified by the Next Header field of the previous header. Typically, the final extension header has a Next Header field of a transport-layer protocol, such as TCP or UDP.

Figure 16: IPv6 Extension Header Format




Table 7: IPv6 Extension Header Types

Header Type	Next Header Value	Description
Hop-by-Hop options header	0	Header that is processed by all hops in the path of a packet. When present, the hop-by-hop options header always follows immediately after the base IPv6 packet header.
Destination Header Options	60	Header that can follow any hop-by-hop options header. The header is processed at the final destination and at each visited address specified by a routing header. Alternatively, the destination options header can follow any Encapsulating Security Payload (ESP) header. The destination options header is processed only at the final destination.
Routing Header	43	Header that is used for source routing.



Note In IPv6, the minimum link MTU is 1280 octets. We recommend that you use an MTU value of 1500 octets for IPv6 links.

CDP IPv6 Address Support

You can use the Cisco Discovery Protocol (CDP) IPv6 address support for the neighbor information feature to transfer IPv6 addressing information between two Cisco devices. Cisco Discovery Protocol support for IPv6 addresses provides IPv6 information to network management products and troubleshooting tools.

ICMP for IPv6

You can use ICMP in IPv6 to provide information about the health of the network. ICMPv6, the version that works with IPv6, reports errors if packets cannot be processed correctly and sends informational messages about the status of the network. For example, if a router cannot forward a packet because it is too large to be sent out on another network, the router sends out an ICMPv6 message to the originating host. Additionally, ICMP packets in IPv6 are used in IPv6 neighbor discovery and path MTU discovery. The path MTU discovery process ensures that a packet is sent using the largest possible size that is supported on a specific route.

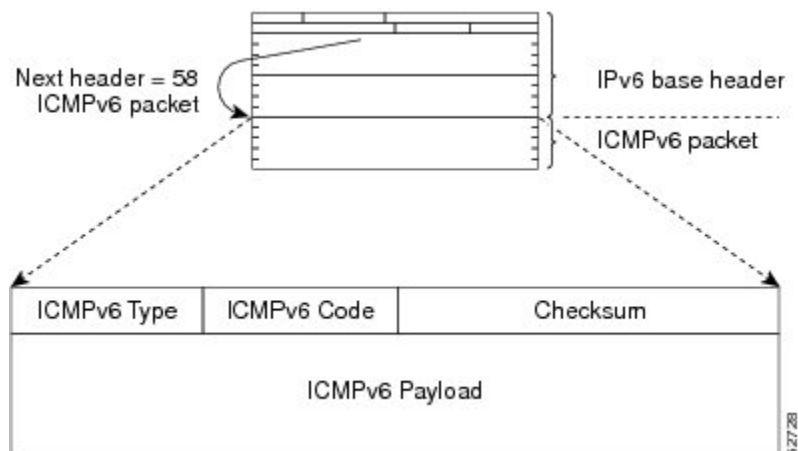
A value of 58 in the Next Header field of the base IPv6 packet header identifies an IPv6 ICMP packet. The ICMP packet follows all the extension headers and is the last piece of information in the IPv6 packet. Within the IPv6 ICMP packets, the ICMPv6 Type and ICMPv6 Code fields identify IPv6 ICMP packet specifics, such as the ICMP message type. The value in the Checksum field is computed by the sender and checked by the receiver from the fields in the IPv6 ICMP packet and the IPv6 pseudo header.



Note The IPv6 header does not have a checksum. But a checksum on the transport layer can determine if packets have not been delivered correctly. All checksum calculations that include the IP address in the calculation must be modified for IPv6 to accommodate the new 128-bit address. A checksum is generated using a pseudo header.

The ICMPv6 Payload field contains error or diagnostic information that relates to IP packet processing.

Figure 17: IPv6 ICMP Packet Header Format



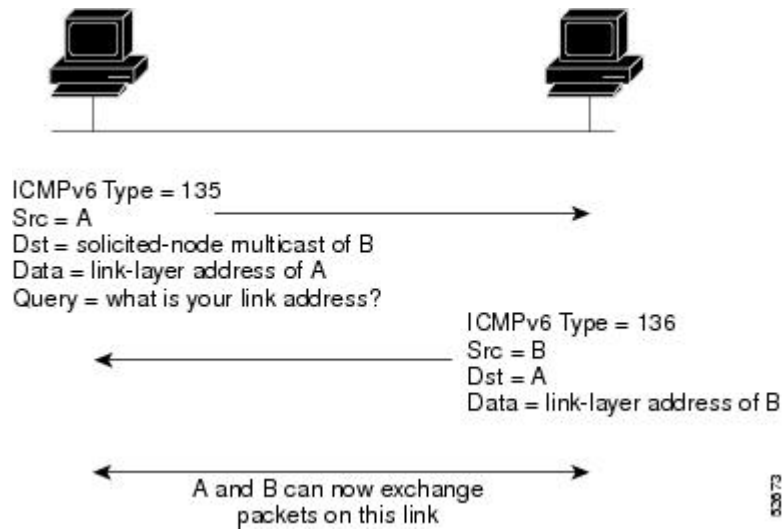
IPv6 Neighbor Discovery

You can use the IPv6 Neighbor Discovery Protocol (NDP) to determine whether a neighboring router is reachable. IPv6 nodes use neighbor discovery to determine the addresses of nodes on the same network (local link), to find neighboring routers that can forward their packets, to verify whether neighboring routers are reachable or not, and to detect changes to link-layer addresses. NDP uses ICMP messages to detect whether packets are sent to neighboring routers that are unreachable.

IPv6 Neighbor Solicitation Message

A node sends a neighbor solicitation message, which has a value of 135 in the Type field of the ICMP packet header, on the local link when it wants to determine the link-layer address of another node on the same local link. The source address is the IPv6 address of the node that sends the neighbor solicitation message. The destination address is the solicited-node multicast address that corresponds to the IPv6 address of the destination node. The neighbor solicitation message also includes the link-layer address of the source node.

Figure 18: IPv6 Neighbor Discovery-Neighbor Solicitation Message



After receiving the neighbor solicitation message, the destination node replies by sending a neighbor advertisement message, which has a value of 136 in the Type field of the ICMP packet header, on the local link. The source address is the IPv6 address of the node (the IPv6 address of the node interface that sends the neighbor advertisement message). The destination address is the IPv6 address of the node that sends the neighbor solicitation message. The data portion includes the link-layer address of the node that sends the neighbor advertisement message.

After the source node receives the neighbor advertisement, the source node and destination node can communicate.

Neighbor solicitation messages can verify the reachability of a neighbor after a node identifies the link-layer address of a neighbor. When a node wants to verify the reachability of a neighbor, it uses the destination address in a neighbor solicitation message as the unicast address of the neighbor.

Neighbor advertisement messages are also sent when there is a change in the link-layer address of a node on a local link. When there is a change, the destination address for the neighbor advertisement is the all-nodes multicast address.

Neighbor unreachability detection identifies the failure of a neighbor or the failure of the forward path to the neighbor and is used for all paths between hosts and neighboring nodes (hosts or routers). Neighbor unreachability detection is performed for neighbors to which only unicast packets are being sent and is not performed for neighbors to which multicast packets are being sent.

A neighbor is considered reachable when a positive acknowledgment is returned from the neighbor (indicating that packets previously sent to the neighbor have been received and processed). A positive acknowledgment—from an upper-layer protocol (such as TCP)—indicates that a connection is making forward progress (reaching its destination). If packets are reaching the peer, they are also reaching the next-hop neighbor of the source. Forward progress is also a confirmation that the next-hop neighbor is reachable.

For destinations that are not on the local link, forward progress implies that the first-hop router is reachable. When acknowledgments from an upper-layer protocol are not available, a node probes the neighbor using unicast neighbor solicitation messages to verify that the forward path is still working. The return of a solicited neighbor advertisement message from the neighbor is a positive acknowledgment that the forward path is still working (neighbor advertisement messages that have the solicited flag set to a value of 1 are sent only in response to a neighbor solicitation message). Unsolicited messages confirm only the one-way path from the

source to the destination node; solicited neighbor advertisement messages indicate that a path is working in both directions.



Note A neighbor advertisement message that has the solicited flag set to a value of 0 is not considered as a positive acknowledgment that the forward path is still working.

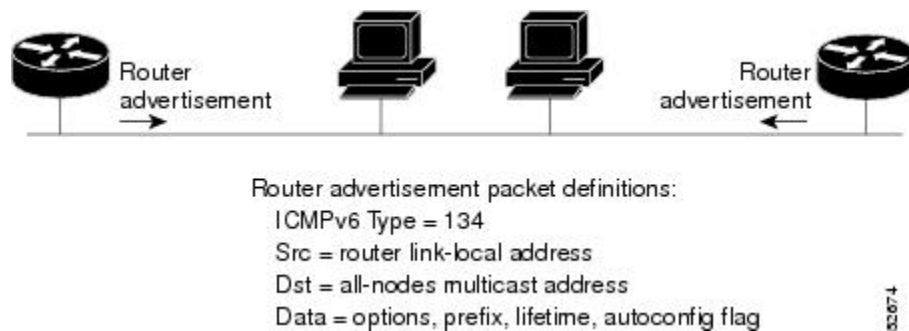
Neighbor solicitation messages are also used in the stateless autoconfiguration process to verify the uniqueness of unicast IPv6 addresses before the addresses are assigned to an interface. Duplicate address detection is performed first on a new, link-local IPv6 address before the address is assigned to an interface (the new address remains in a tentative state while duplicate address detection is performed). A node sends a neighbor solicitation message with an unspecified source address and a tentative link-local address in the body of the message. If another node is already using that address, the node returns a neighbor advertisement message that contains the tentative link-local address. If another node is simultaneously verifying the uniqueness of the same address, that node also returns a neighbor solicitation message. If no neighbor advertisement messages are received in response to the neighbor solicitation message and no neighbor solicitation messages are received from other nodes that are attempting to verify the same tentative address, the node that sent the original neighbor solicitation message considers the tentative link-local address to be unique and assigns the address to the interface.

IPv6 Router Advertisement Message

Router advertisement (RA) messages, which have a value of 134 in the Type field of the ICMP packet header, are periodically sent out to each configured interface of an IPv6 router. For stateless autoconfiguration to work properly, the advertised prefix length in RA messages must always be 64 bits.

The RA messages are sent to the all-nodes multicast address.

Figure 19: Neighbor Discovery—RA Message



RA messages typically include the following information:

- One or more onlink IPv6 prefixes that nodes on the local link can use to automatically configure their IPv6 addresses
- Life-time information for each prefix included in the advertisement
- Sets of flags that indicate the type of autoconfiguration (stateless or stateful) that can be completed
- Default router information (whether the router sending the advertisement should be used as a default router and, if so, the amount of time in seconds that the router should be used as a default router)

- Additional information for hosts, such as the hop limit and MTU that a host should use in packets that it originates

RAs are also sent in response to router solicitation messages. Router solicitation messages, which have a value of 133 in the Type field of the ICMP packet header, are sent by hosts at system startup so that the host can immediately autoconfigure without needing to wait for the next scheduled RA message. The source address is usually the unspecified IPv6 address (0:0:0:0:0:0:0:0). If the host has a configured unicast address, the unicast address of the interface that sends the router solicitation message is used as the source address in the message. The destination address is the all-routers multicast address with a scope of the link. When an RA is sent in response to a router solicitation, the destination address in the RA message is the unicast address of the source of the router solicitation message.

You can configure the following RA message parameters:

- The time interval between periodic RA messages
- The router life-time value, which indicates the usefulness of a router as the default router (for use by all nodes on a given link)
- The network prefixes in use on a given link
- The time interval between neighbor solicitation message retransmissions (on a given link)
- The amount of time that a node considers a neighbor reachable (for use by all nodes on a given link)

The configured parameters are specific to an interface. The sending of RA messages (with default values) is automatically enabled on Ethernet interfaces. For other interface types, you must enter the **no ipv6 nd suppress-ra** command to send RA messages. You can disable the RA message feature on individual interfaces by entering the **ipv6 nd suppress-ra** command.

IPv6 Router Advertisement Options for DNS Configuration

Most of the internet services are identified by a Domain Name Server (DNS) name. Cisco NX-OS IPv6 Router Advertisement (RA) provides the following two options to allow IPv6 hosts to perform automatic DNS configuration:

- Recursive DNS Server (RDNSS)
- DNS Search List (DNSSL)


RDNSS contains the address of recursive DNS servers that help in DNS name resolution in IPv6 hosts. DNS Search List is a list of DNS suffix domain names used by IPv6 hosts when they perform DNS query searches.

For more information on RA options for DNS configuration, refer IETF RFC 6106.

IPv6 Neighbor Redirect Message

Routers send neighbor redirect messages to inform hosts of better first-hop nodes on the path to a destination. A value of 137 in the Type field of the ICMP packet header identifies an IPv6 neighbor redirect message.

Figure 20: IPv6 Neighbor Discovery—Neighbor Redirect Message



Note A router must be able to determine the link-local address for each of its neighboring routers in order to ensure that the target address (the final destination) in a redirect message identifies the neighbor router by its link-local address. For static routing, you should specify the address of the next-hop router using the link-local address of the router. For dynamic routing, you must configure all IPv6 routing protocols to exchange the link-local addresses of neighboring routers.

After forwarding a packet, a router sends a redirect message to the source of the packet under the following circumstances:

- The destination address of the packet is not a multicast address.
- The packet was not addressed to the router.
- The packet is about to be sent out the interface on which it was received.
- The router determines that a better first-hop node for the packet resides on the same link as the source of the packet.
- The source address of the packet is a global IPv6 address of a neighbor on the same link or a link-local address.

Virtualization Support for IPv6

IPv6 supports virtual routing and forwarding (VRF) instances. VRFs exist within virtual device contexts (VDCs). By default, Cisco NX-OS places you in the default VDC and default VRF unless you specifically configure another VDC and VRF. For more information, see the *Cisco Nexus 7000 Series NX-OS Virtual Device Context Configuration Guide*.

Prerequisites for IPv6

IPv6 has the following prerequisites:

- You must be familiar with IPv6 basics such as IPv6 addressing, IPv6 header information, ICMPv6, and the IPv6 Neighbor Discovery (ND) Protocol.
- Ensure that you follow the memory/processing guidelines when you make a device a dual-stack device (IPv4/IPv6).

Guidelines and Limitations for Configuring IPv6

IPv6 has the following configuration guidelines and limitations:

- IPv6 packets are transparent to Layer 2 LAN switches because the switches do not examine Layer 3 packet information before forwarding IPv6 frames. IPv6 hosts can be directly attached to Layer 2 LAN switches.
- You can configure multiple IPv6 global addresses within the same prefix on an interface. However, multiple IPv6 link-local addresses on an interface are not supported.
- It supports contiguous masks only for both IPv4 and IPv6 addresses and does not support discontinuous masks IPv6 and IPv4 filters.
- Each interface can be configured with a maximum of 255 global IPv6 addresses and a maximum of 255 anycast IPv6 addresses.
- Because RFC 3879 deprecates the use of site-local addresses, you should configure private IPv6 addresses according to the recommendations of unique local addressing (ULA) in RFC 4193.
- F2 Series modules do not support IPv6 tunnels.
- On F2 Series modules, you must disable IGMP optimized multicast flooding (OMF) on any VLANs that require any IPv6 packet forwarding (unicast or multicast). IPv6 neighbor discovery functions correctly only in a VLAN with the OMF feature disabled. To disable OMF, use the **no ip igmp snooping optimised-multicast-flood** command in VLAN configuration mode. With OMF disabled, unknown IPv4 multicast traffic (as well as all IPv6 multicast traffic) is flooded to all ports in the VLAN. Note that unknown multicast traffic refers to multicast packets with an active source but no receivers (and therefore no group forwarding entry in the hardware) in the ingress VLAN.
- IPv6 static route next hop link-local address cannot be configured at any local interface.

Default Settings for IPv6

Parameters	Default
ND reachable time	0 milliseconds
neighbor solicitation retransmit interval	1000 milliseconds

By default, IPv6 source routing is enabled on the switch. To disable IPv6 source routing configure no **ipv6 source-route** command on the switch.

Configuring IPv6

Configuring IPv6 Addressing

You must configure an IPv6 address on an interface so that the interface can forward IPv6 traffic. When you configure a global IPv6 address on an interface, it automatically configures a link-local address and activates IPv6 for that interface.



Note Each interface can be configured with a maximum of 255 global IPv6 addresses and a maximum of 255 anycast IPv6 addresses.

Before you begin

Ensure that you are in the correct VDC (or use the **switchto vdc** command).

Procedure

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# interface ethernet number	Enters interface configuration mode.
Step 3	switch(config-if)# ipv6 address { <i>address</i> [eui64] [route-preference preference] [secondary] tag tag-id] or switch(config-if)# ipv6 address ipv6-address use-link-local-only	Specifies an IPv6 address assigned to the interface and enables IPv6 processing on the interface. Entering the ipv6 address command configures global IPv6 addresses with an interface identifier (ID) in the low-order 64 bits of the IPv6 address. Only the 64-bit network prefix for the address needs to be specified; the last 64 bits are automatically computed from the interface ID.

	Command or Action	Purpose
		Entering the ipv6 address use-link-local-only command configures a link-local address on the interface that is used instead of the link-local address that is automatically configured when IPv6 is enabled on the interface. This command enables IPv6 processing on an interface without configuring an IPv6 address.
Step 4	(Optional) switch(config-if)# show ip interface	Displays interfaces configured for IPv4.
Step 5	(Optional) switch(config)# copy running-config startup-config	Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

Example

This example shows how to assign an IPv6 address:

```
switch# configure terminal
switch(config)# interface ethernet 2/3
switch(config-if)# ipv6 address 2001:db8::/64 eui64
switch(config-if)# copy running-config startup-config
switch(config-if)#
```

This example shows how to display an IPv6 interface:

```
switch# configure terminal
switch(config)# show ipv6 interface ethernet 3/1
Ethernet3/1, Interface status: protocol-down/link-down/admin-down, iod: 36
IPv6 address: 0dc3:0dc3:0000:0000:0218:baff:fed8:239d
IPv6 subnet: 0dc3:0dc3:0000:0000:0000:0000:0000:0000/64
IPv6 link-local address: fe80::0218:baff:fed8:239d (default)
IPv6 multicast routing: disabled
IPv6 multicast groups locally joined:
ff02::0001:ffd8:239d ff02::0002 ff02::0001 ff02::0001:ffd8:239d
IPv6 multicast (S,G) entries joined: none
IPv6 MTU: 1500 (using link MTU)
IPv6 RP inbound packet-filtering policy: none
IPv6 RP outbound packet-filtering policy: none
IPv6 inbound packet-filtering policy: none
IPv6 outbound packet-filtering policy: none
IPv6 interface statistics last reset: never
IPv6 interface RP-traffic statistics: (forwarded/originated/consumed)
Unicast packets: 0/0/0
Unicast bytes: 0/0/0
Multicast packets: 0/0/0
Multicast bytes: 0/0/0
```

Configuring IPv6 Neighbor Discovery

You can configure IPv6 neighbor discovery on the router. Neighbor Discovery (ND) enables IPv6 nodes and routers to determine the link-layer address of a neighbor on the same link, find neighboring routers, and keep track of neighbors.

Before you begin

Ensure that you are in the correct VDC (or use the **switchto vdc** command).

Procedure

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# interface ethernet <i>number</i>	Enters interface configuration mode.
Step 3	switch(config-if)# ipv6 nd [hop-limit <i>hop-limit</i> managed-config-flag mtu <i>mtu</i> ns-interval <i>interval</i> other-config-flag prefix ra-interval <i>interval</i> ra-lifetime <i>lifetime</i> reachable-time <i>time</i> redirects retrans-timer <i>time</i> suppress-ra]	<p>Specifies an IPv6 address assigned to the interface and enables IPv6 processing on the interface.</p> <ul style="list-style-type: none"> • hop-limit <i>hop-limit</i>— Advertises the hop limit in IPv6 neighbor discovery packets. The range is from 0 to 255. • managed-config-flag— Advertises in ICMPv6 router-advertisement messages to use stateful address auto-configuration to obtain address information. • mtu <i>mtu</i>—Advertises the maximum transmission unit (MTU) in ICMPv6 router-advertisement messages on this link. The range is from 1280 to 65535 bytes. • ns-interval <i>interval</i>—Configures the retransmission interval between IPv6 neighbor solicitation messages. The range is from 1000 to 3600000 milliseconds. • other-config-flag—Indicates in ICMPv6 router-advertisement messages that hosts use stateful auto-configuration to obtain nonaddress related information. • prefix—Advertises the IPv6 prefix in the router-advertisement messages. • ra-interval <i>interval</i>—Configures the interval between sending ICMPv6 router-advertisement messages. The range is from 4 to 1800 seconds. • ra-lifetime <i>lifetime</i>—Advertises the lifetime of a default router in ICMPv6 router-advertisement messages. The range is from 0 to 9000 seconds. • reachable-time <i>time</i>—Advertises the time when a node considers a neighbor up after receiving a reachability confirmation in

	Command or Action	Purpose
		<p>ICMPv6 router-advertisement messages. The range is from 0 to 9000 seconds.</p> <ul style="list-style-type: none"> • redirects—Enables sending ICMPv6 redirect messages. • retrans-timer <i>time</i>—Advertises the time between neighbor-solicitation messages in ICMPv6 router-advertisement messages. The range is from 0 to 9000 seconds. • suppress-ra— Disables sending ICMPv6 router-advertisement messages.
Step 4	<p>Required: switch(config-if)# ipv6 nd prefix {<i>ipv6-address/prefix-length</i> default} {valid-lifetime infinite no-advertise} {preferred-lifetime infinite} [no-autoconfig] [no-onlink] [off-link]</p>	<p>Advertises the IPv6 prefix in the router advertisement messages.</p> <ul style="list-style-type: none"> • <i>valid-lifetime</i>—The amount of time (in seconds) that the specified IPv6 prefix is advertised as being valid. • infinite—Specifies that the valid lifetime is infinite. • no-advertise—Specifies that the prefix is not advertised. • <i>preferred-lifetime</i>—The amount of time (in seconds) that the specified IPv6 prefix is advertised as being preferred. • no-autoconfig—Indicates to hosts on the local link that the specified prefix cannot be used for IPv6 autoconfiguration. The prefix will be advertised with the A-bit clear. • no-onlink—Configures the specified prefix as not on-link. The prefix will be advertised with the L-bit clear. • off-link—Configures the specified prefix as off-link. The prefix will be advertised with the L-bit clear. The prefix will not be inserted into the routing table as a connected prefix. If the prefix is already present in the routing table as a connected prefix (for example, because the prefix was also configured using the ipv6 address command), it will be removed.
Step 5	(Optional) switch(config-if)# show ip nd interface	Displays interfaces configured for IPv6 neighbor discovery.

	Command or Action	Purpose
Step 6	(Optional) switch(config)# copy running-config startup-config	Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

Example

This example shows how to configure IPv6 neighbor discovery reachable time:

```
switch# configure terminal
switch(config)# interface ethernet 3/1
switch(config-if)# ipv6 nd reachable-time 10
switch(config-if)# copy running-config startup-config
switch(config-if)#
```

This example shows how to display an IPv6 interface:

```
switch# configure terminal
switch(config)# show ipv6 nd interface ethernet 3/1

ICMPv6 ND Interfaces for VRF "default"
Ethernet3/1, Interface status: protocol-down/link-down/admin-down
IPv6 address: 0dc3:0dc3:0000:0000:0218:baff:fed8:239d
ICMPv6 active timers:
Last Neighbor-Solicitation sent: never
Last Neighbor-Advertisement sent: never
Last Router-Advertisement sent:never
Next Router-Advertisement sent in: 0.000000
Router-Advertisement parameters:
Periodic interval: 200 to 600 seconds
Send "Managed Address Configuration" flag: false
Send "Other Stateful Configuration" flag: false
Send "Current Hop Limit" field: 64
Send "MTU" option value: 1500
Send "Router Lifetime" field: 1800 secs
Send "Reachable Time" field: 10 ms
Send "Retrans Timer" field: 0 ms
Neighbor-Solicitation parameters:
NS retransmit interval: 1000 ms
ICMPv6 error message parameters:
Send redirects: false
Send unreachable: false
```

This example shows how to include the IPv6 prefix 2001:0DB8::/35 in router advertisements that are sent out Ethernet interface 0/0 with a valid lifetime of 1000 seconds and a preferred lifetime of 900 seconds:

```
switch(config)# interface ethernet 0/0
switch(config-if)# ipv6 nd prefix 2001:0DB8::/35 1000 900
```

Configuring Optional IPv6 Neighbor Discovery

You can use the following optional IPv6 neighbor discovery commands:

Command	Purpose
ipv6 nd cache limit <i>max-nd-adj</i> [syslog <i>syslogs-per-second</i>]	<p>Configures the maximum number of entries in the neighbor adjacency table. The range is from 1 to 409600.</p> <p>The syslog keyword configures the number of system logs per second. The range is from 1 to 1000.</p> <p>If you configure a limit for IPv6 neighbor discovery entries, system logs appear if you try to add an adjacency after reaching the configured limit.</p> <p>Note You cannot unconfigure the cache limit until the total number of current adjacencies is less than 131,072.</p>
ipv6 nd dad attempts <i>number</i>	Sets the number of consecutive neighbor solicitation messages that the device sends from the IPv6 interface for duplicate address detection (DAD) validation. The default value is 1 attempt.
ipv6 nd fast-path	Improves the performance of glean packets by reducing the processing of the packets in the supervisor. It applies to glean packets where the destination IP address is part of the same subnet and does not apply to packets where the destination IP address is in a different subnet. The default is enabled.
ipv6 nd hop-limit	Configures the maximum number of hops used in router advertisements and all IPv6 packets that are originated by the router.
ipv6 nd managed-config-flag	Sets the managed address configuration flag in IPv6 router advertisements.
ipv6 nd mtu	Sets the maximum transmission unit (MTU) size of IPv6 packets sent on an interface.
ipv6 nd ns-interval	Configures the interval between IPv6 neighbor solicitation retransmissions on an interface.
ipv6 nd other-config-flag	Configures the other stateful configuration flag in IPv6 router advertisements.
ipv6 nd ra-interval	Configures the interval between IPv6 router advertisement (RA) transmissions on an interface.
ipv6 nd ra-lifetime	Configures the router lifetime value in IPv6 router advertisements on an interface.
ipv6 nd reachable-time	Configures the amount of time that a remote IPv6 node is considered reachable after some reachability confirmation event has occurred.

Command	Purpose
ipv6 nd redirects	Enables ICMPv6 redirect messages to be sent.
ipv6 nd retrans-timer	Configures the advertised time between neighbor solicitation messages in router advertisements.
ipv6 nd suppress-ra	Suppresses IPv6 router advertisement transmissions on a LAN interface.

Configuring Recursive DNS Server (RDNSS)

You can configure up to eight DNS servers to advertise with Router Advertisement. You can also remove one or more DNS servers from the advertising list by using the **no** form of the command.

Before you begin

Ensure that you are in the correct VDC (or use the **switchto vdc** command).

Procedure

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters the global configuration mode.
Step 2	switch(config)# interface ethernet <i>number</i>	Enters interface configuration mode.
Step 3	switch(config-if)# ipv6 nd ra dns server <i>ipv6-addr</i> [<i>rdns-life</i> infinite] sequence <i>sequence-num</i>	Configures the recursive DNS server. You can specify the life time and the sequence of the server.
Step 4	switch(config-if)# show ipv6 nd ra dns server [interface <i>interface</i>]	(Optional) Displays the configured RDNSS list.
Step 5	switch(config-if)# ipv6 nd ra dns server suppress	(Optional) Disables the configured server list.
Step 6	switch(config-if)# no ipv6 nd ra dns server <i>ipv6-addr</i> [<i>rdns-life</i> infinite] sequence <i>sequence-num</i>	Removes a server from the RDNSS list.

Example

The following example shows how to configure Recursive DNS Server list on Ethernet 3/3 and verify the same.

```
switch# configure terminal
switch(config)# interface ethernet 3/3
switch(config-if)# ipv6 nd ra dns server 1::1 1000 sequence 0
switch(config-if)# ipv6 nd ra dns server 2::1 infinite sequence 1

switch(config)# show ipv6 nd ra dns server

Recursive DNS Server List on: mgmt0
```

```

Suppress DNS Server List: No
Recursive DNS Server List on: Ethernet3/3
  Suppress DNS Server List: No
  DNS Server 1: 1::1 Lifetime:1000 seconds Sequence:0
  DNS Server 2: 2::1 Infinite Sequence:1

```

Configuring DNS Search List (DNSSL)

You can configure up to eight DNS search lists to advertise with Router Advertisement. You can also remove one or more DNS search lists from the advertising list by using the **no** form of the command.

Before you begin

Ensure that you are in the correct VDC (or use the **switchto vdc** command).

Procedure

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters the global configuration mode.
Step 2	switch(config)# interface ethernet <i>number</i>	Enters interface configuration mode.
Step 3	switch(config-if)# ipv6 nd ra dns search-list <i>list</i> [<i>dnssl-life</i> infinite] sequence <i>sequence-num</i>	Configures the DNS search list. You can specify the life time and the sequence of the search list.
Step 4	switch(config-if)# show ipv6 nd ra dns search-list [<i>interface interface</i>]	(Optional) Displays the configured DNS search list.
Step 5	switch(config-if)# ipv6 nd ra dns search-list suppress	(Optional) Disables the configured search list.
Step 6	switch(config-if)# no ipv6 nd ra dns search-list <i>list</i> [<i>dnssl-life</i> infinite] sequence <i>sequence-num</i>	(Optional) Removes a search list from the RA.

Example

The following example shows how to configure DNS Search list on Ethernet 3/3 and verify the same.

```

switch# configure terminal
switch(config)# interface ethernet 3/3
switch(config-if)# ipv6 nd ra dns search-list cisco.com 100 sequence 1
switch(config-if)# ipv6 nd ra dns search-list ind.cisco.com 100 sequence 2

switch(config)# show ipv6 nd ra dns search-list

DNS Search List on: mgmt0
Suppress DNS Search List: No
DNS Search List on: Ethernet3/3
  Suppress DNS Search List: No
  DNS Server 1:cisco.com 100 Sequence:1
  DNS Server 2:ind.cisco.com 100 Sequence:2

```


Configuring IPv6 Packet Verification

Cisco NX-OS supports an Intrusion Detection System (IDS) that checks for IPv6 packet verification. You can enable or disable these IDS checks.

Procedure

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# hardware ip verify address { destination zero identical reserved source multicast }	Performs the following IDS checks on the IPv6 address: <ul style="list-style-type: none"> • destination zero—Drops IPv6 packets if the destination IP address is ::. • identical—Drops IPv6 packets if the source IPv6 address is identical to the destination IPv6 address. • reserved—Drops IPv6 packets if the IPv6 address is ::1. • source multicast—Drops IPv6 packets if the IPv6 source address is in the FF00::/8 range (multicast).
Step 3	switch(config)# hardware ip verify length { consistent maximum { max-frag max-tcp udp }}	Performs the following IDS checks on the IPv6 address: <ul style="list-style-type: none"> • consistent—Drops IPv6 packets where the Ethernet frame size is greater than or equal to the IPv6 packet length plus the Ethernet header. • maximum max-frag—Drops IPv6 packets if the formula (IPv6 Payload Length - IPv6 Extension Header Bytes) + (Fragment Offset * 8) is greater than 65536.. • maximum max-tcp—Drops IPv6 packets if the TCP length is greater than the IP payload length. • maximum max-udp—Drops IPv6 packets if the TCP length is less than the UDP packet length.
Step 4	switch(config)# hardware ipv6 verify tcp tiny-frag	Drops TCP packets if the IPv6 fragment offset is 1, or if the IPv6 fragment offset is 0 and the IP payload length is less than 16.
Step 5	switch(config)# hardware ipv6 verify version	Drops TCP packets if the EtherType is not set to 6 (IPv6).

	Command or Action	Purpose
Step 6	switch(config)# show hardware forwarding ip verify	Displays the IPv6 packet verification configuration.
Step 7	switch(config)# copy running-config startup-config	(Optional) Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

Verifying the IPv6 Configuration

Use one of the following commands to verify the configuration:

Command	Purpose
show hardware forwarding ip verify	Shows the IPv4 and IPv6 packet verification configuration.
show ipv6 interface	Displays IPv6-related interface information.
show ipv6 adjacency	Displays the adjacency table.
show ipv6 icmp	Displays ICMP IPv6 information.
show ipv6 nd	Displays IPv6 neighbor discovery information.
show ipv6 neighbor	Displays IPv6 neighbor entry.

Configuration Example for IPv6

```
switch# configure terminal
switch(config)# interface ethernet 3/1
switch(config-if)# ipv6 address 2001:db8::/64 eui64
switch(config-if)# ipv6 nd reachable-time 10
switch(config-if)#
```

Related Documents for IPv6

For more information related to IP CLI commands, see the *Cisco Nexus 7000 Series NX-OS Unicast Routing Command Reference*.

Standards for IPv6

No new or modified standards are supported by this feature, and support for existing standards has not been modified by this feature.

Feature History for IPv6

This table includes only the updates for those releases that have resulted in additions or changes to the feature.

Table 9: Feature History for IPv6

Feature Name	Release	Feature Information
Duplicate address detection	6.2(2)	Added the ability to set the number of consecutive neighbor solicitation messages that the device sends from the IPv6 interface.
Glean optimization	6.2(2)	Added the fast-path keyword to the ipv6 nd command to improve the performance of glean packets by reducing the processing of the packets in the supervisor.
IPv6	6.2(2)	Added the ability to configure the maximum number of neighbor discovery entries in the neighbor adjacency table.
IPv6	6.0(1)	Updated for F2 Series modules.
IPv6	5.0(2)	Added support for IPv6 path MTU discovery.
IPv6	4.1(3)	Changed platform {ip ipv6} verify command to the hardware {ip ipv6} verify command.
IPv6	4.0(3)	Added the tag keyword to the ipv6 address command.
IPv6	4.0(1)	This feature was introduced.



CHAPTER 5

Configuring DNS

This chapter contains the following sections:

- [Finding Feature Information, on page 75](#)
- [Information About DNS Clients, on page 75](#)
- [Prerequisites for DNS Clients, on page 76](#)
- [Guidelines and Limitations for DNS Clients, on page 77](#)
- [Default Settings for DNS Client Parameters, on page 77](#)
- [Configuring DNS Clients, on page 77](#)
- [Verifying the DNS Client Configuration, on page 79](#)
- [Configuration Examples for DNS Clients, on page 79](#)
- [Related Documents for DNS Clients, on page 79](#)
- [Standards for DNS Clients, on page 79](#)
- [Feature History for DNS, on page 79](#)

Finding Feature Information

Your software release might not support all the features documented in this module. For the latest caveats and feature information, see the Bug Search Tool at <https://tools.cisco.com/bugsearch/> and the release notes for your software release. To find information about the features documented in this module, and to see a list of the releases in which each feature is supported, see the "New and Changed Information" chapter or the Feature History table in this chapter.

Information About DNS Clients

DNS Client Overview

If your network devices require connectivity with devices in networks for which you do not control the name assignment, you can assign device names that uniquely identify your devices within the entire internetwork using the domain name server (DNS). DNS uses a hierarchical scheme for establishing host names for network nodes, which allows local control of the segments of the network through a client-server scheme. The DNS system can locate a network device by translating the hostname of the device into its associated IP address.

On the Internet, a domain is a portion of the naming hierarchy tree that refers to general groupings of networks based on the organization type or geography. Domain names are pieced together with periods (.) as the delimiting characters. For example, Cisco is a commercial organization that the Internet identifies by a *com* domain, so its domain name is *cisco.com*. A specific hostname in this domain, the File Transfer Protocol (FTP) system, for example, is identified as *ftp.cisco.com*.

DNS Name Servers

Name servers keep track of domain names and know the parts of the domain tree for which they have complete information. A name server may also store information about other parts of the domain tree. To map domain names to IP addresses in Cisco NX-OS, you must identify the hostnames, specify a name server, and enable the DNS service.

Cisco NX-OS allows you to statically map IP addresses to domain names. You can also configure Cisco NX-OS to use one or more domain name servers to find an IP address for a host name.

DNS Operation

A name server handles client-issued queries to the DNS server for locally defined hosts within a particular zone as follows:

- An authoritative name server responds to DNS user queries for a domain name that is under its zone of authority by using the permanent and cached entries in its own host table. If the query is for a domain name that is under its zone of authority but for which it does not have any configuration information, the authoritative name server replies that no such information exists.
- A name server that is not configured as the authoritative name server responds to DNS user queries by using information that it has cached from previously received query responses. If no router is configured as the authoritative name server for a zone, queries to the DNS server for locally defined hosts receive nonauthoritative responses.

Name servers answer DNS queries (forward incoming DNS queries or resolve internally generated DNS queries) according to the forwarding and lookup parameters configured for the specific domain.

High Availability for DNS Clients

Cisco NX-OS supports stateless restarts for the DNS client. After a reboot or supervisor switchover, Cisco NX-OS applies the running configuration.

Virtualization Support for DNS Clients

Cisco NX-OS supports multiple instances of the DNS clients that run on the same system. You can configure a DNS client in each virtual device connect (VDC). You can optionally have a different DNS client configuration in each virtual routing and forwarding (VRF) instance within a VDC. By default, Cisco NX-OS places you in the default VDC and default VRF unless you specifically configure another VDC and VRF. See the *Cisco NX-OS Virtual Device Context Configuration Guide*.

Prerequisites for DNS Clients

- You must have a DNS name server on your network.

Guidelines and Limitations for DNS Clients

- You configure the DNS client in a specific VRF. If you do not specify a VRF, Cisco NX-OS uses the default VRF.
- Cisco NX-OS does not support underscore in a DNS name. Hence do not use underscore in a DNS name.
- If you are familiar with the Cisco IOS CLI, be aware that the Cisco NX-OS commands for this feature might differ from the Cisco IOS commands that you would use.

Default Settings for DNS Client Parameters

The table below lists the default settings for DNS client parameters.

Table 10: Default DNS Client Parameters

Parameters	Default
DNS client	Enabled

Configuring DNS Clients

Configuring the DNS Client

Before you begin

- Ensure that you have a domain name server on your network.
- Ensure that you are in the correct VDC (or use the `switchto vdc` command).

Procedure

	Command or Action	Purpose
Step 1	<code>switch# configure terminal</code>	Enters global configuration mode.
Step 2	<code>switch(config)# ip host name address1 [address2... address6]</code>	Defines up to six static hostname-to-address mappings in the hostname cache. The address can be either an IPv4 address or an IPv6 address.
Step 3	(Optional) <code>switch(config)# ip domain-name name [use-vrf vrf-name]</code>	Defines the default domain name that Cisco NX-OS uses to complete unqualified host names. You can optionally define a VRF that Cisco NX-OS uses to resolve this domain name if it cannot be resolved in the VRF that you configured this domain name under.

	Command or Action	Purpose
		<p>Cisco NX-OS appends the default domain name to any hostname that does not contain a complete domain name before starting a domain-name lookup.</p> <p>use-vrf is used as a DNS query supposed to be sending on a different VRF and listening for the reply on a different VRF. Example: DNS query is sent over VRF RED while the response should come on VRF Default.</p>
Step 4	switch(config)# ip dns source-interface [<i>loopback X different interface</i>]	Defines what will be the source IP for the DNS Query which will be sent out. When DNS server tries to answer back, it will use the Loopback0 as the destination and there should be a valid return route.
Step 5	(Optional) switch(config)# ip domain-list name [use-vrf vrf-name]	<p>Defines additional domain names that Cisco NX-OS can use to complete unqualified hostnames. You can optionally define a VRF that Cisco NX-OS uses to resolve these domain names if they cannot be resolved in the VRF that you configured this domain name under.</p> <p>Cisco NX-OS uses each entry in the domain list to append that domain name to any hostname that does not contain a complete domain name before starting a domain-name lookup. Cisco NX-OS continues this process for each entry in the domain list until it finds a match.</p>
Step 6	(Optional) switch(config)# ip name-server address1 [<i>address2... address6</i>] [use-vrf vrf-name]	<p>Defines up to six name servers. The address can be either an IPv4 address or an IPv6 address.</p> <p>You can optionally define a VRF that Cisco NX-OS uses to reach this name server if it cannot be reached in the VRF that you configured this name server under.</p>
Step 7	(Optional) switch(config)# ip domain-lookup	Enables DNS-based address translation. This feature is enabled by default.
Step 8	(Optional) switch(config)# show hosts	Displays information about DNS.
Step 9	(Optional) switch(config)# copy running-config startup-config	Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

Verifying the DNS Client Configuration

To display the DNS client configuration, perform the following task:

Command	Purpose
show hosts	Displays information about DNS.

Configuration Examples for DNS Clients

This example shows how to establish a domain list with several alternate domain names:

```
ip domain list csi.com
ip domain list telecomprog.edu
ip domain list merit.edu
```

This example shows how to configure the hostname-to-address mapping process and specify IP DNS-based translation. The example also configures the addresses of the name servers and the default domain name.

```
ip domain lookup
ip name-server 192.168.1.111 192.168.1.2
ip domain name cisco.com
```

Related Documents for DNS Clients

Related Topic	
DNS Client CLI commands	<i>Cisco Nexus 7000 Series NX-OS Unicast Routing Command Reference</i>
VDCs and VRFs	<i>Cisco Nexus 7000 Series NX-OS Virtual Device Context Configuration Guide, Release 5.x</i>

Standards for DNS Clients

No new or modified standards are supported by this feature, and support for existing standards has not been modified by this feature.

Feature History for DNS

This table includes only the updates for those releases that have resulted in additions or changes to the feature.

Feature Name	Release	Feature Information
DNS	4.0(1)	This feature was introduced.



CHAPTER 6

Configuring WCCPv2

This chapter contains the following sections:

- [Finding Feature Information, on page 81](#)
- [Information About WCCPv2, on page 81](#)
- [Prerequisites for WCCPv2, on page 88](#)
- [Guidelines and Limitations for WCCPv2, on page 88](#)
- [WCCPv2 Default Settings, on page 90](#)
- [Configuring WCCPv2, on page 90](#)
- [Verifying the WCCPv2 Configuration, on page 96](#)
- [Configuration Examples for WCCPv2, on page 96](#)
- [Related Documents for WCCPv2, on page 97](#)
- [Standards for the WCCPv2, on page 97](#)
- [Feature History for WCCPv2, on page 97](#)

Finding Feature Information

Your software release might not support all the features documented in this module. For the latest caveats and feature information, see the Bug Search Tool at <https://tools.cisco.com/bugsearch/> and the release notes for your software release. To find information about the features documented in this module, and to see a list of the releases in which each feature is supported, see the "New and Changed Information" chapter or the Feature History table in this chapter.

Information About WCCPv2

WCCPv2 Overview

WCCPv2 enables the Cisco NX-OS router to transparently redirect packets to cache engines. WCCPv2 does not interfere with normal router operations. Using WCCPv2, the router can redirect requests on configured interfaces to cache engines rather than to intended host sites. With WCCPv2, the router can balance traffic loads across a cluster of cache engines (cache cluster) and ensure fault-tolerant and fail-safe operation in the cluster. As you add or delete cache engines from a cache cluster, WCCPv2 dynamically redirects the packets to the currently available cache engines.

WCCPv2 accepts the traffic at the cache engine and establishes the connection with the traffic originator (the client). The cache engine acts as if it were the original destination server. If the requested object is not available on the cache engine, the cache engine establishes its own connection out to the original destination server to retrieve the object.

Until Release 8.1(2), WCCPv2 is supported only on the Layer3 or SVI interfaces, for Cisco Nexus 7000 Series Switches.

WCCPv2 communicates between routers and cache engines on UDP port 2048.

By allowing a cache cluster to connect to multiple routers, WCCPv2 provides redundancy and a distributed architecture for instances when a cache engine must connect to many interfaces. In addition, WCCPv2 allows you to keep all the cache engines in a single cluster, which avoids the unnecessary duplication of web pages across several clusters.

WCCPv2 Service Types

A service is a defined traffic type that the router redirects to a cache engine with the WCCPv2 protocol.


You can configure the router to run one of the following cache-related services:

- Well-known —The router and the cache engine know the traffic type, for example the web cache service on TCP port 80 for HTTP.
- Dynamic service—A service in which the cache engine describes the type of redirected traffic to the router.

WCCPv2 Service Groups

A service group is a subset of cache engines within a cluster and the routers connected to the cluster that are running the same service. The figure shows a service group within a cache cluster. The cache engines and the routers can be a part of multiple service groups.

Figure 21: WCCPv2 Cache Cluster and Service Group



You can configure a service group as open or closed. An open service group forwards traffic without redirection if there is no cache engine to redirect the traffic to. A closed service group drops traffic if there is no cache engine to redirect the traffic to.

The service group defines the traffic that is redirected to individual cache engines in that service group. The service group definition consists of the following:

- Service ID (0–255)
- Service Type
- Priority of the service group
- Protocol (TCP or UDP) of redirected traffic
- Service flags
- Up to eight TCP or UDP port numbers (either all source or all destination port numbers)

WCCPv2 Service Group Lists

WCCPv2 requires that each cache engine be aware of all the routers in the service group. You can configure a list of router addresses for each of the routers in the group on each cache engine.

The following sequence of events details how WCCPv2 configuration works:

1. You configure each cache engine with a list of routers.
2. Each cache engine announces its presence and generates a list of all routers with which it has established communications.

- The routers reply with their view (list) of cache engines in the group.

The cache engines and routers exchange control messages every 10 seconds by default.

WCCPv2 Designated Cache Engine

WCCPv2 designates one cache engine as the lead. If there is a group of cache engines, the one seen by all routers and the one that has the lowest IP address becomes the designated cache engine. The designated cache engine determines how traffic should be allocated across cache engines. The traffic assignment method is passed to the entire service group from the designated cache engine so that the routers of the group can redirect the packets and the cache engines of the group can manage their traffic load better.


Cisco NX-OS uses the mask method to assign traffic. The designated cache engine assigns the mask and value sets to the router in the WCCP Redirect Assignment message. The router matches these mask and value sets to the source IP address, destination IP address, source port, and destination port of each packet. The router redirects the packet to the cache engine if the packet matches an assigned mask and value set. If the packet does not match an assigned mask and value set, the router forwards the packet without any redirection.

WCCPv2 Redirection

You can use an IP access list as a redirect list to specify a subset of traffic to redirect with WCCPv2. You can apply this access list for ingress or egress traffic on an interface. The figure shows how redirection applies to ingress or egress traffic.

You can also exclude ingress traffic on an interface but allow egress redirection on that interface.

Figure 22: WCCPv2 Redirection



Supported Modules for WCCPv2 Redirection

The following tables show the supported modules in Cisco NX-OS for WCCPv2 redirection.

Redirect-In

Table 11: Supported Modules for WCCPv2 Redirect-In—Same Module Type

Ingress module	Egress Module	Module used to connect to WCCPv2 enabled device
M	M	M
F2	F2	F2

Ingress module	Egress Module	Module used to connect to WCCPv2 enabled device
F2e	F2e	F2e
F3	F3	F3

Table 12: Supported Modules for WCCPv2 Redirect-In—Mixed Module Type

Ingress module	Egress Module	Module used to connect to WCCPv2 enabled device
M	F2e	F2e
M2/M3	M2/M3	F3
M2/M3	F3	M2/M3
F3	M2/M3	M2/M3
M2/M3	F3	F3
F3	M2/M3	F3
F3	F3	M2/M3
F2e	F2e	F3
F2e	F3	F2e
F3	F2e	F2e
F3	F3	F2e
F3	F2e	F3
F2e	F3	F3

Redirect-Out

Table 13: Supported Modules for WCCPv2 Redirect-Out—Same Module Type

Ingress module	Egress Module	Module used to connect to WCCPv2 enabled device
M	M	M
F2	F2 ¹	F2 ¹
F2e	F2e ¹	F2e ¹
F3	F3	F3



Note ¹ Redirect-out and exclude-in are not supported on interface VLANs (SVIs).

Table 14: Supported Modules for WCCPv2 Redirect-Out—Mixed Module Type

Ingress module	Egress Module	Module used to connect to WCCPv2 enabled device
F2e	M	M
F2e	F2e	M
M	F2e	M
M2/M3	M2/M3	F3
F3 ²	M2/M3	M2/M3
F3	M2/M3	F3
F2e ³	F2e ⁴	F3
F2e ³	F3	F2e ⁵
F3 ³	F2e ⁴	F2e ⁵
F3 ³	F3	F2e ⁵
F3 ³	F2e ⁴	F3
F2e ³	F3	F3



Note ² Will not work if the F3 port is a FabricPath core port.

³ WCCP redirect-out will not work if the ingress traffic is on a FabricPath VLAN.

⁴ WCCP redirect-out is not supported on an F2e SVI.

⁵ WCCP exclude-in is not supported on an F2e SVI.

WCCPv2 Authentication

WCCPv2 can authenticate a device before it adds that device to the service group. Message Digest (MD5) authentication allows each WCCPv2 service group member to use a secret key to generate a keyed MD5 digest string that is part of the outgoing packet. At the receiving end, a keyed digest of an incoming packet is generated. If the MD5 digest within the incoming packet does not match the generated digest, WCCP ignores the packet.

WCCPv2 rejects packets in any of the following cases:

- The authentication schemes differ on the router and in the incoming packet.

- The MD5 digests differ on the router and in the incoming packet.

You must configure the same authentication on all members of a WCCPv2 service group.

WCCPv2 Redirection Method

WCCPv2 negotiates the packet redirection method between the router and the cache engine. Cisco NX-OS uses this traffic redirection method for all cache engines in a service group.

WCCPv2 redirects packets using Layer 2 Destination MAC rewrite method, where WCCPv2 replaces the destination MAC address of the packet with the MAC address of the cache engine that needs to handle the packet. The cache engine and the router must be adjacent to Layer 2.

You can also configure an access control list (ACL), called a redirect list, for a WCCPv2 service group. This ACL can either permit a packet to go through the WCCPv2 redirection process or deny the WCCP redirection and send the packet through the normal packet forwarding procedure.

WCCPv2 Packet Return Method

WCCPv2 filters packets to determine which redirected packets have been returned from the cache engine and which packets have not. WCCPv2 does not redirect the returned packets, because the cache engine has determined that these packets should not be cached. WCCPv2 returns packets that the cache engine does not service to the router that transmitted them.

A cache engine may return a packet for one of the following reasons:

- The cache engine is overloaded and cannot service the packets.
- The cache engine is filtering certain conditions that make caching packets counterproductive, for example, when IP authentication has been turned on.

WCCPv2 negotiates the packet return method between the router and the cache engine. Cisco NX-OS uses this traffic return method for all cache engines in a service group.

WCCPv2 returns packets using the Destination MAC rewrite method, where WCCPv2 replaces the destination MAC address of the packet with the MAC address of the router that originally redirected the packet. The cache engine and the router must be adjacent to Layer 2.

High Availability for WCCPv2

WCCPv2 supports stateful restarts and stateful switchovers. A stateful restart occurs when the WCCPv2 process fails and is restarted. A stateful switchover occurs when the active supervisor switches to the standby supervisor. Cisco NX-OS applies the running configuration after a switchover.

Virtualization Support for WCCPv2

WCCPv2 supports virtual routing and forwarding (VRF) instances. VRFs exist within virtual device contexts (VDCs). By default, Cisco NX-OS places you in the default VDC and default VRF unless you specifically configure another VDC and VRF.

WCCP redirection occurs within a VRF. You must configure the WCCP cache engine so that the forward and return traffic to and from the cache engine occurs from interfaces that are a part of the same VRF.

The VRF used for the WCCP on an interface should match the VRF configured on that interface.

If you change the VRF membership of an interface, Cisco NX-OS removes all layer 3 configuration, including WCCPv2.

For more information, see the *Cisco Nexus 7000 Series NX-OS Virtual Device Context Configuration Guide*.

WCCPv2 Error Handling for SPM Operations

The Service Policy Manager (SPM) supervisor component acts as a data path manager for the WCCP Manager. The WCCP manager is shielded from the underlying platform specifics by the SPM and is portable to platform variations. The WCCP manager has a set of SPM APIs to pass the configurations that are mapped and programmed in the hardware. These APIs can process and parse the application data that is implemented and maintained in one single handler.

The interface redirects that failed to be programmed by the SPM are stored until there is a service group configuration change through the CLI or an RA message. The WCCP manager retries programming policies that failed previously.

The WCCP manager sends policy updates to the SPM in intervals to program TCAM entries in the hardware. These policy updates can be triggered by the CLI or through RA (Redirect-Assign) messages. When the WCCP is notified of an SPM error, a syslog message appears.

Prerequisites for WCCPv2

WCCPv2 has the following prerequisites:

- You must globally enable the WCCPv2 feature.
- You can only configure WCCPv2 on Layer 3 or VLAN interfaces (see the *Cisco Nexus 7000 Series NX-OS Interfaces Configuration Guide*).

Guidelines and Limitations for WCCPv2

WCCPv2 has the following configuration guidelines and limitations:

- A WCCPv2 service group supports up to 32 routers and 32 cache engines.
- All cache engines in a cluster must include all routers that service the cluster in its configuration. If a cache engine within a cluster does not include one or more of the routers in its configuration, the service group detects the inconsistency and the cache engine is not allowed to operate within the service group.
- The cache engine cannot be on the same SVI with a redirect out statement.
- WCCPv2 works with IPv4 networks only.
- Any traffic that is coming from an M1-Series or M2-Series I/O module interface and going towards a Traffic Engineering (TE) Class-based Tunnel Selection (CBTS) tunnel will be dropped if you have configured the **ip wccp redirect exclude in** command on the inbound M1-Series or M2-Series I/O module interface or Switch Virtual Interface (SVI).

- WCCPv2 supports multiple service groups in the same direction (either inbound or outbound) on any Layer 3 interface, under the following conditions:
 - The access-list used must not have **deny ip any any** entry.
 - The access-list used for multiple service groups must not contain overlapping entries.

The following is an example of an overlapping entry:

```
ip access-list wccp_acl1
  permit tcp 10.0.0.0/8 10.0.0.0/8
ip access-list wccp_acl2
  permit tcp 10.10.10.1/32 10.10.10.10/32
```

- Cisco NX-OS removes all Layer 3 configuration on an interface when you change the VDC, interface VRF membership, port-channel membership, or the port mode to Layer 2.
- Cisco NX-OS does not support WCCPv2 on tunnel interfaces.
- WCCPv2 is supported on all types of FEX devices.
- WCCP requires the client, server, and WCCP client to be on separate interfaces. If you migrate a topology from a Cisco Catalyst 6500 Series switch deployment, it might not be supported.
- F2 Series, F2e Series, M1 Series, and M2 Series modules support WCCPv2. However, F2 and F2e Series modules do not support egress WCCPv2 on an SVI including “exclude in” on SVI. F1 Series modules do not support WCCPv2.
- WCCPv2 redirect-in and redirect-out is fully supported in Cisco NX-OS Release 6.2 in non-mixed module VDCs. WCCPv2 is also supported in mixed module VDC scenarios for most module combinations.
- For egress WCCPv2, traffic is not redirected when the ingress includes F2 series modules, and the next-hop is pointing to an SVI interface or subinterface of any module. If the egress WCCP policy is applied on a SVI or subinterface and if the packet ingresses on a F2 module, the same limitation applies.
- Beginning with Cisco NX-OS Release 5.2(4), policy-based routing and WCCPv2 are supported on the same interface. However, policy-based routing with statistics and WCCPv2 is supported on the same interface only if bank chaining is disabled.
- GRE redirection/return and hash assignment are not supported on a Cisco Nexus 7000 Series switch.
- Traffic might encounter a vPC loop and drop if you have Web Cache Control Protocol (WCCP) and vPC on your Cisco Nexus 7000 Series switch and the traffic migrates from a Cisco Nexus 65xx switch to your switch. Traffic that comes from a vPC member port and crosses a vPC peer-link is not permitted to egress any vPC member port. However, it can egress any other type of port, such a Layer 3 port or an orphan port. This behavior is expected.

If traffic drops after you configure WCCP and vPC on your Nexus 7000 Series switch and based on your design, you can perform one of the following tasks to avoid the vPC loop:

- Configure a Layer 2 trunk to carry the traffic in question.
 - Enable a peer gateway.
 - Shut down one of the member ports in the vPC.
- If you are familiar with the Cisco IOS CLI, be aware that the Cisco NX-OS commands for this feature might differ from the Cisco IOS commands that you would use.

- The following restrictions apply to the redirect-list, ACL:
 - Permit statements in the redirect ACL will consume more security TCAM entries compared to deny statements. Ensure the TCAM does not become oversubscribed.
 - The ACL must be an IPV4 simple ACL.
 - The protocol must be IP or TCP.
 - Only individual source or destination port numbers may be specified; port ranges cannot be specified.
 - The use of fragments or options is not permitted.

WCCPv2 Default Settings

Parameters	Default
Authentication	No authentication
WCCPv2	Disable

Configuring WCCPv2

To configure WCCPv2, perform these tasks in this chapter:

Procedure

-
- Step 1** Enable the WCCPv2 feature.
 - Step 2** Configure a WCCPv2 service group.
 - Step 3** Apply WCCPv2 redirection to an interface.
-

Enabling and Disabling WCCPv2

Before you begin

- Enable the WCCPv2 feature.
- Ensure you are in the correct VDC (or use the **switchto vdc** command)

Procedure

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.

	Command or Action	Purpose
Step 2	(config)# [no] feature wccp	Enables or disables the WCCPv2 feature in a VDC. Use the no form of the command to disable the feature.
Step 3	(Optional) switch(config)# copy running-config startup-config	Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

Configuring a WCCPv2 Service Group



Note You must enter the **ip wccp** command with all your required parameters. Any subsequent entry of the **ip wccp** command overwrites the earlier configuration.

Before you begin

- Enable the WCCPv2 feature.
- Ensure you are in the correct VDC (or use the **switchto vdc** command)

Procedure

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# ip wccp { <i>service-number</i> web-cache } [mode { open [redirect-list <i>acl-name</i>] closed service-list <i>acl-name</i> }] [password [0-7] <i>pwstring</i>]	Creates an open or closed mode service group. The service list identifies a named extended IP access list that defines the packets that will match the service. This list is required only when the service is defined as closed mode Optional parameters are as follows: <ul style="list-style-type: none"> • mode—Configures the service group in open or closed mode. On a service list, the mode controls the traffic type that the service group handles. The default is open. For closed mode, use this keyword to configure an IP access list to define the traffic type that matches this service. <p>Closed mode for dynamic service groups requires a service list ACL that specifies the protocol and port information that is used for the service group. If there are no members in the service group, packets matching the service-list ACL are dropped.</p>

	Command or Action	Purpose
		<ul style="list-style-type: none"> • password—Configures MD5 authentication for a service group. Use password 0 <i>pwstring</i> to store the password in clear text. Use password 7 <i>pwstring</i> to store the password in encrypted form. You can use the password 7 keywords for an already encrypted password. • redirect-list—Configures a global WCCPv2 redirection list for the service group to control the traffic that is redirected to the cache engine. • service-list—Configures an IP access list that defines the traffic type redirected by the service group. • The <i>service-number</i> range is from 1 to 255. The <i>acl-name</i> can be any case-sensitive, alphanumeric string up to 64 characters. The <i>pwstring</i> can be any case-sensitive, alphanumeric string up to eight characters.
Step 3	(Optional) switch(config)# copy running-config startup-config	Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

Applying WCCPv2 Redirection to an Interface

Procedure

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# interface ethernet <i>number</i>	Enters interface configuration mode.
Step 3	switch(config-if)# ip wccp {<i>service-number</i> redirect {in out} web-cache redirect {in out}}	<p>Applies the specified type of WCCPv2 redirection to the interface. The command examples show the following:</p> <ul style="list-style-type: none"> • WCCPv2 redirection applied on the ingress or egress traffic for this interface. • WCCPv2 redirection applied on the ingress or egress web cache traffic for this interface. • Ingress traffic excluded from WCCP redirection on this interface.

	Command or Action	Purpose
		Note <code>ip wccp web-cache redirect out</code> command is not supported in WCCP on BDI interface.
Step 4	<code>switch(config)# copy running-config startup-config</code>	(Optional) Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

Example

This example shows how to configure a router to redirect web-related packets without a destination of 19.20.2.1 to the web cache:

```
switch(config)# access-list 100
switch(config-acl)# deny ip any host 192.0.2.1
switch(config-acl)# permit ip any any
switch(config-acl)# exit
switch(config)# ip wccp web-cache redirect-list 100
switch(config)# interface ethernet 2/1
switch(config-if)# ip wccp web-cache redirect out
```

This example shows sample configuration for un-supported features:

```
switch# configure terminal
switch(config)# interface Bdi555
switch(config-if)# ip wccp redirect exclude in
This will remove all redirect-in on the interface. Proceed (y/n)? [no] y
ERROR: Exclude in not supported on BDI

switch(config-if)# ip wccp 62 redirect out
ERROR: Redirect out not supported on BDI
```

This example shows a running-configuration, followed by a verification command that displays the L3VNI-BDI configuration details. Replace the placeholders with relevant values for your setup. The example considers that interface 555 is configured for BDI.

```
switch (config)# show running-configuration interface bdi 555

!Command: show running-config wccp
!Time: Thu Sep 25 02:46:02 2017

version 8.2(1)
interface Bdi555
  description L3VNI-BDI
  no shutdown
  vrf member vrf5000
  no ip redirects
  ip forward
  ip pin sparse-mode
  ip wccp 61 redirect in
```

This example show running-configuration for WCCP configuration on BDI interface. Replace the placeholders with relevant values for your setup.

```
switch (config)# show running-configuration wccp

!Command: show running-config wccp
```

```

!Time: Thu Sep 25 02:46:02 2017

version 8.2(1)
feature wccp

vrf context vrf5000
 ip wccp web-cache
 ip wccp 61
 ip wccp 62

interface Bdi555
 vrf member vrf5000
 ip wccp 61 redirect in

```

Configuring WCCPv2 in a VRF



Note You must enter the **ip wccp** command with all your required parameters. Any subsequent entry of the **ip wccp** command overwrites the earlier configuration.

Before you begin

- Enable the WCCPv2 feature.
- Ensure you are in the correct VDC (or use the **switchto vdc** command)

Procedure

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# vrf context <i>vrf-name</i>	Enters VRF configuration mode. The <i>vrf-name</i> can be any case-sensitive, alphanumeric string up to 63 characters.
Step 3	switch(config)# ip wccp { <i>service-number</i> web-cache } [mode { open [redirect-list <i>acl-name</i>] closed service-list <i>acl-name</i> }] [password [0-7] <i>pwstring</i>]	<p>Creates an open or closed mode service group. The service list identifies a named extended IP access list that defines the packets that will match the service. This list is required only when the service is defined as closed mode</p> <p>Optional parameters are as follows:</p> <ul style="list-style-type: none"> • mode—Configures the service group in open or closed mode. On a service list, the mode controls the traffic type that the service group handles. The default is open. For closed mode, use this keyword to configure an IP access list to define the traffic type that matches this service. <p>Closed mode for dynamic service groups requires a service list ACL that specifies</p>

	Command or Action	Purpose
		<p>the protocol and port information that is used for the service group. If there are no members in the service group, packets matching the service-list ACL are dropped.</p> <ul style="list-style-type: none"> • password—Configures MD5 authentication for a service group. Use password 0 <i>pwstring</i> to store the password in clear text. Use password 7 <i>pwstring</i> to store the password in encrypted form. You can use the password 7 keywords for an already encrypted password. • redirect-list—Configures a global WCCPv2 redirection list for the service group to control the traffic that is redirected to the cache engine. • service-list—Configures an IP access list that defines the traffic type redirected by the service group. • The <i>service-number</i> range is from 1 to 255. The <i>acl-name</i> can be any case-sensitive, alphanumeric string up to 64 characters. The <i>pwstring</i> can be any case-sensitive, alphanumeric string up to eight characters.
Step 4	(Optional) switch(config-vrf)# show ip wccp [vrf <i>vrf-name</i>]	Displays information about WCCPv2. The <i>vrf-name</i> can be any case-sensitive, alphanumeric string up to 64 characters.
Step 5	(Optional) switch(config)# copy running-config startup-config	Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

Example

This example shows how to configure WCCPv2 in VRF Red on interface Ethernet 2/1:

```
switch# configure terminal
switch(config)# vrf context Red
switch(config-vrf)# ip wccp web-cache password Test1 redirect-list httpTest
switch(config-vrf)# interface ethernet 2/1
switch(config-if)# vrf member Red
switch(config-if)# ip wccp web-cache redirect out
```

Verifying the WCCPv2 Configuration

Use one of the following commands to verify the configuration:

Command	Purpose
<code>show ip wccp [vrf vrf-name] [service-number web-cache]</code>	Displays the WCCPv2 status for all groups or one group in a VRF.
<code>show ip interface [ethernet-number]</code>	Displays the WCCPv2 interface information.
<code>show ip wccp [service-number web-cache]</code>	Displays the WCCPv2 service group status.
<code>show ip wccp [service-number web-cache] detail</code>	Displays the clients in a WCCPv2 service group.
<code>show ip wccp [service-number web-cache] mask</code>	Displays the WCCPv2 mask assignment.
<code>show ip wccp [service-number web-cache] service</code>	Displays the WCCPv2 service group definition.
<code>show ip wccp [service-number web-cache] view</code>	Displays the WCCPv2 group membership.

Configuration Examples for WCCPv2

This example shows how to configure WCCPv2 authentication on router redirect web-related packets without a destination of 192.0.2.1 to the web cache:

```
access-list 100
 deny ip any host 192.0.2.1
 permit ip any any
 feature wccp
 ip wccp web-cache password 0 Test1 redirect-list 100
 interface ethernet 1/2
 ip wccp web-cache redirect out
 no shutdown
```

This example shows the sample output when WCCP is configuration in a VRF.

```
switch(config)# show ip wccp vrf vrf5000

VRF vrf5000 WCCP information:
  Router information:
    Router Identifier:          50.50.50.1
    Protocol Version:          2.0
  Service Identifier: web-cache
    Number of Service Group Clients: 1
    Number of Service Group Routers: 1
    Service mode:               Open
    Service Access-list:        -none-
    Redirect Access-list:       -none-
  Service Identifier: 61
    Number of Service Group Clients: 1
    Number of Service Group Routers: 1
    Service mode:               Open
    Service Access-list:        -none-
    Redirect Access-list:       -none-
  Service Identifier: 62
```

```

Number of Service Group Clients:      1
Number of Service Group Routers:     1
Service mode:                        Open
Service Access-list:                 -none-
Redirect Access-list:                 -none-

```

The following example shows a verification command to display the kind of service for WCCP.

```

switch(config)# show ip wccp vrf vrf5000 61 service
WCCP service information definition:
  Type:           Dynamic
  Id:             61
  Priority:       34
  Protocol:      6
  Options:       0x00000501
  -----
  Mask/Value sets: 1
  Value elements : 16
  Ports:          -none-

```

The following example shows a verification command to display cache engine information, after the connection with the cache engine is established

```

switch(config)# show ip wccp vrf vrf5000 61 view

WCCP Router Informed of:
50.50.50.1

WCCP Cache Engines Visible:
10.10.10.3

WCCP Cache Engines Not Visible:
-none-

```

Related Documents for WCCPv2

Related Topic	Document Title
WCCPv2 CLI commands	<i>Cisco Nexus 7000 Series NX-OS Unicast Routing Command Reference</i>
IP ACLs	<i>Cisco Nexus 7000 Series NX-OS Security Configuration Guide, Release 6.x</i>

Standards for the WCCPv2

No new or modified standards are supported by this feature, and support for existing standards has not been modified by this feature.

Feature History for WCCPv2

This table includes only the updates for those releases that have resulted in additions or changes to the feature.

Feature Name	Release	Feature Information
WCCPv2 on BDI	8.2(1)	Added support on BDI interface.
WCCPv2 Redirection	7.3(0)DX(1)	Added support for M3 module.
WCCPv2	5.2(4)	Added support for policy-based routing and WCCPv2 on the same interface if bank chaining is disabled.
WCCPv2 Error Handling for SPM Operations	5.1(1)	This feature was added.
WCCPv2	4.2(1)	This feature was introduced.



PART II

Routing

- [Configuring OSPFv2, on page 101](#)
- [Configuring OSPFv3, on page 145](#)
- [Configuring EIGRP, on page 193](#)
- [Configuring IS-IS, on page 227](#)
- [Configuring Basic BGP, on page 253](#)
- [Configuring Advanced BGP, on page 283](#)
- [Configuring RIP, on page 337](#)
- [Configuring Static Routing, on page 355](#)
- [Configuring the Interoperability of Modules for Unicast Routing, on page 367](#)
- [Configuring Layer 3 Virtualization, on page 371](#)
- [Managing the Unicast RIB and FIB, on page 383](#)
- [Configuring Route Policy Manager, on page 397](#)
- [Configuring Policy-Based Routing, on page 419](#)



CHAPTER 7

Configuring OSPFv2

This chapter contains the following sections:

- [Finding Feature Information](#), on page 101
- [Information About OSPFv2](#), on page 101
- [Prerequisites for OSPFv2](#), on page 112
- [Guidelines and Limitations for OSPFv2](#), on page 112
- [Default Settings for OSPFv2](#), on page 114
- [Configuring Basic OSPFv2](#), on page 115
- [Configuring Advanced OSPFv2](#), on page 124
- [Verifying the OSPFv2 Configuration](#), on page 141
- [Monitoring OSPFv2](#), on page 142
- [Configuration Examples for OSPFv2](#), on page 142
- [Related Documents for OSPFv2](#), on page 143
- [Feature History for OSPFv2](#), on page 143

Finding Feature Information

Your software release might not support all the features documented in this module. For the latest caveats and feature information, see the Bug Search Tool at <https://tools.cisco.com/bugsearch/> and the release notes for your software release. To find information about the features documented in this module, and to see a list of the releases in which each feature is supported, see the "New and Changed Information" chapter or the Feature History table in this chapter.

Information About OSPFv2

OSPFv2 is an IETF link-state protocol for IPv4 networks. An OSPFv2 router sends a special message, called a hello packet, out each OSPF-enabled interface to discover other OSPFv2 neighbor routers. Once a neighbor is discovered, the two routers compare information in the to determine if the routers have compatible configurations. The neighbor routers try to establish , which means that the routers synchronize their link-state databases to ensure that they have identical OSPFv2 routing information. Adjacent routers share (LSAs) that include information about the operational state of each link, the cost of the link, and any other neighbor information. The routers then flood these received LSAs out every OSPF-enabled interface so that all OSPFv2 routers eventually have identical link-state databases. When all OSPFv2 routers have identical link-state

databases, the network is converged. Each router then uses Dijkstra's Shortest Path First (SPF) algorithm to build its route table.

You can divide OSPFv2 networks into areas. Routers send most LSAs only within one area, which reduces the CPU and memory requirements for an OSPF-enabled router.

OSPFv2 supports IPv4, while OSPFv3 supports IPv6. For more information, see the "Configuring OSPFv3" chapter.



Note OSPFv2 on Cisco NX-OS supports RFC 2328. This RFC introduced a different method to calculate route summary costs which is not compatible with the calculation used by RFC1583. RFC 2328 also introduced different selection criteria for AS-external paths. It is important to ensure that all routers support the same RFC. Use the **rfc1583compatibility** command if your network includes routers that are only compliant with RFC1583. The default supported RFC standard for OSPFv2 may be different for Cisco NX-OS and Cisco IOS. You must make adjustments to set the values identically. For more information, see the "OSPF RFC Compatibility Mode Example" section.

Hello Packet

OSPFv2 routers periodically send Hello packets on every OSPF-enabled interface. The hello interval determines how frequently the router sends these Hello packets and is configured per interface. OSPFv2 uses Hello packets for the following tasks:

- Neighbor discovery
- Keepalives
- Bidirectional communications
- Designated router election

The Hello packet contains information about the originating OSPFv2 interface and router, including the assigned OSPFv2 cost of the link, the , and optional capabilities of the originating router. An OSPFv2 interface that receives these Hello packets determines if the settings are compatible with the receiving interface settings. Compatible interfaces are considered neighbors and are added to the neighbor table.

Hello packets also include a list of router IDs for the routers that the originating interface has communicated with. If the receiving interface sees its own router ID in this list, then bidirectional communication has been established between the two interfaces.

OSPFv2 uses Hello packets as a keepalive message to determine if a neighbor is still communicating. If a router does not receive a Hello packet by the configured (usually a multiple of the hello interval), then the neighbor is removed from the local neighbor table.

Neighbors

An OSPFv2 interface must have a compatible configuration with a remote interface before the two can be considered neighbors. The two OSPFv2 interfaces must match the following criteria:

- Hello interval
- Dead interval

- Area ID
- Authentication
- Optional capabilities

If there is a match, the following information is entered into the neighbor table:

- Neighbor ID—The router ID of the neighbor.
- Priority—Priority of the neighbor. The priority is used for designated router election.
- State—Indication of whether the neighbor has just been heard from, is in the process of setting up bidirectional communications, is sharing the link-state information, or has achieved full adjacency.
- Dead time—Indication of the time since the last Hello packet was received from this neighbor.
- IP Address—The IP address of the neighbor.
- Designated Router—Indication of whether the neighbor has been declared as the designated router or as the backup designated router.
- Local interface—The local interface that received the Hello packet for this neighbor.

Adjacency

Not all neighbors establish adjacency. Depending on the network type and designated router establishment, some neighbors become fully adjacent and share LSAs with all their neighbors, while other neighbors do not.

Adjacency is established using Database Description packets, Link State Request packets, and Link State Update packets in OSPF. The Database Description packet includes just the LSA headers from the link-state database of the neighbor. The local router compares these headers with its own link-state database and determines which LSAs are new or updated. The local router sends a Link State Request packet for each LSA that it needs new or updated information on. The neighbor responds with a Link State Update packet. This exchange continues until both routers have the same link-state information.

Designated Routers

Networks with multiple routers present a unique situation for OSPF. If every router floods the network with LSAs, the same link-state information is sent from multiple sources. Depending on the type of network, OSPFv2 might use a single router, the (DR), to control the LSA floods and represent the network to the rest of the OSPFv2 area. If the DR fails, OSPFv2 selects a (BDR). If the DR fails, OSPFv2 uses the BDR.

Network types are as follows:

- Point-to-point—A network that exists only between two routers. All neighbors on a point-to-point network establish adjacency and there is no DR.
- Broadcast—A network with multiple routers that can communicate over a shared medium that allows broadcast traffic, such as Ethernet. OSPFv2 routers establish a DR and BDR that controls LSA flooding on the network. OSPFv2 uses the well-known IPv4 multicast addresses 224.0.0.5 and a MAC address of 0100.5300.0005 to communicate with neighbors.

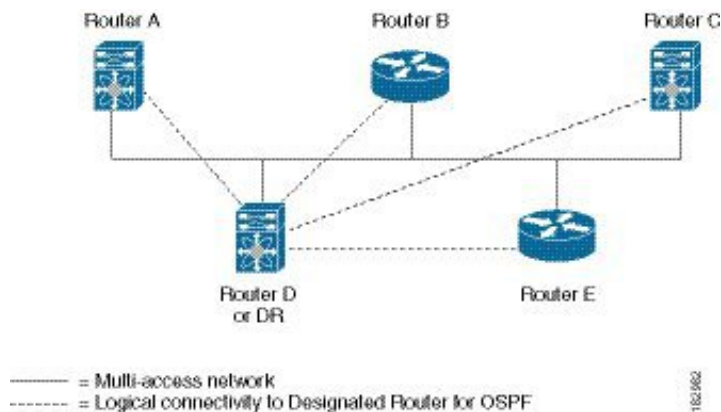
The DR and BDR are selected based on the information in the Hello packet. When an interface sends a Hello packet, it sets the priority field and the DR and BDR field if it knows who the DR and BDR are. The routers

follow an election procedure based on which routers declare themselves in the DR and BDR fields and the priority field in the Hello packet. As a final tie breaker, OSPFv2 chooses the highest router IDs as the DR and BDR.

All other routers establish adjacency with the DR and the BDR and use the IPv4 multicast address 224.0.0.6 to send LSA updates to the DR and BDR. Figure 3-1 shows this adjacency relationship between all routers and the DR.

DRs are based on a router interface. A router might be the DR for one network and not for another network on a different interface.

Figure 23: DR in Multi-Access Network




Areas

You can limit the CPU and memory requirements that OSPFv2 puts on the routers by dividing an OSPFv2 network into . An area is a logical division of routers and links within an OSPFv2 domain that creates separate subdomains. LSA flooding is contained within an area, and the link-state database is limited to links within the area. You can assign an area ID to the interfaces within the defined area. The Area ID is a 32-bit value that you can enter as a number or in dotted decimal notation, such as 10.2.3.1.

Cisco NX-OS always displays the area in dotted decimal notation.

If you define more than one area in an OSPFv2 network, you must also define the backbone area, which has the reserved area ID of 0. If you have more than one area, then one or more routers become (ABRs). An ABR connects to both the backbone area and at least one other defined area.

Figure 24: OSPFv2 Areas



The ABR has a separate link-state database for each area to which it connects. The ABR sends Network Summary (type 3) from one connected area to the backbone area. The backbone area sends summarized information about one area to another area. In the OSPFv2 Areas Figure, Area 0 sends summarized information about Area 5 to Area 3.

OSPFv2 defines one other router type: the autonomous system boundary router (ASBR). This router connects an OSPFv2 area to another autonomous system. An autonomous system is a network controlled by a single technical administration entity. OSPFv2 can redistribute its routing information into another autonomous system or receive redistributed routes from another autonomous system.

Link-State Advertisements

Link-State Advertisements Types

OSPFv2 uses link-state advertisements (LSAs) to build its routing table.

Names	Description
Router LSA	LSA sent by every router. This LSA includes the state and the cost of all links and a list of all OSPFv2 neighbors on the link. Router LSAs trigger an SPF recalculation. Router LSAs are flooded to local OSPFv2 area.
Network LSA	LSA sent by the DR. This LSA lists all routers in the multi-access network. Network LSAs trigger an SPF recalculation.
Network Summary LSA	LSA sent by the area border router to an external area for each destination in the local area. This LSA includes the link cost from the area border router to the local destination.
ASBR Summary LSA	LSA sent by the area border router to an external area. This LSA advertises the link cost to the ASBR only.

Names	Description
AS External LSA	LSA generated by the ASBR. This LSA includes the link cost to an external autonomous system destination. AS External LSAs are flooded throughout the autonomous system.
NSSA External LSA	LSA generated by the ASBR within a not-so-stubby area (NSSA). This LSA includes the link cost to an external autonomous system destination. NSSA External LSAs are flooded only within the local NSSA.
Opaque LSAs	LSA used to extend OSPF.

Link Cost

Each OSPFv2 interface is assigned a . The cost is an arbitrary number. By default, Cisco NX-OS assigns a cost that is the configured reference bandwidth divided by the interface bandwidth. By default, the reference bandwidth is 40 Gb/s. The link cost is carried in the LSA updates for each link.

Flooding and LSA Group Pacing

When an OSPFv2 router receives an LSA, it forwards that LSA out every OSPF-enabled interface, flooding the OSPFv2 area with this information. This LSA flooding guarantees that all routers in the network have identical routing information. LSA flooding depends on the OSPFv2 area configuration. The LSAs are flooded based on the (every 30 minutes by default). Each LSA has its own link-state refresh time.

You can control the flooding rate of LSA updates in your network by using the LSA group pacing feature. LSA group pacing can reduce high CPU or buffer utilization. This feature groups LSAs with similar link-state refresh times to allow OSPFv2 to pack multiple LSAs into an OSPFv2 Update message.

By default, LSAs with link-state refresh times within 10 seconds of each other are grouped together. You should lower this value for large link-state databases or raise it for smaller databases to optimize the OSPFv2 load on your network.

Link-State Database

Each router maintains a link-state database for the OSPFv2 network. This database contains all the collected LSAs, and includes information on all the routes through the network. OSPFv2 uses this information to calculate the best path to each destination and populates the routing table with these best paths.

LSAs are removed from the link-state database if no LSA update has been received within a set interval, called the MaxAge. Routers flood a repeat of the LSA every 30 minutes to prevent accurate link-state information from being aged out. Cisco NX-OS supports the LSA grouping feature to prevent all LSAs from refreshing at the same time.

Opaque LSAs

Opaque LSAs allow you to extend OSPF functionality. Opaque LSAs consist of a standard LSA header followed by application-specific information. This information might be used by OSPFv2 or by other applications. OSPFv2 uses Opaque LSAs to support OSPFv2 Graceful Restart capability. Three Opaque LSA types are defined as follows:

- LSA type 9—Flooded to the local network.

- LSA type 10—Flooded to the local area.
- LSA type 11—Flooded to the local autonomous system.

OSPFv2 and the Unicast RIB

OSPFv2 runs the Dijkstra shortest path first algorithm on the link-state database. This algorithm selects the best path to each destination based on the sum of all the link costs for each link in the path. The resultant shortest path for each destination is then put in the OSPFv2 route table. When the OSPFv2 network is converged, this route table feeds into the unicast RIB. OSPFv2 communicates with the unicast RIB to do the following:

- Add or remove routes
- Handle route redistribution from other protocols
- Provide convergence updates to remove stale OSPFv2 routes and for stub router advertisements

OSPFv2 also runs a modified Dijkstra algorithm for fast recalculation for summary and external (type 3, 4, 5, and 7) LSA changes.

Authentication

You can configure authentication on OSPFv2 messages to prevent unauthorized or invalid routing updates in your network. Cisco NX-OS supports two authentication methods:

- Simple password authentication
- MD5 authentication digest

You can configure the OSPFv2 authentication for an OSPFv2 area or per interface.

Simple Password Authentication

Simple password authentication uses a simple clear-text password that is sent as part of the OSPFv2 message. The receiving OSPFv2 router must be configured with the same clear-text password to accept the OSPFv2 message as a valid route update. Because the password is in clear text, anyone who can watch traffic on the network can learn the password.

MD5 Authentication

You should use MD5 authentication to authenticate OSPFv2 messages. You configure a password that is shared at the local router and all remote OSPFv2 neighbors. For each OSPFv2 message, Cisco NX-OS creates an MD5 one-way message digest based on the message itself and the encrypted password. The interface sends this digest with the OSPFv2 message. The receiving OSPFv2 neighbor validates the digest using the same encrypted password. If the message has not changed, the digest calculation is identical and the OSPFv2 message is considered valid.

MD5 authentication includes a sequence number with each OSPFv2 message to ensure that no message is replayed in the network.

Advanced Features for OSPFv2

Cisco NX-OS supports advanced OSPFv2 features that enhance the usability and scalability of OSPFv2 in the network.


Stub Area

You can limit the amount of external routing information that floods an area by making it a stub area. A stub area is an area that does not allow AS External (type 5) LSAs. These LSAs are usually flooded throughout the local autonomous system to propagate external route information. Stub areas have the following requirements:

- All routers in the stub area are stub routers.
- No ASBR routers exist in the stub area.
- You cannot configure virtual links in the stub area.

The following figure shows an example of an OSPFv2 autonomous system where all routers in area 0.0.0.10 have to go through the ABR to reach external autonomous systems. Area 0.0.0.10 can be configured as a stub area.

Figure 25: Stub Area



Stub areas use a default route for all traffic that needs to go through the backbone area to the external autonomous system. The default route is 0.0.0.0 for IPv4.

Not-So-Stubby Area

A Not-so-Stubby Area (NSSA) is similar to a stub area, except that an NSSA allows you to import autonomous system external routes within an NSSA using redistribution. The NSSA ASBR redistributes these routes and generates NSSA External (type 7) LSAs that it floods throughout the NSSA. You can optionally configure the ABR that connects the NSSA to other areas to translate this NSSA External LSA to AS External (type 5) LSAs. The ABR then floods these AS External LSAs throughout the OSPFv2 autonomous system. Summarization and filtering are supported during the translation.

You can, for example, use NSSA to simplify administration if you are connecting a central site using OSPFv2 to a remote site that is using a different routing protocol. Before NSSA, the connection between the corporate site border router and a remote router could not be run as an OSPFv2 stub area because routes for the remote site could not be redistributed into a stub area. With NSSA, you can extend OSPFv2 to cover the remote connection by defining the area between the corporate router and remote router as an NSSA.

The backbone Area 0 cannot be an NSSA.



Note OSPF became compliant with RFC 3101 section 2.5(3). When an Area Border Router attached to a Not-so-Stubby Area receives a default route LSA with P-bit clear, it should be ignored. OSPF had been previously adding the default route under these conditions.


If you have already designed your networks with RFC non-compliant behavior and expect a default route to be added on NSSA ABR, you will see a change in behavior when you upgrade.

If you decide to continue with the old behavior, you have the option to enable it with the **default-route nssa-abr pbit-clear** command.

Virtual Links

Virtual links allow you to connect an OSPFv2 area ABR to a backbone area ABR when a direct physical connection is not available. The figure shows a virtual link that connects Area 3 to the backbone area through Area 5.

Figure 26: Virtual Links



You can also use virtual links to temporarily recover from a partitioned area, which occurs when a link within the area fails, isolating part of the area from reaching the designated ABR to the backbone area.

Route Redistribution

OSPFv2 can learn routes from other routing protocols by using route redistribution. You configure OSPFv2 to assign a link cost for these redistributed routes or a default link cost for all redistributed routes.

Route redistribution uses route maps to control which external routes are redistributed. You must configure a route map with the redistribution to control which routes are passed into OSPFv2. A route map allows you to filter routes based on attributes such as the destination, origination protocol, route type, route tag, and so on. You can use route maps to modify parameters in the AS External (type 5) and NSSA External (type 7) LSAs before these external routes are advertised in the local OSPFv2 autonomous system.

OSPFv2 sets the type-5 LSA's forwarding address as described below:

- If the next-hop for the route is an attached-route then the forwarding address is the next-hop address for that route.
- If the next-hop for the route is a recursive route and next-hop's next-hop is an attached route then the forwarding address is the next-hop's next-hop address.

Route Summarization

Because OSPFv2 shares all learned routes with every OSPF-enabled router, you might want to use route summarization to reduce the number of unique routes that are flooded to every OSPF-enabled router. Route summarization simplifies route tables by replacing more-specific addresses with an address that represents all the specific addresses. For example, you can replace 10.1.1.0/24, 10.1.2.0/24, and 10.1.3.0/24 with one summary address, 10.1.0.0/16.

Typically, you would summarize at the boundaries of area border routers (ABRs). Although you could configure summarization between any two areas, it is better to summarize in the direction of the backbone so that the backbone receives all the aggregate addresses and injects them, already summarized, into other areas. The two types of summarization are as follows

- Inter-area route summarization
- External route summarization

You configure inter-area route summarization on ABRs, summarizing routes between areas in the autonomous system. To take advantage of summarization, you should assign network numbers in areas in a contiguous way to be able to lump these addresses into one range.

External route summarization is specific to external routes that are injected into OSPFv2 using route redistribution. You should make sure that external ranges that are being summarized are contiguous. Summarizing overlapping ranges from two different routers could cause packets to be sent to the wrong destination. Configure external route summarization on ASBRs that are redistributing routes into OSPF.

When you configure a summary address, Cisco NX-OS automatically configures a discard route for the summary address to prevent routing black holes and route loops.

High Availability and Graceful Restart

Cisco NX-OS provides a multilevel high-availability architecture. OSPFv2 supports stateful restart, which is also referred to as non-stop routing (NSR). If OSPFv2 experiences problems, it attempts to restart from its previous run-time state. The neighbors do not register any neighbor event in this case. If the first restart is not successful and another problem occurs, OSPFv2 attempts a graceful restart.

A graceful restart, or nonstop forwarding (NSF), allows OSPFv2 to remain in the data forwarding path through a process restart. When OSPFv2 needs to perform a graceful restart, it sends a link-local opaque (type 9) LSA, called a grace LSA. This restarting OSPFv2 platform is called NSF capable.

The grace LSA includes a grace period, which is a specified time that the neighbor OSPFv2 interfaces hold onto the LSAs from the restarting OSPFv2 interface. (Typically, OSPFv2 tears down the adjacency and discards all LSAs from a down or restarting OSPFv2 interface.) The participating neighbors, which are called NSF helpers, keep all LSAs that originate from the restarting OSPFv2 interface as if the interface was still adjacent.

When the restarting OSPFv2 interface is operational again, it rediscovers its neighbors, establishes adjacency, and starts sending its LSA updates again. At this point, the NSF helpers recognize that the graceful restart has finished.

Stateful restart is used in the following scenarios:

- First recovery attempt after the process experiences problems
- ISSU
- User-initiated switchover using the **system switchover** command
- Active supervisor reload using the **reload module active-sup** command

Graceful restart is used in the following scenarios:

- Second recovery attempt after the process experiences problems within a 4-minute interval
- Manual restart of the process using the **restart ospf** command
- Active supervisor removal



Note The Cisco Nexus 7000 series devices support the Internet Engineering Task Force (IETF) version only. As a result, NSF IETF must be explicitly configured under the routing protocols in the Virtual Switching System (VSS). Use the **nsf ietf** command in router configuration mode for NSF IETF configuration. No additional configuration is required on the Cisco Nexus 7000 pairs because they run NSF IETF graceful-restart by default. However, each neighbor device that will become Layer 3 adjacent must have NSF configured and the same mode of NSF must be enabled to successfully operate a graceful failover.

OSPFv2 Stub Router Advertisements

You can configure an OSPFv2 interface to act as a stub router using the OSPFv2 Stub Router Advertisements feature. Use this feature when you want to limit the OSPFv2 traffic through this router, such as when you want to introduce a new router to the network in a controlled manner or limit the load on a router that is already overloaded. You might also want to use this feature for various administrative or traffic engineering reasons.

OSPFv2 stub router advertisements do not remove the OSPFv2 router from the network topology, but they do prevent other OSPFv2 routers from using this router to route traffic to other parts of the network. Only the traffic that is destined for this router or directly connected to this router is sent.

OSPFv2 stub router advertisements mark all stub links (directly connected to the local router) to the cost of the local OSPFv2 interface. All remote links are marked with the maximum cost (0xFFFF).

Multiple OSPFv2 Instances

Cisco NX-OS supports multiple instances of the OSPFv2 protocol that run on the same node. You cannot configure multiple instances over the same interface. By default, every instance uses the same system router ID. You must manually configure the router ID for each instance if the instances are in the same OSPFv2 autonomous system.

SPF Optimization

Cisco NX-OS optimizes the SPF algorithm in the following ways:

- Partial SPF for Network (type 2) LSAs, Network Summary (type 3) LSAs, and AS External (type 5) LSAs—When there is a change on any of these LSAs, Cisco NX-OS performs a faster partial calculation rather than running the whole SPF calculation.

- SPF timers—You can configure different timers for controlling SPF calculations. These timers include exponential backoff for subsequent SPF calculations. The exponential backoff limits the CPU load of multiple SPF calculations.

BFD

This feature supports bidirectional forwarding detection (BFD). BFD is a detection protocol that provides fast forwarding-path failure detection times. BFD provides subsecond failure detection between two adjacent devices and can be less CPU-intensive than protocol hello messages because some of the BFD load can be distributed onto the data plane on supported modules. See the *Cisco Nexus 7000 Series NX-OS Interfaces Configuration Guide* for more information.

Virtualization Support for OSPFv2

OSPFv2 supports virtual routing and forwarding (VRF) instances. VRFs exist within virtual device contexts (VDCs). By default, Cisco NX-OS places you in the default VDC and default VRF unless you specifically configure another VDC and VRF.

Cisco NX-OS Release 6.1 or later supports more than four process instances for OSPFv2 per VDC. However, only the first four configured OSPFv2 instances are supported with MPLS LDP and MPLS TE. Each OSPFv2 instance can support multiple VRFs, up to the system limit. For more information, see the *Cisco Nexus 7000 Series NX-OS Virtual Device Context Configuration Guide* and the *Cisco Nexus 7000 Series NX-OS Verified Scalability Guide*.

Prerequisites for OSPFv2

OSPFv2 has the following prerequisites:

- You must be familiar with routing fundamentals to configure OSPF.
- You are logged on to the switch.
- You have configured at least one interface for IPv4 that can communicate with a remote OSPFv2 neighbor.
- You have completed the OSPFv2 network strategy and planning for your network. For example, you must decide whether multiple areas are required.
- You have enabled the OSPF feature.
- You have installed the appropriate license and entered the desired VDC (see the *Cisco Nexus 7000 Series NX-OS Virtual Device Context Configuration Guide* for configuration information and the Cisco NX-OS Licensing Guide for licensing information) if you are configuring VDCs.

Guidelines and Limitations for OSPFv2

OSPFv2 has the following configuration guidelines and limitations:

- CE devices install type 3 LSAs with DN-bit or Type 5 LSAs with DN-bit and VPN Route TAG in the RIB (non-default VRF). This behaviour is applicable prior to Cisco NX-OS Release 8.3(2).

- The default-information originate command must be configured so that the MPLS default route is advertised to the CE-VRF. When using default-information originate command, the DN-bit in type 3 5 LSAs options and Route TAGs in Type 5 LSAs are not set for the default route only.
- The Cisco Nexus 7000 supports the Internet Engineering Task Force (IETF) version only. As a result, NSF IETF must be explicitly configured under the routing protocols in the Virtual Switching System (VSS). No additional configuration is required on the Cisco Nexus 7000 pairs because they run NSF IETF graceful-restart by default. However, each neighbor device that will become Layer 3 adjacent must have NSF configured and the same mode of NSF must be enabled to successfully operate a graceful failover.
- Cisco NX-OS displays areas in dotted decimal notation regardless of whether you enter the area in decimal or dotted decimal notation.
- All OSPFv2 routers must operate in the same RFC compatibility mode. OSPFv2 for Cisco NX-OS complies with RFC 2328. Use the **rfc1583compatibility** command in router configuration mode if your network includes routers that support only RFC 1583.
- In scaled scenarios, when the number of interfaces and link-state advertisements in an OSPF process is large, the snmp-walk on OSPF MIB objects is expected to time out with a small-values timeout at the SNMP agent. If you observe a timeout on the querying SNMP agent while polling OSPF MIB objects, increase the timeout value on the polling SNMP agent.
- MTU configured at interface level works in either the data plane or in the control plane but not at both planes at the same time.

When you configure MTU with a size lower than the supported size in data and control planes a few features that have minimum MTU requirements may not work in both the planes.

For example, MPLS VPN is supported in the data plane since this plane supports the MTU of 1500 bytes that the MPLS VPN requires. But control plane does not support MPLS VPN because this plane cannot handle the 1500-byte packets.

To make the configured MTU work in control plane for MPLS VPN, you need to manually configure the OSPF packet size (by using the **packet-size** *size* command) so that OSPF works on the control plane. This is applicable from Cisco NX-OS Release 8.3(2) onwards.

The **packet-size** *size* command is supported on the Ethernet, SVI, and GRE tunnel interfaces.

- Cisco NX-OS Release 6.1 or later supports more than four process instances for OSPFv2 per VDC. However, only the first four configured OSPFv2 instances are supported with MPLS LDP and MPLS TE.
- The **default-information-originate always** command advertises the OSPF default route from Cisco NX-OS Release 7.3(5)D1(1) and later releases and from Cisco NX-OS Release 8.0(1) and later releases in 8.x release train.
- The following guidelines and limitations apply to the administrative distance feature, which is supported beginning with Cisco NX-OS Release 6.1:
 - When an OSPF route has two or more equal cost paths, configuring the administrative distance is non-deterministic for the **match ip route-source** command.
 - Configuring the administrative distance is supported only for the **match route-type**, **match ip address prefix-list**, and **match ip route-source prefix-list** commands. The other match statements are ignored.

- There is no preference among the **match route-type**, **match ip address**, and **match ip route-source** commands for setting the administrative distance of OSPF routes. In this way, the behavior of the table map for setting the administrative distance in Cisco NX-OS OSPF is different from that in Cisco IOS OSPF.
- The discard route is always assigned an administrative distance of 220. No configuration in the table map applies to OSPF discard routes.
- In Cisco NX-OS Release 6.2(6a) and later releases, you can filter next-hop paths for an OSPF route to prevent the path from being added to the RIB. Before Cisco NX-OS Release 6.2(6a), filtering on a specific path was ignored and the entire route was not added to the RIB.



Note If you are familiar with the Cisco IOS CLI, be aware that the Cisco NX-OS commands for this feature might differ from the Cisco IOS commands that you would use.

Default Settings for OSPFv2

Table 15: Default OSPFv2 Parameters

Parameters	Default
Administrative distance	110
Hello interval	10 seconds
Dead interval	40 seconds
Discard routes	Enabled
Graceful restart grace period	60 seconds
OSPFv2 feature	Disabled
Stub router advertisement announce time	600 seconds
Reference bandwidth for link cost calculation	40 Gb/s
LSA minimal arrival time	1000 milliseconds
LSA group pacing	10 seconds
SPF calculation initial delay time	200 milliseconds
SPF minimum hold time	5000 milliseconds
SPF calculation initial delay time	1000 milliseconds

Configuring Basic OSPFv2

Enabling OSPFv2

You must enable the OSPFv2 feature before you can configure OSPFv2.

Before you begin

Ensure that you are in the correct VDC (or use the **switchto vdc** command).

Procedure

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# [no] feature ospf	Enables the OSPFv2 feature. Note Use the no form of this command to disable the OSPFv2 feature and remove all associated configuration.
Step 3	(Optional) switch(config)# show feature	Displays enabled and disabled features.
Step 4	(Optional) switch(config)# copy running-config startup-config	Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

Creating an OSPFv2 Instance

The first step in configuring OSPFv2 is to create an OSPFv2 instance. You assign a unique instance tag for this OSPFv2 instance. The instance tag can be any string.



- Note** The OSPF router ID changes without a restart on a Cisco Nexus 7000 switch when you have not configured a manual router ID in the following cases:
- Configuring an SVI or physical interface with a higher IP address than the current router ID on a setup without any configured loopback interfaces.
 - Configuring a loopback interface with any given IP address on a setup without any previously configured loopback interfaces.
 - Configuring a loopback interface with a higher IP address than the IP address of an existing configured loopback interface.

When a router ID changes, OSPF has to re-advertise all LSAs with the new router ID. To avoid this issue, you can configure a manual OSPF router ID.

Before you begin

Ensure that you have enabled the OSPF feature.

Use the **show ip ospf instance-tag** command to verify that the instance tag is not in use.

OSPFv2 must be able to obtain a router identifier (for example, a configured loopback address) or you must configure the router ID option.

Ensure that you are in the correct VDC (or use the `switchto vdc` command).

Procedure

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# router ospf instance-tag	Creates a new OSPFv2 instance with the configured instance tag. Note Use the no form of this command in global configuration mode to remove the OSPFv2 instance and all associated configurations. Using the no form of this command in the interface configuration mode does not remove the OSPF configuration. You must manually remove any OSPFv2 commands configured in interface mode.
Step 3	(Optional) switch(config-router)# router-id ip-address	Configures the OSPFv2 router ID. This IP address identifies this OSPFv2 instance and must exist on a configured interface in the system. This command restarts the OSPF process automatically and changes the router id after it is configured.
Step 4	(Optional) switch(config-router)# show ip ospf instance-tag	Displays OSPF information.
Step 5	(Optional) switch(config)# copy running-config startup-config	Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

Configuring OSPF Packet Size

MTU configured at interface level works in either the data plane or in the control plane but not at both planes at the same time.

When you configure MTU with a size lower than the supported size in data and control planes a few features that have minimum MTU requirements may not work in both the planes.

For example, MPLS VPN is supported in the data plane since this plane supports the MTU of 1500 bytes that the MPLS VPN requires. But control plane does not support MPLS VPN because this plane cannot handle the 1500-byte packets.

To make the configured MTU work in control plane for MPLS VPN, you need to manually configure the OSPF packet size so that OSPF works on the control plane. This is applicable from Cisco NX-OS Release 8.3(2) onwards.

Procedure

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# [no] router ospf <i>instance-tag</i>	Creates a new OSPF instance with the configured instance tag. Note Use the no form of this command in global configuration mode to remove the OSPFv2 instance and all associated configurations. Using the no form of this command in the interface configuration mode does not remove the OSPF configuration. You must manually remove any OSPFv2 commands configured in interface mode.
Step 3	switch(config-router)# router-id <i>ip-address</i>	Configures the OSPFv2 router ID. This IP address identifies this OSPFv2 instance and must exist on a configured interface in the system. This command restarts the OSPF process automatically and changes the router id after it is configured.
Step 4	switch(config-router)# packet-size <i>size</i>	<ul style="list-style-type: none"> • Configures the OSPFv2 packet size. The size range is from 572 to 9212 bytes. • You can configure the packet-size in the interface configuration mode also. • You can configure the packet-size <i>size</i> command even if their ospf mtu-ignore command is already configured in the network.
Step 5	(Optional) switch(config-router)# show ip ospf interface <i>interface-number</i>	Displays OSPF information.

Example

This example shows how to configure the OSPF packet-size:

```
router ospf 1
  router-id 3.3.3.3
  [no] packet-size 2000
```

This example shows the display of the OSPF packet-size:

```
Switch (config-router)# show ip ospf interface ethernet 1/25
Ethernet1/25 is up, line protocol is up
  IP address 1.0.0.1/24
  ----- snip -----
  Number of opaque link LSAs: 0, checksum sum 0
  Max Packet Size: 2000
```

Configuring Optional Parameters on an OSPFv2 Instance

You can configure optional parameters for OSPF. The following commands are available in the router configuration mode.

For more information about OSPFv2 instance parameters, see the “Configuring Advanced OSPFv2” section

Before you begin

Ensure that you have enabled the OSPF feature.

OSPFv2 must be able to obtain a router identifier (for example, a configured loopback address) or you must configure the router ID option.

Ensure that you are in the correct VDC (or use the `switchto vdc` command).

Procedure

	Command or Action	Purpose
Step 1	switch(config-router)# distance <i>number</i>	Configures the administrative distance for this OSPFv2 instance. The range is from 1 to 255. The default is 110.
Step 2	switch(config-router)# log-adjacency-changes [detail]	Generates a system message whenever a neighbor changes state.
Step 3	switch(config-router)# maximum-paths <i>path-number</i>	Configures the maximum number of equal OSPFv2 paths to a destination in the route table. This command is used for load balancing. The range is from 1 to 16. The default is 8.
Step 4	switch(config-router)# [no]name-lookup <i>path-number</i>	Enables the translation of OSPF router IDs to host names, either by looking up the local hosts database or querying DNS names in IPv6. This command makes it easier to identify a device because it displays the device by name rather than by its router ID or neighbor ID.

	Command or Action	Purpose
		Note To stop displaying OSPF router IDs as DNS names, use the no form of this command.
Step 5	switch(config-router)# passive-interface default	Suppresses routing updates on all interfaces. This command is overridden by the VRF or interface command mode configuration.

Example

This example shows how to create an OSPFv2 instance:

```
switch# configure terminal
switch(config)# router ospf 201
switch(config-router)# copy running-config startup-config
```

Configuring Networks in OSPFv2

You can configure a network to OSPFv2 by associating it through the interface that the router uses to connect to that network. You can add all networks to the default backbone area (Area 0), or you can create new areas using any decimal number or an IP address.



Note All areas must connect to the backbone area either directly or through a virtual link.



Note OSPF is not enabled on an interface until you configure a valid IP address for that interface.

Before you begin

Ensure that you have enabled the OSPF feature

Ensure that you are in the correct VDC (or use the switchto vdc command).

Procedure

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# interface <i>interface-type slot/port</i>	Enters interface configuration mode.
Step 3	switch(config-if)# ip address <i>ip-prefix/length</i>	Assigns an IP address and subnet mask to this interface.
Step 4	switch(config-if)# ip router ospf <i>instance-tag area area-id</i> [secondaries none]	Adds the interface to the OSPFv2 instance and area.

	Command or Action	Purpose
Step 5	(Optional) switch(config-if)# show ip ospf instance-tag interface interface-type slot/port	Displays OSPF information.
Step 6	(Optional) switch(config)# copy running-config startup-config	Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.
Step 7	(Optional) switch(config)# ip ospf cost number	Configures the OSPFv2 cost metric for this interface. The default is to calculate cost metric, based on reference bandwidth and interface bandwidth. The range is from 1 to 65535.
Step 8	(Optional) switch(config)# ip ospf dead-interval seconds	Configures the OSPFv2 dead interval, in seconds. The range is from 1 to 65535. The default is four times the hello interval, in seconds.
Step 9	(Optional) switch(config)# ip ospf hello-interval seconds	Configures the OSPFv2 hello interval, in seconds. The range is from 1 to 65535. The default is 10 seconds.
Step 10	(Optional) switch(config)# ip ospf mtu-ignore	Configures OSPFv2 to ignore any IP MTU mismatch with a neighbor. The default is to not establish adjacency if the neighbor MTU does not match the local interface MTU.
Step 11	(Optional) switch(config)# [default no] ip ospf passive-interface	Suppresses routing updates on the interface. This command overrides the router or VRF command mode configuration. The default option removes this interface mode command and reverts to the router or VRF configuration, if present.
Step 12	(Optional) switch(config)# ip ospf priority number	Configures the OSPFv2 priority, used to determine the DR for an area. The range is from 0 to 255. The default is 1.
Step 13	(Optional) switch(config)# ip ospf shutdown	Shuts down the OSPFv2 instance on this interface.

Example

This example shows how to add a network area 0.0.0.10 in OSPFv2 instance 201:

```
switch# configure terminal
switch(config)# interface ethernet 1/2

switch(config-if)# ip address 192.0.2.1/16
switch(config-if)# ip router ospf 201 area 0.0.0.10
switch(config-if)# copy running-config startup-config
```

Use the **show ip ospf interface** command to verify the interface configuration. Use the **show ip ospf neighbor** command to see the neighbors for this interface.

Configuring Authentication for an Area

You can configure authentication for all networks in an area or for individual interfaces in the area. Interface authentication configuration overrides area authentication.

Before you begin

Ensure that you have enabled the OSPF feature.

Ensure that all neighbors on an interface share the same authentication configuration, including the shared authentication key.

Create the key chain for this authentication configuration. See the *Cisco Nexus 7000 Series NX-OS Security Configuration Guide*

Ensure that you are in the correct VDC (or use the `switchto vdc` command).



Note For OSPFv2, the key identifier in the **key key-id** command supports values from 0 to 255 only.

Procedure

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# router ospf instance-tag	Creates a new OSPFv2 instance with the configured instance tag.
Step 3	switch(config-router)# area area-id authentication [message-digest]	Configures the authentication mode for an area.
Step 4	switch(config-router)# interface interface-type slot/port	Enters interface configuration mode.
Step 5	(Optional) Configure one of the following commands: <ul style="list-style-type: none"> • ip ospf authentication-key [0 3] key • ip ospf message-digest-key key-id md5 [0 3] key 	The first command configures simple password authentication for this interface. Use this command if the authentication is not set to key-chain or message-digest. The 0 keyword configures the password in clear text. The 3 keyword configures the password as 3DES encrypted. The second command configures message digest authentication for this interface. Use this command if the authentication is set to message-digest. The key-id range is from 1 to 255. The MD5 option 0 configures the password in clear text and 3 configures the pass key as 3DES encrypted.

	Command or Action	Purpose
Step 6	(Optional) switch(config)# show ip ospf instance-tag interface interface-type slot/port	Displays OSPF information.
Step 7	(Optional) switch(config)# copy running-config startup-config	Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

Configuring Authentication for an Interface

You can configure authentication for all networks in an area or for individual interfaces in the area. Interface authentication configuration overrides area authentication.

Before you begin

Ensure that you have enabled the OSPF feature.

Ensure that all neighbors on an interface share the same authentication configuration, including the shared authentication key.

Create the key chain for this authentication configuration. See the *Cisco Nexus 7000 Series NX-OS Security Configuration Guide*.

Ensure that you are in the correct VDC (or use the `switchto vdc` command).



Note For OSPFv2, the key identifier in the `key key-id` command supports values from 0 to 255 only.

Procedure

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# interface interface-type slot/port	Enters interface configuration mode.
Step 3	switch(config-if)# ip ospf authentication [message-digest]	Enables interface authentication mode for OSPFv2 for either cleartext or message-digest type. Overrides area-based authentication for this interface. All neighbors must share this authentication type.
Step 4	(Optional) switch(config-if)# ip ospf authentication key-chain key-name	Configures interface authentication to use key chains for OSPFv2. For details on key chains, see the <i>Cisco Nexus 7000 Series NX-OS Security Configuration Guide</i> .
Step 5	(Optional) switch(config-if)# ip ospf authentication-key [0 3 7] key	Configures simple password authentication for this interface. Use this command if the authentication is not set to key-chain or message-digest.

	Command or Action	Purpose
		<p>The options are as follows:</p> <ul style="list-style-type: none"> • 0—configures the password in clear text. • 3—configures the pass key as 3DES encrypted. • 7—configures the key as Cisco type 7 encrypted.
Step 6	(Optional) switch(config-if)# ip ospf message-digest-key <i>key-id</i> md5 [0 3 7] <i>key</i>	<p>Configures message digest authentication for this interface. Use this command if the authentication is set to message-digest. The key-id range is from 1 to 255. The MD5 options are as follows:</p> <ul style="list-style-type: none"> • 0—configures the password in clear text. • 3—configures the pass key as 3DES encrypted. • 7—configures the key as Cisco type 7 encrypted.
Step 7	(Optional) switch(config-if)# show ip ospf instance-tag interface <i>interface-type slot/port</i>	Displays OSPF information.
Step 8	(Optional) switch(config)# copy running-config startup-config	Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

Example

This example shows how to set an interface for simple, unencrypted passwords and set the password for Ethernet interface 1/2:

```
switch# configure terminal
switch(config)# router ospf 201
switch(config-router)# exit
switch(config)# interface ethernet 1/2

switch(config-if)# ip router ospf 201 area 0.0.0.10
switch(config-if)# ip ospf authentication
switch(config-if)# ip ospf authentication-key 0 mypass
switch(config-if)# copy running-config startup-config
```

Configuring Advanced OSPFv2

Configuring Filter Lists for Border Routers

You can separate your OSPFv2 domain into a series of areas that contain related networks. All areas must connect to the backbone area through an area border router (ABR). OSPFv2 domains can connect to external domains as well, through an autonomous system border router (ASBR).

ABRs have the following optional configuration parameters:

- Area range—Configures route summarization between areas.
- Filter list—Filters the Network Summary (type 3) LSAs that are allowed in from an external area.

ASBRs also support filter lists.

Before you begin

Ensure that you have enabled the OSPF feature.

Create the route map that the filter list uses to filter IP prefixes in incoming or outgoing Network Summary (type 3) LSAs.

Ensure that you are in the correct VDC (or use the `switchto vdc` command).

Procedure

	Command or Action	Purpose
Step 1	<code>switch# configure terminal</code>	Enters global configuration mode.
Step 2	<code>switch(config)# router ospf instance-tag</code>	Creates a new OSPFv2 instance with the configured instance tag.
Step 3	<code>switch(config-router)# area area-id filter-list route-map map-name {in out}</code>	Filters incoming or outgoing Network Summary (type 3) LSAs on an ABR.
Step 4	(Optional) <code>switch(config-if)# show ip ospf policy statistics area id filter-list {in out}</code>	Displays OSPF policy information.
Step 5	(Optional) <code>switch(config)# copy running-config startup-config</code>	Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

Example

This example shows how to configure a filter list in area 0.0.0.10:

```
switch# configure terminal
switch(config)# router ospf 201
switch(config-router)# area 0.0.0.10 filter-list route-map FilterLSAs in
switch(config-router)# copy running-config startup-config
```

Configuring Stub Areas

You can configure a stub area for part of an OSPFv2 domain where external traffic is not necessary. Stub areas block AS External (type 5) LSAs and limit unnecessary routing to and from selected networks. You can optionally block all summary routes from going into the stub area.

Before you begin

Ensure that you have enabled the OSPF feature.

Ensure that there are no virtual links or ASBRs in the proposed stub area.

Ensure that you are in the correct VDC (or use the `switchto vdc` command).

Procedure

	Command or Action	Purpose
Step 1	<code>switch# configure terminal</code>	Enters global configuration mode.
Step 2	<code>switch(config)# router ospf instance-tag</code>	Creates a new OSPFv2 instance with the configured instance tag.
Step 3	<code>switch(config-router)# area area-id stub</code>	Creates this area as a stub area.
Step 4	(Optional) <code>switch(config-router)# area area-id default-cost cost</code>	Sets the cost metric for the default summary route sent into this stub area. The range is from 0 to 16777215. The default is 1.
Step 5	(Optional) <code>switch(config-if)# show ip ospf instance-tag</code>	Displays OSPF information.
Step 6	(Optional) <code>switch(config)# copy running-config startup-config</code>	Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

Example

This example shows how to create a stub area:

```
switch# configure terminal
switch(config)# router ospf 201
switch(config-router)# area 0.0.0.10 stub
switch(config-router)# copy running-config startup-config
```

Configuring a Totally Stubby Area

You can create a totally stubby area and prevent all summary route updates from going into the stub area. To create a totally stubby area, use the following command in router configuration mode:

Command	Purpose
<code>router ospf instance-tag</code>	Creates this area as a totally stubby area.

Configuring NSSA

You can configure an NSSA for part of an OSPFv2 domain where limited external traffic is required. You can optionally translate this external traffic to an AS External (type 5) LSA and flood the OSPFv2 domain with this routing information. An NSSA can be configured with the following optional parameters:

- No redistribution—Redistributed routes bypass the NSSA and are redistributed to other areas in the OSPFv2 autonomous system. Use this option when the NSSA ASBR is also an ABR.
- Default information originate—Generates an NSSA External (type 7) LSA for a default route to the external autonomous system. Use this option on an NSSA ASBR if the ASBR contains the default route in the routing table. This option can be used on an NSSA ABR whether or not the ABR contains the default route in the routing table.
- Route map—Filters the external routes so that only those routes that you want are flooded throughout the NSSA and other areas.
- Translate—Translates NSSA External LSAs to AS External LSAs for areas outside the NSSA. Use this command on an NSSA ABR to flood the redistributed routes throughout the OSPFv2 autonomous system. You can optionally suppress the forwarding address in these AS External LSAs. If you choose this option, the forwarding address is set to 0.0.0.0.
- No summary—Blocks all summary routes from flooding the NSSA. Use this option on the NSSA ABR.

Before you begin

Ensure that you have enabled the OSPF feature.

Ensure that there are no virtual links in the proposed NSSA and that it is not the backbone area.

Ensure that you are in the correct VDC (or use the `switchto vdc` command).

Procedure

	Command or Action	Purpose
Step 1	<code>switch# configure terminal</code>	Enters global configuration mode.
Step 2	<code>switch(config)# router ospf instance-tag</code>	Creates a new OSPFv2 instance with the configured instance tag.
Step 3	<code>switch(config-router)# area area-id nssa [no-redistribution] [default-information-originate]originate [route-map map-name] [no-summary] [translate type7 {always never} [suppress-fa]]</code>	Creates this area as an NSSA.
Step 4	(Optional) <code>switch(config-router)# area area-id default-cost cost</code>	Sets the cost metric for the default summary route sent into this NSSA.
Step 5	(Optional) <code>switch(config-if)# show ip ospf instance-tag</code>	Displays OSPF information.

	Command or Action	Purpose
Step 6	(Optional) <code>switch(config)# copy running-config startup-config</code>	Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

Example

This example shows how to create an NSSA that blocks all summary route updates:

```
switch# configure terminal
switch(config)# router ospf 201
switch(config-router)# area 0.0.0.10 nssa no-summary
switch(config-router)# copy running-config startup-config
```

This example shows how to create an NSSA that generates a default route:

```
switch# configure terminal
switch(config)# router ospf 201
switch(config-router)# area 0.0.0.10 nssa default-info-originate
switch(config-router)# copy running-config startup-config
```

This example shows how to create an NSSA that filters external routes and blocks all summary route updates:

```
switch# configure terminal
switch(config)# router ospf 201
switch(config-router)# area 0.0.0.10 nssa route-map ExternalFilter no-summary
switch(config-router)# copy running-config startup-config
```

This example shows how to create an NSSA that always translates NSSA External (type 5) LSAs to AS External (type 7) LSAs:

```
switch# configure terminal
switch(config)# router ospf 201
switch(config-router)# area 0.0.0.10 nssa translate type 7 always
switch(config-router)# copy running-config startup-config
```

Configuring Virtual Links

A virtual link connects an isolated area to the backbone area through an intermediate area. You can configure the following optional parameters for a virtual link:

- Authentication—Sets a simple password or MD5 message digest authentication and associated keys.
- Dead interval—Sets the time that a neighbor waits for a Hello packet before declaring the local router as dead and tearing down adjacencies.
- Hello interval—Sets the time between successive Hello packets.
- Retransmit interval—Sets the estimated time between successive LSAs.
- Transmit delay—Sets the estimated time to transmit an LSA to a neighbor.



Note You must configure the virtual link on both routers involved before the link becomes active.

You cannot add a virtual link to a stub area.

Before you begin

Ensure that you have enabled the OSPF feature.

Ensure that you are in the correct VDC (or use the `switchto vdc` command).

Procedure

	Command or Action	Purpose
Step 1	<code>switch# configure terminal</code>	Enters global configuration mode.
Step 2	<code>switch(config)# router ospf instance-tag</code>	Creates a new OSPFv2 instance with the configured instance tag.
Step 3	<code>switch(config-router)# area area-id virtual link router-id</code>	Creates one end of a virtual link to a remote router. You must create the virtual link on that remote router to complete the link.
Step 4	(Optional) <code>switch(config-router-vlink)# show ip ospf virtual-link [brief]</code>	Displays OSPF virtual link information.
Step 5	(Optional) <code>switch(config-router-vlink)# authentication [key-chain key-id message-digest null]</code>	Overrides area-based authentication for this virtual link.
Step 6	(Optional) <code>switch(config-router-vlink)# authentication-key [0 3] key</code>	Configures a simple password for this virtual link. Use this command if the authentication is not set to key-chain or message-digest. 0 configures the password in clear text. 3 configures the password as 3DES encrypted.
Step 7	(Optional) <code>switch(config-router-vlink)# dead-interval seconds</code>	Configures the OSPFv2 dead interval, in seconds. The range is from 1 to 65535. The default is four times the hello interval, in seconds.
Step 8	(Optional) <code>switch(config-router-vlink)# hello-interval seconds</code>	Configures the OSPFv2 hello interval, in seconds. The range is from 1 to 65535. The default is 10 seconds.
Step 9	(Optional) <code>switch(config-router-vlink)# message-digest-key key-id md5 [0 3] key</code>	Configures message digest authentication for this virtual link. Use this command if the authentication is set to message-digest. 0 configures the password in cleartext. 3 configures the pass key as 3DES encrypted.
Step 10	(Optional) <code>switch(config-router-vlink)# retransmit-interval seconds</code>	Configures the OSPFv2 retransmit interval, in seconds. The range is from 1 to 65535. The default is 5.

	Command or Action	Purpose
Step 11	(Optional) switch(config-router-vlink)# transmit-delay <i>seconds</i>	Configures the OSPFv2 transmit-delay, in seconds. The range is from 1 to 450. The default is 1.
Step 12	(Optional) switch(config)# copy running-config startup-config	Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

Example

This example shows how to create a simple virtual link between two ABRs.

The configuration for ABR 1 (router ID 27.0.0.55) is as follows:

```
switch# configure terminal
switch(config)# router ospf 201
switch(config-router)# area 0.0.0.10 virtual-link 10.1.2.3
switch(config-router)# copy running-config startup-config
```

The configuration for ABR 2 (Router ID 10.1.2.3) is as follows:

```
switch# configure terminal
switch(config)# router ospf 101
switch(config-router)# area 0.0.0.10 virtual-link 27.0.0.55
switch(config-router)# copy running-config startup-config
```

Configuring Redistribution

You can redistribute routes learned from other routing protocols into an OSPFv2 autonomous system through the ASBR.

You can configure the following optional parameters for route redistribution in OSPF:

- **Default information originate**—Generates an AS External (type 5) LSA for a default route to the external autonomous system.



Note Default information originate ignores match statements in the optional route map.

- **Default metric**—Sets all redistributed routes to the same cost metric.



Note If you redistribute static routes, Cisco NX-OS requires the configuration of **default-information originate** command under the router OSPF process to successfully redistribute or generate the default static route.

Before you begin

Ensure that you have enabled the OSPF feature.

Create the necessary route maps used for redistribution.

Ensure that you are in the correct VDC (or use the `switchto vdc` command).

Procedure

	Command or Action	Purpose
Step 1	<code>switch# configure terminal</code>	Enters global configuration mode.
Step 2	<code>switch(config)# router ospf instance-tag</code>	Creates a new OSPFv2 instance with the configured instance tag.
Step 3	<code>switch(config-router)# redistribute {bgp id direct eigrp id isis id ospf id rip id static} route-map map-name</code>	Redistributes the selected protocol into OSPF through the configured route map. Note If you redistribute static routes, Cisco NX-OS also redistributes the default static route.
Step 4	<code>switch(config-router)# default-information originate [always] [route-map map-name]</code>	Creates a default route into this OSPF domain if the default route exists in the RIB. Use the following optional keywords: <ul style="list-style-type: none"> • always—Always generate the default route of 0.0.0. even if the route does not exist in the RIB • route-map—Generate the default route if the route map returns true. Note This command ignores match statements in the route map.
Step 5	<code>switch(config-router)# default-metric [cost]</code>	Sets the cost metric for the redistributed routes. This command does not apply to directly connected routes. Use a route map to set the default metric for directly connected routes.
Step 6	(Optional) <code>switch(config)# copy running-config startup-config</code>	Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

Example

This example shows how to create an NSSA that blocks all summary route updates:

```
switch# configure terminal
switch(config)# router ospf 201
switch(config-router)# redistribute bgp route-map FilterExternalBGP
switch(config-router)# copy running-config startup-config
```

Limiting the Number of Redistributed Routes

Route redistribution can add many routes to the OSPFv2 route table. You can configure a maximum limit to the number of routes accepted from external protocols. OSPFv2 provides the following options to configure redistributed route limits:

- **Fixed limit**—Logs a message when OSPFv2 reaches the configured maximum. OSPFv2 does not accept any more redistributed routes. You can optionally configure a threshold percentage of the maximum where OSPFv2 logs a warning when that threshold is passed.
- **Warning only**—Logs a warning only when OSPFv2 reaches the maximum. OSPFv2 continues to accept redistributed routes.
- **Withdraw**—Starts the timeout period when OSPFv2 reaches the maximum. After the timeout period, OSPFv2 requests all redistributed routes if the current number of redistributed routes is less than the maximum limit. If the current number of redistributed routes is at the maximum limit, OSPFv2 withdraws all redistributed routes. You must clear this condition before OSPFv2 accepts more redistributed routes.
- You can optionally configure the timeout period.

Before you begin

Ensure that you have enabled the OSPF feature.

Ensure that you are in the correct VDC (or use the `switchto vdc` command).

Procedure

	Command or Action	Purpose
Step 1	<code>switch# configure terminal</code>	Enters global configuration mode.
Step 2	<code>switch(config)# router ospf instance-tag</code>	Creates a new OSPFv2 instance with the configured instance tag.
Step 3	<code>switch(config-router)# redistribute {bgp id direct eigrp id isis id ospf id rip id static} route-map map-name</code>	Redistributes the selected protocol into OSPF through the configured route map.
Step 4	<code>switch(config-router)# redistribute maximum-prefix max [threshold] [warning-only withdraw [num-retries timeout]]</code>	Specifies a maximum number of prefixes that OSPFv2 distributes. The range is from 0 to 65536. Optionally specifies the following: <ul style="list-style-type: none"> • threshold—Percent of maximum prefixes that trigger a warning message. • warning-only—Logs an warning message when the maximum number of prefixes is exceeded. • withdraw—Withdraws all redistributed routes. Optionally tries to retrieve the redistributed routes. The <i>num-retries</i> range is from 1 to 12. The <i>timeout</i> is 60 to 600 seconds. The default is 300 seconds. Use

	Command or Action	Purpose
		the clear ip ospf redistribution command if all routes are withdrawn.
Step 5	(Optional) switch(config-router)# show running-config ospf	Displays the OSPFv2 configuration.
Step 6	(Optional) switch(config)# copy running-config startup-config	Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

Example

This example shows how to limit the number of redistributed routes into OSPF:

```
switch# configure terminal
switch(config)# router ospf 201
switch(config-router)# redistribute bgp route-map FilterExternalBGP
switch(config-router)# redistribute maximum-prefix 1000 75
```

Configuring Route Summarization

You can configure route summarization for inter-area routes by configuring an address range that is summarized. You can also configure route summarization for external, redistributed routes by configuring a summary address for those routes on an ASBR.

Before you begin

Ensure that you have enabled the OSPF feature.

Ensure that you are in the correct VDC (or use the `switchto vdc` command).

Procedure

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# router ospf instance-tag	Creates a new OSPFv2 instance with the configured instance tag.
Step 3	Configure one of the following commands: <ul style="list-style-type: none"> • area area-id range ip-prefix/length [no-advertise] [cost cost] • summary-address ip-prefix/length [no-advertise tag tag] 	<p>The first command creates a summary address on an ABR for a range of addresses and optionally does not advertise this summary address in a Network Summary (type 3) LSA. The cost range is from 0 to 16777215.</p> <p>The second command creates a summary address on an ASBR for a range of addresses and optionally assigns a tag for this summary address that can be used for redistribution with route maps.</p>

	Command or Action	Purpose
Step 4	(Optional) switch(config-router)# [no] discard route {internal external}	When you configure a summary address, Cisco NX-OS software automatically configures a discard route for the summary address to prevent routing black holes and route loops. You can use the no form of this command to prevent the discard routes from being created.
Step 5	(Optional) switch(config-router)# show ip ospf summary-address	Displays information about OSPF summary addresses.
Step 6	(Optional) switch(config)# copy running-config startup-config	Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

Example

This example shows how to create summary addresses between areas on an ABR:

```
switch# configure terminal
switch(config)# router ospf 201
switch(config-router)# area 0.0.0.10 range 10.3.0.0/16
switch(config-router)# copy running-config startup-config
```

This example shows how to create summary addresses on an ASBR:

```
switch# configure terminal
switch(config)# router ospf 201
switch(config-router)# summary-address 10.5.0.0/16
switch(config-router)# copy running-config startup-config
```

Configuring Stub Route Advertisements

Use stub route advertisements when you want to limit the OSPFv2 traffic through this router for a short time.

Stub route advertisements can be configured with the following optional parameters:

- **on startup**—Sends stub route advertisements for the specified announce time.
- **wait-for bgp**—Sends stub router advertisements until BGP converges.



Note You should not save the running configuration of a router when it is configured for a graceful shutdown because the router continues to advertise a maximum metric after it is reloaded.

Before you begin

Ensure that you have enabled the OSPF feature.

Ensure that you are in the correct VDC (or use the switchto vdc command).

Procedure

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# router ospf <i>instance-tag</i>	Creates a new OSPFv2 instance with the configured instance tag.
Step 3	switch(config-router)# max-metric router-lsa [external-lsa [<i>max-metric-value</i>]] [include-stub] [on-startup [<i>seconds</i>]] [wait-for bgp tag] [summary-lsa [<i>max-metric-value</i>]]	Configures OSPFv2 stub route advertisements.
Step 4	(Optional) switch(config)# copy running-config startup-config	Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

Example

This example shows how to enable the stub router advertisements on startup for the default 600 seconds:

```
switch# configure terminal
switch(config)# router ospf 201
switch(config-router)# max-metric router-lsa on-startup
switch(config-router)# copy running-config startup-config
```

Configuring the Administrative Distance of Routes

Beginning with Cisco NX-OS Release 6.1, you can set the administrative distance of routes added by OSPFv2 into the RIB.

The administrative distance is a rating of the trustworthiness of a routing information source. A higher value indicates a lower trust rating. Typically, a route can be learned through more than one routing protocol. The administrative distance is used to discriminate between routes learned from more than one routing protocol. The route with the lowest administrative distance is installed in the IP routing table.

Before you begin

Ensure that you have enabled OSPF.

Ensure that you are in the correct VDC (or use the **switchto vdc** command).

See the guidelines and limitations for this feature.

Procedure

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# router ospf <i>instance-tag</i>	Creates a new OSPFv2 instance with the configured instance tag.

	Command or Action	Purpose
Step 3	switch(config-router)# [no] table-map <i>map-name</i> [filter]	Configures the policy for filtering or modifying OSPFv2 routes before sending them to the RIB. You can enter up to 63 alphanumeric characters for the map name. The filter keyword specifies that only routes that are permitted by the route map(<i>map-name</i>) configuration are downloaded to the routing information base (RIB).
Step 4	switch(config-router)# exit	Exits router configuration mode.
Step 5	switch(config)# route-map <i>map-name</i> [permit deny] [<i>seq</i>]	Creates a route map or enters route-map configuration mode for an existing route map. Use <i>seq</i> to order the entries in a route map. Note The permit option enables you to set the distance. If you use the deny option, the default distance is applied.
Step 6	switch(config-route-map)# match route-type <i>route-type</i>	Matches against one of the following route types: <ul style="list-style-type: none"> • external: The external route (BGP, EIGRP, and OSPF type 1 or 2) • inter-area: OSPF inter-area route • internal: The internal route (including the OSPF intra- or inter-area) • intra-area: OSPF intra-area route • nssa-external: The NSSA external route (OSPF type 1 or 2) • type-1: The OSPF external type 1 route • type-2: The OSPF external type 2 route
Step 7	switch(config-route-map)# match ip route-source prefix-list <i>name</i>	Matches the IPv4 route source address or router ID of a route to one or more IP prefix lists. Use the ip prefix-list command to create the prefix list.
Step 8	switch(config-route-map)# match ip address prefix-list <i>name</i>	Matches against one or more IPv4 prefix lists. Use the ip prefix-list command to create the prefix list.
Step 9	switch(config-route-map)# set distance <i>value</i>	Sets the administrative distance of routes for OSPFv2. The range is from 1 to 255.

	Command or Action	Purpose
Step 10	(Optional) switch(config)# copy running-config startup-config	Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

Example

This example shows how to configure the OSPFv2 administrative distance for inter-area routes to 150, for external routes to 200, and for all prefixes in prefix list p1 to 190:

```
switch# configure terminal
switch(config)# router ospf 201
switch(config-router)# table-map foo
switch(config-router)# exit
switch(config)# route-map foo permit 10
switch(config-route-map)# match route-type inter-area
switch(config-route-map)# set distance 150
switch(config)# route-map foo permit 20
switch(config-route-map)# match route-type external
switch(config-route-map)# set distance 200
switch(config)# route-map foo permit 30
switch(config-route-map)# match ip route-source prefix-list p1
switch(config-route-map)# match ip address prefix-list p1
switch(config-route-map)# set distance 190
```

The following example shows how to configure a route map for blocking the next hops that are learned through VLAN 10:

```
switch(config)# route-map Filter-OSPF 10 deny
switch(config-route-map)# match interface VLAN 10
switch(config-route-map)# exit
switch(config)# route-map Filter-OSPF 20 permit
```

The following example shows how to configure the **table-map** command with the **filter** keyword to use a route map (Filter-OSPF) to remove the next-hop path that is learned through VLAN 10 but not the next-hop path that is learned through VLAN 20:

```
switch(config)# route ospf p1
switch(config-router)# table-map Filter-OSPF filter
```

Modifying the Default Timers

OSPFv2 includes a number of timers that control the behavior of protocol messages and shortest path first (SPF) calculations. OSPFv2 includes the following optional timer parameters:

- LSA arrival time—Sets the minimum interval allowed between LSAs that arrive from a neighbor. LSAs that arrive faster than this time are dropped.
- Pacing LSAs—Sets the interval at which LSAs are collected into a group and refreshed, checksummed, or aged. This timer controls how frequently LSA updates occur and optimizes how many are sent in an LSA update message.
- Throttle LSAs—Sets the rate limits for generating LSAs. This timer controls how frequently LSAs are generated after a topology change occurs.
- Throttle SPF calculation—Controls how frequently the SPF calculation is run.

At the interface level, you can also control the following timers:

- Retransmit interval—Sets the estimated time between successive LSAs
- Transmit delay—Sets the estimated time to transmit an LSA to a neighbor.

Before you begin

Ensure that you have enabled the OSPF feature.

Ensure that you are in the correct VDC (or use the `switchto vdc` command).

Procedure

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# router ospf <i>instance-tag</i>	Creates a new OSPFv2 instance with the configured instance tag.
Step 3	switch(config-router)# timers lsa-arrival msec	Sets the LSA arrival time in milliseconds. The range is from 10 to 600000. The default is 1000 milliseconds.
Step 4	switch(config-router)# timers lsa-group-pacing seconds	Sets the interval in seconds for grouping LSAs. The range is from 1 to 1800. The default is 240 seconds.
Step 5	switch(config-router)# timers throttle lsa <i>start-time hold-interval max-time</i>	Sets the rate limit in milliseconds for generating LSAs with the following timers: <ul style="list-style-type: none"> • <i>start-time</i>—The range is from 50 to 5000 milliseconds. The default value is 50 milliseconds. • <i>hold-interval</i>—The range is from 50 to 30,000 milliseconds. The default value is 5000 milliseconds. • <i>max-time</i>—The range is from 50 to 30,000 milliseconds. The default value is 5000 milliseconds
Step 6	switch(config-router)# timers throttle spf <i>delay-time hold-time max-wait</i>	Sets the SPF best path schedule initial delay time and the minimum hold time in seconds between SPF best path calculations. The range is from 1 to 600000. The default is no delay time and 5000 millisecond hold time.
Step 7	switch(config)# interface <i>type slot/port</i>	Enters interface configuration mode.
Step 8	switch(config-if)# ip ospf hello-interval <i>seconds</i>	Sets the hello interval this interface. The range is from 1 to 65535. The default is 10.

	Command or Action	Purpose
Step 9	switch(config-if)# ip ospf dead-interval <i>seconds</i>	Sets the dead interval for this interface. The range is from 1 to 65535.
Step 10	switch(config-if)# ip ospf retransmit-interval <i>seconds</i>	Sets the estimated time in seconds between LSAs transmitted from this interface. The range is from 1 to 65535. The default is 5.
Step 11	switch(config-if)# ip ospf transmit-delay <i>seconds</i>	Sets the estimated time in seconds to transmit an LSA to a neighbor. The range is from 1 to 450. The default is 1.
Step 12	(Optional) switch(config-if)# show ip ospf	Displays information about OSPF.
Step 13	(Optional) switch(config)# copy running-config startup-config	Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

Example

This example shows how to control LSA flooding with the `lsa-group-pacing` option:

```
switch# configure terminal
switch(config)# router ospf 201
switch(config-router)# timers lsa-group-pacing 300
switch(config-router)# copy running-config startup-config
```

Configuring Graceful Restart

Graceful restart is enabled by default. You can configure the following optional parameters for graceful restart in an OSPFv2 instance:

- **Grace period**—Configures how long neighbors should wait after a graceful restart has started before tearing down adjacencies.
- **Helper mode disabled**—Disables helper mode on the local OSPFv2 instance. OSPFv2 does not participate in the graceful restart of a neighbor.
- **Planned graceful restart only**—Configures OSPFv2 to support graceful restart only in the event of a planned restart.

Before you begin

Ensure that you have enabled OSPF.

Ensure that all neighbors are configured for graceful restart with matching optional parameters set.

Ensure that you are in the correct VDC (or use the `switchto vdc` command).

Procedure

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# router ospf instance-tag	Creates a new OSPFv2 instance with the configured instance tag.
Step 3	switch(config-router)# graceful-restart	Enables a graceful restart. A graceful restart is enabled by default.
Step 4	(Optional) switch(config-router)# graceful-restart grace-period seconds	Sets the grace period, in seconds. The range is from 5 to 1800. The default is 60 seconds.
Step 5	(Optional) switch(config-router)# graceful-restart helper-disable	Disables helper mode. This feature is enabled by default.
Step 6	(Optional) switch(config-router)# graceful-restart planned-only	Configures a graceful restart for planned restarts only.
Step 7	(Optional) switch(config-if)# show ip ospf instance-tag	Displays OSPF information.
Step 8	(Optional) switch(config)# copy running-config startup-config	Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

Example

This example shows how to enable a graceful restart if it has been disabled and set the grace period to 120 seconds:

```
switch# configure terminal
switch(config)# router ospf 201
switch(config-router)# graceful-restart
switch(config-router)# graceful-restart grace-period 120
switch(config-router)# copy running-config startup-config
```

Restarting an OSPFv2 Instance

You can restart an OSPv2 instance. This action clears all neighbors for the instance.

To restart an OSPFv2 instance and remove all associated neighbors, use the following command:

Command	Purpose
restart ospf instance-tag	Restarts the OSPFv2 instance and removes all neighbors.

Configuring OSPFv2 with Virtualization

You can configure multiple OSPFv2 instances in each VDC. You can also create multiple VRFs within each VDC and use the same or multiple OSPFv2 instances in each VRF. You assign an OSPFv2 interface to a VRF.



Note Configure all other parameters for an interface after you configure the VRF for an interface. Configuring a VRF for an interface deletes all the configuration for that interface

Before you begin

Create the VDCs.

Ensure that you have enabled the OSPF feature.

Ensure that you are in the correct VDC (or use the `switchto vdc` command).

Procedure

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# vrf context <i>vrf-name</i>	Creates a new VRF and enters VRF configuration mode.
Step 3	switch(config)# router ospf <i>instance-tag</i>	Creates a new OSPFv2 instance with the configured instance tag.
Step 4	switch(config-router)# vrf <i>vrf-name</i>	Enters VRF configuration mode.
Step 5	(Optional) switch(config-router-vrf)# maximum-paths <i>path</i>	Configures the maximum number of equal OSPFv2 paths to a destination in the route table for this VRF. Used for load balancing.
Step 6	switch(config-router-vrf)# interface <i>interface-type slot/port</i>	Enters interface configuration mode.
Step 7	switch(config-if)# vrf member <i>vrf-name</i>	Adds this interface to a VRF.
Step 8	switch(config-if)# ip address <i>ip-prefix/length</i>	Configures an IP address for this interface. You must do this step after you assign this interface to a VRF.
Step 9	switch(config-if)# ip router ospf <i>instance-tag area area-id</i>	Assigns this interface to the OSPFv2 instance and area configured.
Step 10	(Optional) switch(config)# copy running-config startup-config	Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

Example

This example shows how to create summary addresses between areas on an ABR:

```
switch# configure terminal
switch(config)# vrf context NewVRF
switch(config)# router ospf 201
switch(config)# interface ethernet 1/2

switch(config-if)# vrf member NewVRF
switch(config-if)# ip address 192.0.2.1/16
switch(config-if)# ip router ospf 201 area 0
switch(config)# copy running-config startup-config
```

Verifying the OSPFv2 Configuration

To display the OSPFv2 configuration, perform one of the following tasks:

Command	Purpose
show ip ospf [<i>instance-tag</i>] [vrf <i>vrf-name</i>]	Displays the information about one or more OSPFv2 routing instances. The output includes the following area-level counts: <ul style="list-style-type: none"> • Interfaces in this area—A count of all interfaces added to this area (configured interfaces). • Active interfaces—A count of all interfaces considered to be in router link states and SPF (UP interfaces). • Passive interfaces—A count of all interfaces considered to be OSPF passive (no adjacencies will be formed). • Loopback interfaces—A count of all local loopback interfaces.
show ip ospf border-routers [vrf { <i>vrf-name</i> all default management }]	Displays the OSPFv2 link-state database summary.
show ip ospf interface <i>number</i> [vrf { <i>vrf-name</i> all default management }]	Displays the OSPFv2 interface configuration.
show ip ospf lsa-content-changed-list <i>neighbor-id interface-type number</i> [vrf { <i>vrf-name</i> all default management }]	Displays the OSPFv2 LSAs that have changed.
show ip ospf neighbors [<i>neighbor-id</i>] [detail] [<i>interface-type number</i>] [vrf { <i>vrf-name</i> all default management }] [summary]	Displays the list of OSPFv2 neighbors.
show ip ospf request-list <i>neighbor-id interface-type number</i> [vrf { <i>vrf-name</i> all default management }]	Displays the list of OSPFv2 link-state requests.

Command	Purpose
show ip ospf retransmission-list <i>neighbor-id interface-type number</i> [vrf { <i>vrf-name</i> all default management }]	Displays the list of OSPFv2 link-state retransmissions.
show ip ospf route [<i>ospf-route</i>] [summary] [vrf { <i>vrf-name</i> all default management }]	Displays the internal OSPFv2 routes.
show ip ospf summary-address [vrf { <i>vrf-name</i> all default management }]	Displays information about the OSPFv2 summary addresses.
show ip ospf virtual-links [brief] [vrf { <i>vrf-name</i> all default management }]	Displays information about OSPFv2 virtual links.
show ip ospf vrf { <i>vrf-name</i> all default management }	Displays information about VRF-based OSPFv2 configuration.
show running-configuration ospf	Displays the current running OSPFv2 configuration.

Monitoring OSPFv2

To display OSPFv2 statistics, use the following commands:

Command	Purpose
show ip ospf policy statistics area <i>area-id filter-list</i> { in out } [vrf <i>vrf-name</i> all default management]	Displays the OSPFv2 route policy statistics for an area.
show ip ospf policy statistics redistribute { bgp <i>id</i> direct eigrp <i>id</i> isis <i>id</i> ospf <i>id</i> rip <i>id</i> static } [vrf { <i>vrf-name</i> all default management }]	Displays the OSPFv2 route policy statistics.
show ip ospf statistics [vrfnumber [vrf { <i>vrf-name</i> all default management }]	Displays the OSPFv2 event counters.
show ip ospf traffic <i>interface-type number</i> [vrf { <i>vrf-name</i> all default management }]	Displays the OSPFv2 packet counters.

Configuration Examples for OSPFv2

```
feature ospf
router ospf 201
  router-id 290.0.2.1
interface ethernet 1/2

ip router ospf 201 area 0.0.0.10
ip ospf authentication
ip ospf authentication-key 0 mypass
```


Related Documents for OSPFv2

For more information related to OSPFv2 CLI commands, see the *Cisco Nexus 5000 Series Command Reference*

Feature History for OSPFv2

This table includes only the updates for those releases that have resulted in additions or changes to the feature.

Table 16: Feature History for OSPFv2

Feature Name	Release	Feature Information
OSPF Packet-size	8.3(2)	Added support for configuring OSPF packet-size.
OSPF—Distribute List to Filter Paths	6.2(6a)	Added support for filtering next-hop paths for an OSPF route to prevent the path from being added to the RIB.
Administrative distance of routes	6.2(2)	Added the filter keyword to the table-map command to specify that only routes permitted by the route map are downloaded to the RIB.
Route summarization	6.2(2)	Added the ability to prevent discard routes from being created
OSPFv2	6.2(2)	Added support for the optional name lookup parameter for OSPFv2 instances.
OSPFv2	6.1(1)	Added support for more than four process instances for OSPFv2 per VDC.
OSPFv2	6.1(1)	Added support for configuring the administrative distance of routes for OSPFv2.
Passive interface	5.2(1)	Added support for setting the passive interface mode on all interfaces in the router or VRF.
OSPFv2	5.1(2)	Added options for the max-metric router-lsa command.
BFD	5.0(2)	Added support for BFD. See the <i>Cisco Nexus 7000 Series NX-OS Interfaces Configuration Guide</i> for more information.
OSPFv2	4.0(1)	This feature was introduced.



CHAPTER 8

Configuring OSPFv3

This chapter contains the following sections:

- [Finding Feature Information, on page 145](#)
- [Information About OSPFv3, on page 145](#)
- [Advanced Features, on page 154](#)
- [Prerequisites for OSPFv3, on page 158](#)
- [Guidelines and Limitations for OSPFv3, on page 158](#)
- [Default Settings for OSPFv3, on page 160](#)
- [Configuring Basic OSPFv3, on page 160](#)
- [Configuring Advanced OSPFv3, on page 166](#)
- [Configuring OSPFv3 Encryption at Router Level, on page 187](#)
- [Configuring OSPFv3 Encryption at Area Level, on page 188](#)
- [Configuring OSPFv3 Encryption at Interface Level, on page 188](#)
- [Configuring OSPFv3 Encryption for Virtual Links, on page 190](#)
- [Configuration Examples for OSPFv3, on page 191](#)
- [Related Documents for OSPFv3, on page 191](#)
- [Feature History for OSPFv3, on page 192](#)

Finding Feature Information

Your software release might not support all the features documented in this module. For the latest caveats and feature information, see the Bug Search Tool at <https://tools.cisco.com/bugsearch/> and the release notes for your software release. To find information about the features documented in this module, and to see a list of the releases in which each feature is supported, see the "New and Changed Information" chapter or the Feature History table in this chapter.

Information About OSPFv3

OSPFv3 is an IETF link-state protocol. An OSPFv3 router sends a special message, called a hello packet, out each OSPF-enabled interface to discover other OSPFv3 neighbor routers. Once a neighbor is discovered, the two routers compare information in the Hello packet to determine if the routers have compatible configurations. The neighbor routers attempt to establish adjacency, which means that the routers synchronize their link-state databases to ensure that they have identical OSPFv3 routing information. Adjacent routers share link-state advertisements (LSAs) that include information about the operational state of each link, the cost of the link,

and any other neighbor information. The routers then flood these received LSAs out every OSPF-enabled interface so that all OSPFv3 routers eventually have identical link-state databases. When all OSPFv3 routers have identical link-state databases, the network is converged. Each router then uses Dijkstra's Shortest Path First (SPF) algorithm to build its route table.

You can divide OSPFv3 networks into areas. Routers send most LSAs only within one area, which reduces the CPU and memory requirements for an OSPF-enabled router.

OSPFv3 supports IPv6.

Comparison of OSPFv3 and OSPFv2

Much of the OSPFv3 protocol is the same as in OSPFv2. OSPFv3 is described in RFC 2740.

The key differences between the OSPFv3 and OSPFv2 protocols are as follows:

- OSPFv3 expands on OSPFv2 to provide support for IPv6 routing prefixes and the larger size of IPv6 addresses.
- LSAs in OSPFv3 are expressed as prefix and prefix length instead of address and mask.
- The router ID and area ID are 32-bit numbers with no relationship to IPv6 addresses.
- OSPFv3 uses link-local IPv6 addresses for neighbor discovery and other features.
- OSPFv3 can use the IPv6 authentication trailer (RFC 6506) or IPsec (RFC 4552) for authentication. However, Cisco NX-OS does not support RFC 6506.
- OSPFv3 redefines LSA types.

Hello Packet

OSPFv3 routers periodically send Hello packets on every OSPF-enabled interface. The hello interval determines how frequently the router sends these Hello packets and is configured per interface. OSPFv3 uses Hello packets for the following tasks:

- Neighbor discovery
- Keepalives
- Bidirectional communications
- Designated router election

The Hello packet contains information about the originating OSPFv3 interface and router, including the assigned OSPFv3 cost of the link, the hello interval, and optional capabilities of the originating router. An OSPFv3 interface that receives these Hello packets determines if the settings are compatible with the receiving interface settings. Compatible interfaces are considered neighbors and are added to the neighbor table.

Hello packets also include a list of router IDs for the routers that the originating interface has communicated with. If the receiving interface sees its own router ID in this list, then bidirectional communication has been established between the two interfaces.

OSPFv3 uses Hello packets as a keepalive message to determine if a neighbor is still communicating. If a router does not receive a Hello packet by the configured dead interval (usually a multiple of the hello interval), then the neighbor is removed from the local neighbor table.

Neighbors

An OSPFv3 interface must have a compatible configuration with a remote interface before the two can be considered neighbors. The two OSPFv3 interfaces must match the following criteria:

- Hello interval
- Dead interval
- Area ID
- Optional capabilities

If there is a match, the information is entered into the neighbor table:

- If there is a match, the information is entered into the neighbor table:
- Priority—Priority of the neighbor router. The priority is used for designated router election.
- State—Indication of whether the neighbor has just been heard from, is in the process of setting up bidirectional communications, is sharing the link-state information, or has achieved full adjacency.
- Dead time—Indication of how long since the last Hello packet was received from this neighbor.
- Link-local IPv6 Address—The link-local IPv6 address of the neighbor.
- Designated Router—Indication of whether the neighbor has been declared the designated router or backup designated router.
- Local interface—The local interface that received the Hello packet for this neighbor.

When the first Hello packet is received from a new neighbor, the neighbor is entered into the neighbor table in the initialization state. Once bidirectional communication is established, the neighbor state becomes two-way. ExStart and exchange states come next, as the two interfaces exchange their link-state database. Once this is all complete, the neighbor moves into the full state, which signifies full adjacency. If the neighbor fails to send any Hello packets in the dead interval, then the neighbor is moved to the down state and is no longer considered adjacent.

Adjacency

Not all neighbors establish adjacency. Depending on the network type and designated router establishment, some neighbors become fully adjacent and share LSAs with all their neighbors, while other neighbors do not.

Adjacency is established using Database Description packets, Link State Request packets, and Link State Update packets in OSPFv3. The Database Description packet includes the LSA headers from the link-state database of the neighbor. The local router compares these headers with its own link-state database and determines which LSAs are new or updated. The local router sends a Link State Request packet for each LSA that it needs new or updated information on. The neighbor responds with a Link State Update packet. This exchange continues until both routers have the same link-state information.

Designated Routers

Networks with multiple routers present a unique situation for OSPFv3. If every router floods the network with LSAs, the same link-state information is sent from multiple sources. Depending on the type of network,

OSPFv3 might use a single router, the designated router (DR), to control the LSA floods and represent the network to the rest of the OSPFv3 area. If the DR fails, OSPFv3 uses the BDR.

Network types are as follows:


- Point-to-point—A network that exists only between two routers. All neighbors on a point-to-point network establish adjacency and there is no DR.
- Broadcast—A network with multiple routers that can communicate over a shared medium that allows broadcast traffic, such as Ethernet. OSPFv3 routers establish a DR and BDR that controls LSA flooding on the network. OSPFv3 uses the well-known IPv6 multicast addresses, FF02::5, and a MAC address of 0100.5300.0005 to communicate with neighbors.

The DR and BDR are selected based on the information in the Hello packet. When an interface sends a Hello packet, it sets the priority field and the DR and BDR field if it knows who the DR and BDR are. The routers follow an election procedure based on which routers declare themselves in the DR and BDR fields and the priority field in the Hello packet. As a final determinant, OSPFv3 chooses the highest router IDs as the DR and BDR.

All other routers establish adjacency with the DR and the BDR and use the IPv6 multicast address FF02::6 to send LSA updates to the DR and BDR. The Figure shows this adjacency relationship between all routers and the DR.

DRs are based on a router interface. A router might be the DR for one network and not for another network on a different interface.

Figure 27: DR in Multi-Access Network




Areas

You can limit the CPU and memory requirements that OSPFv3 puts on the routers by dividing an OSPFv3 network into areas. An area is a logical division of routers and links within an OSPFv3 domain that creates separate subdomains. LSA flooding is contained within an area, and the link-state database is limited to links within the area. You can assign an area ID to the interfaces within the defined area. The Area ID is a 32-bit value that can be expressed as a number or in dotted decimal notation, such as 10.2.3.1.

Cisco NX-OS always displays the area in dotted decimal notation.

If you define more than one area in an OSPFv3 network, you must also define the backbone area, which has the reserved area ID of 0. If you have more than one area, then one or more routers become area border routers (ABRs). An ABR connects to both the backbone area and at least one other defined area.

Figure 28: OSPFv3 Areas



The ABR has a separate link-state database for each area which it connects to. The ABR sends Inter-Area Prefix (type 3) LSAs from one connected area to the backbone area. The backbone area sends summarized information about one area to another area. In the figure, Area 0 sends summarized information about Area 5 to Area 3.

OSPFv3 defines one other router type: the autonomous system boundary router (ASBR). This router connects an OSPFv3 area to another autonomous system. An autonomous system is a network controlled by a single technical administration entity. OSPFv3 can redistribute its routing information into another autonomous system or receive redistributed routes from another autonomous system.

Link-State Advertisement Types

OSPFv3 uses link-state advertisements (LSAs) to build its routing table.

	Names	Description
1	Router LSA	LSA sent by every router. This LSA includes the state and cost of all links but does not include prefix information. Router LSAs trigger an SPF recalculation. Router LSAs are flooded to the local OSPFv3 area.
2	Network LSA	LSA sent by the DR. This LSA lists all routers in the multi-access network but does not include prefix information. Network LSAs trigger an SPF recalculation.

	Names	Description
3	Inter-Area Prefix LSA	LSA sent by the area border router to an external area for each destination in local area. This LSA includes the link cost from the border router to the local destination.
4	Inter-Area Router LSA	LSA sent by the area border router to an external area. This LSA advertises the link cost to the ASBR only.
5	AS External LSA	LSA generated by the ASBR. This LSA includes the link cost to an external autonomous system destination. AS External LSAs are flooded throughout the autonomous system.
7	Type-7 LSA	LSA generated by the ASBR within an NSSA. This LSA includes the link cost to an external autonomous system destination. Type-7 LSAs are flooded only within the local NSSA.
8	Link LSA	LSA sent by every router, using a link-local flooding scope. This LSA includes the link-local address and IPv6 prefixes for this link.
9	Intra-Area Prefix LSA	LSA sent by every router. This LSA includes any prefix or link state changes. Intra-Area Prefix LSAs are flooded to the local OSPFv3 area. This LSA does not trigger an SPF recalculation.

	Names	Description
11	Grace LSAs	LSA sent by a restarting router, using a link-local flooding scope. This LSA is used for a graceful restart of OSPFv3.

Link Cost

Each OSPFv3 interface is assigned a link cost. The cost is an arbitrary number. By default, Cisco NX-OS assigns a cost that is the configured reference bandwidth divided by the interface bandwidth. By default, the reference bandwidth is 40 Gb/s. The link cost is carried in the LSA updates for each link.

Flooding and LSA Group Pacing

OSPFv3 floods LSA updates to different sections of the network, depending on the LSA type. OSPFv3 uses the following flooding scopes:

- Link-local—LSA is flooded only on the local link. Used for Link LSAs and Grace LSAs.
- Area-local—LSA is flooded throughout a single OSPF area only. Used for Router LSAs, Network LSAs, Inter-Area-Prefix LSAs, Inter-Area-Router LSAs, and Intra-Area-Prefix LSAs.
- AS scope—LSA is flooded throughout the routing domain. An AS scope is used for AS External LSAs.

LSA flooding guarantees that all routers in the network have identical routing information. LSA flooding depends on the OSPFv3 area configuration. The LSAs are flooded based on the link-state refresh time (every 30 minutes by default). Each LSA has its own link-state refresh time.

You can control the flooding rate of LSA updates in your network by using the LSA group pacing feature. LSA group pacing can reduce high CPU or buffer utilization. This feature groups LSAs with similar link-state refresh times to allow OSPFv3 to pack multiple LSAs into an OSPFv3 Update message.

By default, LSAs with link-state refresh times within 10 seconds of each other are grouped together. You should lower this value for large link-state databases or raise it for smaller databases to optimize the OSPFv3 load on your network.

Link-State Database

Each router maintains a link-state database for the OSPFv3 network. This database contains all the collected LSAs and includes information on all the routes through the network. OSPFv3 uses this information to calculate the best path to each destination and populates the routing table with these best paths.

LSAs are removed from the link-state database if no LSA update has been received within a set interval, called the MaxAge. Routers flood a repeat of the LSA every 30 minutes to prevent accurate link-state information from being aged out. Cisco NX-OS supports the LSA grouping feature to prevent all LSAs from refreshing at the same time.

Multi-Area Adjacency

OSPFv3 multi-area adjacency allows you to configure a link on the primary interface that is in more than one area. This link becomes the preferred intra-area link in those areas. Multi-area adjacency establishes a point-to-point unnumbered link in an OSPFv3 area that provides a topological path for that area. The primary

adjacency uses the link to advertise an unnumbered point-to-point link in the Router LSA for the corresponding area when the neighbor state is full.

The multi-area interface exists as a logical construct over an existing primary interface for OSPF; however, the neighbor state on the primary interface is independent of the multi-area interface. The multi-area interface establishes a neighbor relationship with the corresponding multi-area interface on the neighboring router.

OSPFv3 and the IPv6 Unicast RIB

OSPFv3 runs the Dijkstra shortest path first algorithm on the link-state database. This algorithm selects the best path to each destination based on the sum of all the link costs for each link in the path. The shortest path for each destination is then put in the OSPFv3 route table. When the OSPFv3 network is converged, this route table feeds into the IPv6 unicast RIB. OSPFv3 communicates with the IPv6 unicast RIB to do the following:

- Add or remove routes
- Handle route redistribution from other protocols
- Provide convergence updates to remove stale OSPFv3 routes and for stub router advertisements.

OSPFv3 also runs a modified Dijkstra algorithm for fast recalculation for Inter-Area Prefix, Inter-Area Router, AS-External, type-7, and Intra-Area Prefix (type 3, 4, 5, 7, 8) LSA changes.

Address Family Support

Cisco NX-OS supports multiple address families, such as unicast IPv6 and multicast IPv6. OSPFv3 features that are specific to an address family are as follows:

- Default routes
- Route summarization
- Route redistribution
- Filter lists for border routers
- SPF optimization

Use the **address-family ipv6 unicast** command to enter the IPv6 unicast address family configuration mode when configuring these features.

Authentication

You can configure authentication on OSPFv3 messages to prevent unauthorized or invalid routing updates in your network.

RFC 4552 provides authentication to OSPFv3 using an IPv6 authentication header (AH) or encapsulating security payload (ESP) extension header. Cisco NX-OS partially supports RFC 4552 by using the IPv6 AH header to authenticate OSPFv3 packets.

Cisco NX-OS supports the IP security (IPSec) authentication method and the message digest 5 (MD5) or secure hash algorithm 1 (SHA1) algorithm to authenticate OSPFv3 packets. OSPFv3 IPSec authentication supports only static keys.

You can configure IPSec authentication for an OSPFv3 process, area, or interface.

Encryption

Beginning from Cisco Nexus Release 8.4.4, you can encrypt and authenticate OSPFv3 messages. OSPFv3 depends on IPsec for secure connection. IPsec supports two encapsulation types:

- Authentication Header (AH)
- Encapsulating Security Payload (ESP)

ESP configuration provides both encryption and authentication for OSPFv3 messages.

You can configure ESP at the following levels:

- Router
- Area
- Interface
- Virtual Links

Guidelines and Limitations for configuring ESP on OSPFv3

ESP configuration has the following guidelines and limitations:

- ESP configuration supports IPsec Transport Mode only.
- You can configure ESP on OSPFv3 for one SPI at one level, cannot configure two SPIs in one level.
- You cannot configure both encryption and authentication configurations for a same level.
- Supported encryption algorithms in ESP:
 - AES-CBC (128-bit)
 - 3DES-CBC
 - NULL
- Supported authentication algorithms in ESP:
 - SHA-1
 - NULL
- You cannot configure both ESP and AUTH algorithm as null in one ESP CLI.
- If ESP is not configured at local level, it inherits configuration from higher level, if configured:
 - If ESP is not configured at interface level, it inherits configuration from area level.
 - If ESP is not configured at area level, it inherits configuration from router level.
- On local level SPI, inherited data will be removed internally.



Note Ensure that the CoPP policy is customized to allow ESP packets, as default CoPP policy drops ESP packets.

Advanced Features

Cisco NX-OS supports advanced OSPFv3 features that enhance the usability and scalability of OSPFv3 in the network.


Stub Area

You can limit the amount of external routing information that floods an area by making it a stub area. A stub area is an area that does not allow AS External (type 5) LSAs. These LSAs are usually flooded throughout the local autonomous system to propagate external route information. Stub areas have the following requirements:

- All routers in the stub area are stub routers.
- No ASBR routers exist in the stub area.
- You cannot configure virtual links in the stub area.

The figure shows an example an OSPFv3 autonomous system where all routers in area 0.0.0.10 have to go through the ABR to reach external autonomous systems. Area 0.0.0.10 can be configured as a stub area.

Figure 29: Stub Area



Stub areas use a default route for all traffic that needs to go through the backbone area to the external autonomous system. The default route is an Inter-Area-Prefix LSA with the prefix length set to 0 for IPv6.

Not-So-Stubby Area

A Not-So-Stubby Area (NSSA) is similar to the stub area, except that an NSSA allows you to import autonomous system external routes within an NSSA using redistribution. The NSSA ASBR redistributes these routes and generates type-7 LSAs that it floods throughout the NSSA. You can optionally configure the ABR that connects the NSSA to other areas to translate this type-7 LSA to AS External (type 5) LSAs. The ABR then floods these AS External LSAs throughout the OSPFv3 autonomous system. Summarization and filtering are supported during the translation.

You can, for example, use NSSA to simplify administration if you are connecting a central site using OSPFv3 to a remote site that is using a different routing protocol. Before NSSA, the connection between the corporate site border router and a remote router could not be run as an OSPFv3 stub area because routes for the remote site could not be redistributed into a stub area. With NSSA, you can extend OSPFv3 to cover the remote connection by defining the area between the corporate router and remote router as an NSSA.

The backbone Area 0 cannot be an NSSA



Note OSPF became compliant with RFC 3101 section 2.5(3). When an Area Border Router attached to a Not-so-Stubby Area receives a default route LSA with P-bit clear, it should be ignored. OSPF had been previously adding the default route under these conditions.


If you have already designed your networks with RFC non-compliant behavior and expect a default route to be added on NSSA ABR, you will see a change in behavior when you upgrade.

If you decide to continue with the old behavior, you have the option to enable it with the **default-route nssa-abr pbit-clear** command.

Virtual Links

Virtual links allow you to connect an OSPFv3 area ABR to a backbone area ABR when a direct physical connection is not available. The figure shows a virtual link that connects Area 3 to the backbone area through Area 5.

Figure 30: Virtual Links



You can also use virtual links to temporarily recover from a partitioned area, which occurs when a link within the area fails, isolating part of the area from reaching the designated ABR to the backbone area.

Route Redistribution

OSPFv3 can learn routes from other routing protocols by using route redistribution. You configure OSPFv3 to assign a link cost for these redistributed routes or a default link cost for all redistributed routes.

Route redistribution uses route maps to control which external routes are redistributed. You must configure a route map with the redistribution to control which routes are passed into OSPFv3. A route map allows you to filter routes based on attributes such as the destination, origination protocol, route type, route tag, and so on. You can use route maps to modify parameters in the AS External (type 5) and NSSA External (type 7) LSAs before these external routes are advertised in the local OSPFv3 autonomous system.

OSPFv3 sets the type-5 LSA's forwarding address as described below:

- If the next-hop for the route is an attached-route then the forwarding address is the next-hop address for that route.
- If the next-hop for the route is a recursive route and next-hop's next-hop is an attached route then the forwarding address is the next-hop's next-hop address.

Route Summarization

Because OSPFv3 shares all learned routes with every OSPF-enabled router, you might want to use route summarization to reduce the number of unique routes that are flooded to every OSPF-enabled router. Route summarization simplifies route tables by replacing more-specific addresses with an address that represents all the specific addresses. For example, you can replace 2010:11:22:0:1000::1 and 2010:11:22:0:2000:679:1 with one summary address, 2010:11:22::/32.

Typically, you would summarize at the boundaries of area border routers (ABRs). Although you could configure summarization between any two areas, it is better to summarize in the direction of the backbone so that the backbone receives all the aggregate addresses and injects them, already summarized, into other areas. The two types of summarization are as follows:

- Inter-area route summarization
- External route summarization

You configure inter-area route summarization on ABRs, summarizing routes between areas in the autonomous system. To take advantage of summarization, assign network numbers in areas in a contiguous way to be able to lump these addresses into one range.

External route summarization is specific to external routes that are injected into OSPFv3 using route redistribution. You should make sure that external ranges that are being summarized are contiguous. Summarizing overlapping ranges from two different routers could cause packets to be sent to the wrong destination. Configure external route summarization on ASBRs that are redistributing routes into OSPF.

When you configure a summary address, Cisco NX-OS automatically configures a discard route for the summary address to prevent routing black holes and route loops.

High Availability and Graceful Restart

Cisco NX-OS provides a multilevel high-availability architecture. OSPFv3 supports stateful restart, which is also referred to as non-stop routing (NSR). If OSPFv3 experiences problems, it attempts to restart from its previous run-time state. The neighbors do not register any neighbor event in this case. If the first restart is not successful and another problem occurs, OSPFv3 attempts a graceful restart.

A graceful restart, or non-stop forwarding (NSF), allows OSPFv3 to remain in the data forwarding path through a process restart. When OSPFv3 needs to perform a graceful restart, it sends a link-local Grace (type 11) LSA. This restarting OSPFv3 platform is called NSF capable.

The Grace LSA includes a grace period, which is a specified time that the neighbor OSPFv3 interfaces hold onto the LSAs from the restarting OSPFv3 interface. (Typically, OSPFv3 tears down the adjacency and discards all LSAs from a down or restarting OSPFv3 interface.) The participating neighbors, which are called NSF helpers, keep all LSAs that originate from the restarting OSPFv3 interface as if the interface was still adjacent.

When the restarting OSPFv3 interface is operational again, it rediscovers its neighbors, establishes adjacency, and starts sending its LSA updates again. At this point, the NSF helpers recognize that the graceful restart has finished.

Stateful restart is used in the following scenarios:

- First recovery attempt after the process experiences problems
- ISSU
- User-initiated switchover using the **system switchover** command

Graceful restart is used in the following scenarios:

- Second recovery attempt after the process experiences problems within a 4-minute interval
- Manual restart of the process using the **restart ospfv3** command
- Active supervisor removal
- Active supervisor reload using the **reload module active-sup** command

Multiple OSPFv3 Instances

Cisco NX-OS supports multiple instances of the OSPFv3 protocol. By default, every instance uses the same system router ID. You must manually configure the router ID for each instance if the instances are in the same OSPFv3 autonomous system.

The OSPFv3 header includes an instance ID field to identify that OSPFv3 packet for a particular OSPFv3 instance. You can assign the OSPFv3 instance. The interface drops all OSPFv3 packets that do not have a matching OSPFv3 instance ID in the packet header.

Cisco NX-OS allows only one OSPFv3 instance on an interface.

SPF Optimization

Cisco NX-OS optimizes the SPF algorithm in the following ways:

- Partial SPF for Network (type 2) LSAs, Inter-Area Prefix (type 3) LSAs, and AS External (type 5) LSAs—When there is a change on any of these LSAs, Cisco NX-OS performs a faster partial calculation rather than running the whole SPF calculation.
- SPF timers—You can configure different timers for controlling SPF calculations. These timers include exponential backoff for subsequent SPF calculations. The exponential backoff limits the CPU load of multiple SPF calculations.

Virtualization Support

OSPFv3 supports virtual routing and forwarding (VRF) instances. VRFs exist within virtual device contexts (VDCs). By default, Cisco NX-OS places you in the default VDC and default VRF unless you specifically configure another VDC and VRF. Each OSPFv3 instance can support multiple VRFs, up to the system limit. For more information, see the *Cisco Nexus 7000 Series NX-OS Virtual Device Context Configuration Guide*.

Prerequisites for OSPFv3

OSPFv3 has the following prerequisites:

- You must be familiar with routing fundamentals to configure OSPFv3.
- You must be logged on to the switch.
- You have configured at least one interface for IPv6 that is capable of communicating with a remote OSPFv3 neighbor.
- You have completed the OSPFv3 network strategy and planning for your network. For example, you must decide whether multiple areas are required.
- You have enabled OSPF.
- You are familiar with IPv6 addressing and basic configuration.

Guidelines and Limitations for OSPFv3

OSPFv3 has the following configuration guidelines and limitations:

- You can have up to four instances of OSPFv3 in a VDC.
- Before Cisco NX-OS Release 6.2(2), Bidirectional Forwarding Detection (BFD) was not supported for OSPFv3. In Cisco NX-OS Release 6.2(2) and later releases, BFD includes a client for OSPFv3.
- Cisco NX-OS displays areas in dotted decimal notation regardless of whether you enter the area in decimal or dotted decimal notation.
- MTU configured at interface level works in either the data plane or in the control plane but not at both planes at the same time.

When you configure MTU with a size lower than the supported size in data and control planes a few features that have minimum MTU requirements may not work in both the planes.

For example, MPLS VPN is supported in the data plane since this plane supports the MTU of 1500 bytes that the MPLS VPN requires. But control plane does not support MPLS VPN because this plane cannot handle the 1500-byte packets.

To make the configured MTU work in control plane for MPLS VPN, you need to manually configure the OSPF packet size (by using the **packet-size** *size* command) so that OSPF works on the control plane. This is applicable from Cisco NX-OS Release 8.3(2) onwards.

The **packet-size** *size* command is supported on the Ethernet, SVI, and GRE tunnel interfaces.

- If you configure OSPFv3 in a virtual port channel (vPC) environment, use the following timer commands in router configuration mode on the core switch to ensure fast OSPF convergence when a vPC peer link is shut down:

```
switch (config-router)# timers throttle spf 1 50 50
switch (config-router)# timers lsa-arrival 10
```

- The value of object OSPFv3 router ID differs from RFC 5643 for traps `ospfv3NbrRestartHelperStatusChange` and `ospfv3VirtNbrRestartHelperStatusChange`. As per the RFC

5643, the value of object OSPFv3 router ID should be the router ID of the originator of the trap. But the current implementation will provide the router ID of the neighbor for both `ospfv3NbrRestartHelperStatusChange` and `ospfv3VirtNbrRestartHelperStatusChange`.

- Only the first four OSPFv3 instances are supported with MPLS LDP and MPLS TE.
- In scaled scenarios, when the number of interfaces and link-state advertisements in an OSPFv3 process is large, the `snmp-walk` on OSPF MIB objects is expected to time out with a small-value timeout at the SNMP agent. If you observe a timeout on the querying SNMP agent while polling OSPF MIB objects, increase the timeout value on the polling SNMP agent.
- If there is a particular OSPFv3 prefix that is learnt through type-5 as well as type-7, and both have different forwarding addresses, then these two route types are not comparable as per RFC3101, Section 2.5, step 6(e). (This applies only if the same destination/cost/non-zero forwarding addresses are there). OSPF will therefore do ECMP with all available next-hops.
- NXOS OSPF and U6RIB store only one route-type per route. If there is a mix of route-type across next-hops, only one of them, (the new path type) will be shown for all next hops.

Currently, route-type is a route property, and not a next-hop property.

- The **default-information-originate always** command advertises the OSPF default route from Cisco NX-OS Release 7.3(5)D1(1) and later releases and from Cisco NX-OS Release 8.0(1) and later releases in 8.x release train.
- The following guidelines and limitations apply to the administrative distance feature, which is supported beginning with Cisco NX-OS Release 6.1:
 - When an OSPF route has two or more equal cost paths, configuring the administrative distance is non-deterministic for the **match ip route-source** command.
 - For matching route sources in OSPFv3 routes, you must configure **match ip route-source** instead of **match ipv6 route-source** because the route sources and router IDs for OSPFv3 are IPv4 addresses.
 - Configuring the administrative distance is supported only for the **match route-type**, **match ipv6 address prefix-list**, and **match ip route-source prefix-list** commands. The other match statements are ignored.
 - The discard route is always assigned an administrative distance of 220. No configuration in the table map applies to OSPF discard routes.
 - There is no preference among the **match route-type**, **match ipv6 address**, and **match ip route-source** commands for setting the administrative distance of OSPF routes. In this way, the behavior of the table map for setting the administrative distance in Cisco NX-OS OSPF is different from that in Cisco IOS OSPF.
 - In Cisco NX-OS Release 6.2(6a) and later releases, you can filter next-hop paths for an OSPF route to prevent the path from being added to the RIB. Before Cisco NX-OS Release 6.2(6a), filtering on a specific path was ignored and the entire route was not added to the RIB.
- If you are familiar with the Cisco IOS CLI, be aware that the Cisco NX-OS commands for this feature might differ from the Cisco IOS commands that you would use.

Default Settings for OSPFv3

Table 17: Default OSPFv3 Parameters

Parameters	Default
Administrative distance	110
Hello interval	10 seconds
Dead interval	40 seconds
Discard routes	Enabled
Graceful restart grace period	60 seconds
Graceful restart notify period	15 seconds
OSPFv3 feature	Disabled
Stub router advertisement announce time	600 seconds
Reference bandwidth for link cost calculation	40 Gb/s
LSA minimal arrival time	1000 milliseconds
LSA group pacing	10 seconds
SPF calculation initial delay time	0 milliseconds
SPF calculation hold time	5000 milliseconds
SPF calculation initial delay time	0 milliseconds

Configuring Basic OSPFv3

Configure OSPFv3 after you have designed your OSPFv3 network.

Enabling OSPFv3

Before you begin

Ensure that you are in the correct VDC (or use the **switchto vdc** command).

Procedure

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.

	Command or Action	Purpose
Step 2	switch(config)# [no] feature ospfv3	Enables OSPFv3. To disable the OSPFv3 feature and remove all associated configurations, use the no form of the command.
Step 3	(Optional) switch(config)# show feature	Displays enabled and disabled features.
Step 4	(Optional) switch(config)# copy running-config startup-config	Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

Creating an OSPFv3 Instance

The first step in configuring OSPFv3 is to create an instance or OSPFv3 instance. You assign a unique instance tag for this OSPFv3 instance. The instance tag can be any string. For each OSPFv3 instance, you can also configure the following optional parameters:

- Router ID—Configures the router ID for this OSPFv3 instance. If you do not use this parameter, the router ID selection algorithm is used.
- Administrative distance—Rates the trustworthiness of a routing information source.
- Log adjacency changes—Creates a system message whenever an OSPFv3 neighbor changes its state.
- Name lookup—Translates OSPF router IDs to host names, either by looking up the local hosts database or querying DNS names in IPv6.
- Maximum paths—Sets the maximum number of equal paths that OSPFv3 installs in the route table for a particular destination. Use this parameter for load balancing between multiple paths.
- Reference bandwidth—Controls the calculated OSPFv3 cost metric for a network. The calculated cost is the reference bandwidth divided by the interface bandwidth. You can override the calculated cost by assigning a link cost when a network is added to the OSPFv3 instance.



Note The OSPF router ID changes without a restart on a Cisco Nexus 7000 series switch when you have not configured a manual router ID in the following cases:

- Configuring an SVI or physical interface with a higher IP address than the current router ID on a setup without any configured loopback interfaces.
- Configuring a loopback interface with any given IP address on a setup without any previously configured loopback interfaces.
- Configuring a loopback interface with a higher IP address than the IP address of an existing configured loopback interface.

When a router ID changes, OSPF has to re-advertise all LSAs with the new router ID. To avoid this issue, you need to configure a manual OSPF router ID.

Before you begin

You must enable OSPFv3.

Ensure that the OSPFv3 instance tag that you plan on using is not already in use on this router.

Use the **show ospfv3 instance-tag** command to verify that the instance tag is not in use.

OSPFv3 must be able to obtain a router identifier (for example, a configured loopback address) or you must configure the router ID option.

Ensure that you are in the correct VDC (or use the **switchto vdc** command).

Procedure

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# [no] router ospfv3 <i>instance-tag</i>	Creates a new OSPFv3 instance with the configured instance tag. Note The no router ospfv3 instance tag command does not remove OSPF configuration in interface mode. You must manually remove any OSPFv3 commands configured in interface mode.
Step 3	(Optional) switch(config-router)# router-id <i>ip-address</i>	Configures the OSPFv3 router ID. This ID uses the dotted decimal notation and identifies this OSPFv3 instance and must exist on a configured interface in the system. This command restarts the OSPF process automatically and changes the router id after it is configured.
Step 4	(Optional) switch(config-router)# show ipv6 ospfv3 instance-tag	Displays OSPFv3 information.
Step 5	(Optional) switch(config-router)# log-adjacency-changes [detail]	Generates a system message whenever a neighbor changes state.
Step 6	(Optional) switch(config-router)# passive-interface default	Suppresses routing updates on all interfaces. This command is overridden by the VRF or interface command mode configuration.
Step 7	(Optional) switch(config-router-af)# distance numbers	Configures the administrative distance for this OSPFv3 instance. The range is from 1 to 255. The default is 110.
Step 8	switch(config-router-af)# maximum-paths paths	Configures the maximum number of equal OSPFv3 paths to a destination in the route table. The range is from 1 to 16. The default is 8. This command is used for load balancing.

	Command or Action	Purpose
Step 9	(Optional) switch(config)# copy running-config startup-config	Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

Example

This example shows how to create an OSPFv3 instance:

```
switch# configure terminal
switch(config)# router ospfv3 201
switch(config-router)# copy running-config startup-config
```

Configuring OSPFv3 Packet Size

MTU configured at interface level works in either the data plane or in the control plane but not at both planes at the same time.

When you configure MTU with a size lower than the supported size in data and control planes a few features that have minimum MTU requirements may not work in both the planes.

For example, MPLS VPN is supported in the data plane since this plane supports the MTU of 1500 bytes that the MPLS VPN requires. But control plane does not support MPLS VPN because this plane cannot handle the 1500-byte packets.

To make the configured MTU work in control plane for MPLS VPN, you need to manually configure the OSPF packet size so that OSPF works on the control plane. This is applicable from Cisco NX-OS Release 8.3(2) onwards.

Procedure

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# [no] router ospfv3 instance-tag	Creates a new OSPFv3 instance with the configured instance tag. Note The no router ospfv3 instance-tag command does not remove OSPF configuration in interface mode. You must manually remove any OSPFv3 commands configured in interface mode.
Step 3	switch(config-router)# router-id ip-address	Configures the OSPFv3 router ID. This ID uses the dotted decimal notation and identifies this OSPFv3 instance and must exist on a configured interface in the system.

	Command or Action	Purpose
		This command restarts the OSPF process automatically and changes the router id after it is configured.
Step 4	switch(config-router)# ospfv3 packet-size size	<ul style="list-style-type: none"> Configures the OSPFv3 packet size. The size range is from 1280 to 9212 bytes. You can configure the packet-size in the interface configuration mode also. You can configure the packet-size size command even if the ip ospf mtu-ignore command is already configured in the network.
Step 5	(Optional) switch(config-router)# show ospfv3 interface	Displays OSPF information.

Example

This example shows how to configure the OSPFv3 packet-size:

```
router ospf 1
  router-id 3.3.3.3
  [no] packet-size 2000
```

This example shows the display of the configured OSPFv3 packet-size:

```
Switch (config-router)# show ospfv3 interface ethernet 1/25
Ethernet1/25 is up, line protocol is up
  IP address 1.0.0.1/24
  ----- snip -----
  Number of opaque link LSAs: 0, checksum sum 0
  Max Packet Size: 2000
```

Configuring Networks in OSPFv3

You can configure a network to OSPFv3 by associating it through the interface that the router uses to connect to that network. You can add all networks to the default backbone area (Area 0), or you can create new areas using any decimal number or an IP address.



Note All areas must connect to the backbone area either directly or through a virtual link.



Note OSPFv3 is not enabled on an interface until you configure a valid IPv6 address for that interface.

Before you begin

You must enable OSPFv3.

Ensure that you are in the correct VDC (or use the `switchto vdc` command).

Procedure

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# interface <i>interface-type</i> <i>slot/port</i>	Enters interface configuration mode.
Step 3	switch(config-if)# ipv6 address <i>ipv6-prefix/length</i>	Assigns an IPv6 address to this interface.
Step 4	switch(config-if)# ipv6 router ospfv3 <i>instance-tag area area-id [secondaries none]</i>	Adds the interface to the OSPFv3 instance and area.
Step 5	(Optional) switch(config-if)# show ipv6 ospfv3 <i>instance-tag interface interface-type</i> <i>slot/port</i>	Displays OSPFv3 information.
Step 6	(Optional) switch(config-if)# ospfv3 cost <i>number</i>	Configures the OSPFv3 cost metric for this interface. The default is to calculate a cost metric, based on the reference bandwidth and interface bandwidth. The range is from 1 to 65535.
Step 7	(Optional) switch(config-if)# ospfv3 dead-interval <i>seconds</i>	Configures the OSPFv3 dead interval, in seconds. The range is from 1 to 65535. The default is four times the hello interval, in seconds.
Step 8	(Optional) switch(config-if)# ospfv3 hello-interval <i>seconds</i>	Configures the OSPFv3 hello interval, in seconds. The range is from 1 to 65535. The default is 10 seconds.
Step 9	(Optional) switch(config-if)# ospfv3 instance <i>instance</i>	Configures the OSPFv3 instance ID. The range is from 0 to 255. The default is 0. The instance ID is link-local in scope.
Step 10	(Optional) switch(config-if)# ospfv3 mtu-ignore	Configures OSPFv3 to ignore any IP maximum transmission unit (MTU) mismatch with a neighbor. The default is to not establish adjacency if the neighbor MTU does not match the local interface MTU.
Step 11	(Optional) switch(config-if)# ospfv3 network { broadcast point-point }	Sets the OSPFv3 network type.
Step 12	(Optional) switch(config-if)# ospfv3 priority <i>number</i>	Configures the OSPFv3 priority, used to determine the DR for an area. The range is from 0 to 255. The default is 1.

	Command or Action	Purpose
Step 13	(Optional) switch(config-if)# ospfv3 shutdown	Shuts down the OSPFv3 instance on this interface.
Step 14	(Optional) switch(config)# copy running-config startup-config	Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

Example

This example shows how to add a network area 0.0.0.10 in OSPFv3 instance 201:

```
switch# configure terminal
switch(config)# interface ethernet 1/2
switch(config-if)# ipv6 address 2001:0DB8::1/48
switch(config-if)# ipv6 ospfv3 201 area 0.0.0.10
switch(config-if)# copy running-config startup-config
```

Configuring Advanced OSPFv3

Configure OSPFv3 after you have designed your OSPFv3 network.

Configuring Filter Lists for Border Routers

You can separate your OSPFv3 domain into a series of areas that contain related networks. All areas must connect to the backbone area through an area border router (ABR). OSPFv3 domains can connect to external domains as well through an autonomous system border router (ASBR).

ABRs have the following optional configuration parameters:

- Area range—Configures route summarization between areas.
- Filter list—Filters the Inter-Area Prefix (type 3) LSAs on an ABR that are allowed in from an external area.

ASBRs also support filter lists.

Before you begin

Create the route map that the filter list uses to filter IP prefixes in incoming or outgoing Inter-Area Prefix (type 3) LSAs.

You must enable OSPFv3.

Ensure that you are in the correct VDC (or use the **switchto vdc** command).

Procedure

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.

	Command or Action	Purpose
Step 2	switch(config)# router ospfv3 <i>instance-tag</i>	Creates a new OSPFv3 instance with the configured instance tag
Step 3	switch(config-router)# address-family ipv6 unicast	Enters IPv6 unicast address family mode.
Step 4	switch(config-router-af)# area <i>area-id</i> filter-list route-map <i>map-name</i> {in out}	Filters incoming or outgoing Inter-Area Prefix (type 3) LSAs on an ABR.
Step 5	(Optional) switch(config-if)# show ipv6 ospfv3 policy statistics <i>area id</i> filter-list {in out}	Displays OSPFv3 policy information.
Step 6	(Optional) switch(config)# copy running-config startup-config	Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

Example

This example shows how to enable graceful restart if it has been disabled:

```
switch# configure terminal
switch(config)# router ospfv3 201
switch(config-router)# address-family ipv6 unicast
switch(config-router-af)# area 0.0.0.10 filter-list route-map FilterLSAs in
switch(config-router-af)# copy running-config startup-config
```

Configuring Stub Areas

You can configure a stub area for part of an OSPFv3 domain where external traffic is not necessary. Stub areas block AS External (type 5) LSAs, limiting unnecessary routing to and from selected networks. You can optionally block all summary routes from going into the stub area.

Before you begin

You must enable OSPF.

Ensure that there are no virtual links or ASBRs in the proposed stub area.

Ensure that you are in the correct VDC (or use the **switchto vdc** command).

Procedure

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# router ospfv3 <i>instance-tag</i>	Creates a new OSPFv3 instance with the configured instance tag.
Step 3	switch(config-router)# area <i>area-id</i> stub	Creates this area as a stub area.

	Command or Action	Purpose
Step 4	(Optional) switch(config-router)# address-family ipv6 unicast	Enters IPv6 unicast address family mode.
Step 5	(Optional) switch(config-router-af)# area <i>area-id</i> default cost <i>cost</i>	Sets the cost metric for the default summary route sent into this stub area. The range is from 0 to 16777215.
Step 6	(Optional) switch(config)# copy running-config startup-config	Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

Example

This shows how to create a stub area that blocks all summary route updates:

```
switch# configure terminal
switch(config)# router ospfv3 201
switch(config-router)# area 0.0.0.10 stub no-summary
switch(config-router)# copy running-config startup-config
```

Configuring a Totally Stubby Area

You can create a totally stubby area and prevent all summary route updates from going into the stub area.

To create a totally stubby area, use the following command in router configuration mode:

Procedure

	Command or Action	Purpose
Step 1	switch(config-router)# area <i>area-id</i> stub no-summary	Creates this area as a totally stubby area.

Configuring NSSA

You can configure an NSSA for part of an OSPFv3 domain where limited external traffic is required. You can optionally translate this external traffic to an AS External (type 5) LSA and flood the OSPFv3 domain with this routing information. An NSSA can be configured with the following optional parameters:

- **No redistribution**—Redistributes routes that bypass the NSSA to other areas in the OSPFv3 autonomous system. Use this option when the NSSA ASBR is also an ABR.
- **Default information originate**—Generates a Type-7 LSA for a default route to the external autonomous system. Use this option on an NSSA ASBR if the ASBR contains the default route in the routing table. This option can be used on an NSSA ABR whether or not the ABR contains the default route in the routing table.
- **Route map**—Filters the external routes so that only those routes you want are flooded throughout the NSSA and other areas.

- **Translate**—Translates Type-7 LSAs to AS External (type 5) LSAs for areas outside the NSSA. Use this command on an NSSA ABR to flood the redistributed routes throughout the OSPFv3 autonomous system. You can optionally suppress the forwarding address in these AS External LSAs.
- **No summary**—Blocks all summary routes from flooding the NSSA. Use this option on the NSSA ABR.

Before you begin

You must enable OSPF.

Ensure that there are no virtual links in the proposed NSSA and that it is not the backbone area.

Ensure that you are in the correct VDC (or use the **switchto vdc** command).

Procedure

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# router ospfv3 instance-tag	Creates a new OSPFv3 instance with the configured instance tag.
Step 3	switch(config-router)# area area-id nssa [no-redistribution] [default-information-originate] [route-map map-name] [no-summary] [translate type7 {always never}] [suppress-fa]	Creates this area as an NSSA.
Step 4	(Optional) switch(config-router)# address-family ipv6 unicast	Enters IPv6 unicast address family mode.
Step 5	(Optional) switch(config-router-af)# area area-id default cost cost	Sets the cost metric for the default summary route sent into this NSSA. The range is from 0 to 16777215.
Step 6	(Optional) switch(config)# copy running-config startup-config	Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

Example

This example shows how to create an NSSA that blocks all summary route updates:

```
switch# configure terminal
switch(config)# router ospfv3 201
switch(config-router)# area 0.0.0.10 nssa no-summary
switch(config-router)# copy running-config startup-config
```

This example shows how to create an NSSA that generates a default route:

```
switch# configure terminal
switch(config)# router ospfv3 201
switch(config-router)# area 0.0.0.10 nssa default-info-originate
switch(config-router)# copy running-config startup-config
```

This example shows how to create an NSSA that filters external routes and blocks all summary route updates:

```
switch# configure terminal
switch(config)# router ospfv3 201
switch(config-router)# area 0.0.0.10 nssa route-map ExternalFilter no-summary
switch(config-router)# copy running-config startup-config
```

This example shows how to create an NSSA that always translates Type-7 LSAs to AS External (type 5) LSAs:

```
switch# configure terminal
switch(config)# router ospfv3 201
switch(config-router)# area 0.0.0.10 nssa translate type 7 always
switch(config-router)# copy running-config startup-config
```

This example shows how to create an NSSA that blocks all summary route updates:

```
switch# configure terminal
switch(config)# router ospfv3 201
switch(config-router)# area 0.0.0.10 nssa no-summary
switch(config-router)# copy running-config startup-config
```

Configuring Multi-Area Adjacency

You can add more than one area to an existing OSPFv3 interface. The additional logical interfaces support multi-area adjacency.

Before you begin

You must enable OSPF.

Ensure that you are in the correct VDC (or use the **switchto vdc** command).

Ensure that you have configured a primary area for the interface.

Procedure

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# interface <i>interface-type</i> <i>slot/port</i>	Enters interface configuration mode.
Step 3	switch(config-if)# ipv6 router ospfv3 <i>instance-tag</i> multi-area <i>area-id</i>	Adds the interface to another area.
Step 4	(Optional) switch(config-if)# show ipv6 ospfv3 <i>instance-tag</i> interface <i>interface-type</i> <i>slot/port</i>	Displays OSPFv3 information.
Step 5	(Optional) switch(config)# copy running-config startup-config	Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

Example

This example shows how to add a second area to an OSPFv3 interface:

```
switch# configure terminal
switch(config)# interface ethernet 1/2
switch(config-if)# ipv6 address 2001:0DB8::1/48
switch(config-if)# ipv6 ospfv3 201 area 0.0.0.10
switch(config-if)# ipv6 ospfv3 201 multi-area 20
switch(config-if)# copy running-config startup-config
```

Configuring Virtual Links

A virtual link connects an isolated area to the backbone area through an intermediate area. You can configure the following optional parameters for a virtual link:

- Dead interval—Sets the time that a neighbor waits for a Hello packet before declaring the local router as dead and tearing down adjacencies.
- Hello interval—Sets the time between successive Hello packets.
- Retransmit interval—Sets the estimated time between successive LSAs.
- Transmit delay—Sets the estimated time to transmit an LSA to a neighbor.



Note You must configure the virtual link on both routers involved before the link becomes active.

Before you begin

You must enable OSPF.

Ensure that you are in the correct VDC (or use the **switchto vdc** command).

Procedure

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# router ospfv3 <i>instance-tag</i>	Creates a new OSPFv3 instance with the configured instance tag.
Step 3	switch(config-router)# area <i>area-id</i> virtual-link <i>router-id</i>	Creates one end of a virtual link to a remote router. You must create the virtual link on that remote router to complete the link.
Step 4	(Optional) switch(config-if)# show ipv6 ospfv3 virtual-link [brief]	Displays OSPFv3 virtual link information.
Step 5	(Optional) switch(config-router-vlink)# dead-interval <i>seconds</i>	Configures the OSPFv3 dead interval, in seconds. The range is from 1 to 65535. The

	Command or Action	Purpose
		default is four times the hello interval, in seconds.
Step 6	(Optional) switch(config-router-vlink)# hello-interval <i>seconds</i>	Configures the OSPFv3 hello interval, in seconds. The range is from 1 to 65535. The default is 10 seconds.
Step 7	(Optional) switch(config-router-vlink)# retransmit-interval <i>seconds</i>	Configures the OSPFv3 retransmit interval, in seconds. The range is from 1 to 65535. The default is 5.
Step 8	(Optional) switch(config-router-vlink)# transmit-delay <i>seconds</i>	Configures the OSPFv3 transmit-delay, in seconds. The range is from 1 to 450. The default is 1.
Step 9	(Optional) switch(config)# copy running-config startup-config	Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

Example

These examples show how to create a simple virtual link between two ABRs:

Configuration for ABR 1 (router ID 2001:0DB8::1) is as follows:

```
switch# configure terminal
switch(config)# router ospfv3 201
switch(config-router)# area 0.0.0.10 virtual-link 2001:0DB8::10
switch(config-router)# copy running-config startup-config
```

Configuration for ABR 2 (router ID 2001:0DB8::10) is as follows:

```
switch# configure terminal
switch(config)# router ospfv3 201
switch(config-router)# area 0.0.0.10 virtual-link 2001:0DB8::1
switch(config-router)# copy running-config startup-config
```

Configuring Redistribution

You can redistribute routes learned from other routing protocols into an OSPFv3 autonomous system through the ASBR.

You can configure the following optional parameters for route redistribution in OSPF:

- Default information originate—Generates an AS External (type 5) LSA for a default route to the external autonomous system.



Note Default information originate ignores **match** statements in the optional route map.

- Default metric—Sets all redistributed routes to the same cost metric.



Note If you redistribute static routes, Cisco NX-OS also redistributes the default static route.

Before you begin

Create the necessary route maps used for redistribution.

You must enable OSPF.

Ensure that you are in the correct VDC (or use the **switchto vdc** command).

Procedure

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# router ospfv3 <i>instance-tag</i>	Creates a new OSPFv3 instance with the configured instance tag.
Step 3	switch(config-router)# address-family ipv6 unicast	Enters IPv6 unicast address family mode.
Step 4	switch(config-router-af)# redistribute { <i>bgpid</i> direct isis id rip id static } route-map <i>map-name</i>	Redistributes the selected protocol into OSPFv3 through the configured route map. Note If you redistribute static routes, Cisco NX-OS also redistributes the default static route.
Step 5	switch(config-router-af)# default-information originate [always] [route-map <i>map-name</i>]	Creates a default route into this OSPFv3 domain if the default route exists in the RIB. Use the following optional keywords: <ul style="list-style-type: none"> • always —Always generates the default route of 0.0.0, even if the route does not exist in the RIB. • route-map—Generates the default route if the route map returns true. Note This command ignores match statements in the route map.
Step 6	switch(config-router-af)# default-metric <i>cost</i>	Sets the cost metric for the redistributed routes. The range is from 1 to 16777214. This command does not apply to directly connected routes. Use a route map to set the default metric for directly connected routes.

	Command or Action	Purpose
Step 7	(Optional) switch(config)# copy running-config startup-config	Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

Example

This example shows how to redistribute the Border Gateway Protocol (BGP) into OSPFv3:

```
switch# configure terminal
switch(config)# router ospfv3 201
switch(config-router)# address-family ipv6 unicast
switch(config-router-af)# redistribute bgp route-map FilterExternalBGP
switch(config-router-af)# copy running-config startup-config
```

Limiting the Number of Redistributed Routes

Route redistribution can add many routes to the OSPFv3 route table. You can configure a maximum limit to the number of routes accepted from external protocols. OSPFv3 provides the following options to configure redistributed route limits:

- **Fixed limit**—Logs a message when OSPFv3 reaches the configured maximum. OSPFv3 does not accept any more redistributed routes. You can optionally configure a threshold percentage of the maximum where OSPFv3 logs a warning when that threshold is passed.
- **Warning only**—Logs a warning only when OSPFv3 reaches the maximum. OSPFv3 continues to accept redistributed routes.
- **Withdraw**—Starts the configured timeout period when OSPFv3 reaches the maximum. After the timeout period, OSPFv3 requests all redistributed routes if the current number of redistributed routes is less than the maximum limit. If the current number of redistributed routes is at the maximum limit, OSPFv3 withdraws all redistributed routes. You must clear this condition before OSPFv3 accepts more redistributed routes. You can optionally configure the timeout period.

Before you begin

You must enable OSPF.

Ensure that you are in the correct VDC (or use the **switchto vdc** command).

Procedure

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# router ospfv3 instance-tag	Creates a new OSPFv3 instance with the configured instance tag.
Step 3	switch(config-router)# address-family ipv6 unicast	Enters IPv6 unicast address family mode.

	Command or Action	Purpose
Step 4	switch(config-router)# redistribute { <i>bgpid</i> direct isis <i>id</i> rip <i>id</i> static } route-map <i>map-name</i>	Redistributes the selected protocol into OSPFv3 through the configured route map.
Step 5	switch(config-router)# redistribute maximum-prefix <i>max</i> [<i>threshold</i>] [warning-only withdraw [<i>num-retries</i> <i>timeout</i>]]	Specifies a maximum number of prefixes that OSPFv2 distributes. The range is from 0 to 65536. Optionally, specifies the following: <ul style="list-style-type: none"> • threshold—Percent of maximum prefixes that triggers a warning message. • warning-only—Logs a warning message when the maximum number of prefixes is exceeded. • withdraw—Withdraws all redistributed routes and optionally tries to retrieve the redistributed routes. The num-retries range is from 1 to 12. The timeout range is from 60 to 600 seconds. The default is 300 seconds.
Step 6	(Optional) show running-config ospfv3 Example: switch(config-router)# show running-config ospf	Displays the OSPFv3 configuration.
Step 7	(Optional) switch(config)# copy running-config startup-config	Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

Example

This example shows how to limit the number of redistributed routes into OSPF:

```
switch# configure terminal
switch(config)# router ospfv3 201
switch(config-router)# address-family ipv6 unicast
switch(config-router-af)# redistribute bgp route-map FilterExternalBGP
switch(config-router-af)# copy running-config startup-config
```

Configuring Route Summarization

You can configure route summarization for inter-area routes by configuring an address range that is summarized. You can also configure route summarization for external, redistributed routes by configuring a summary address for those routes on an ASBR.

Before you begin

You must enable OSPF.

Ensure that you are in the correct VDC (or use the **switchto vdc** command).

Procedure

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# router ospfv3 instance-tag	Creates a new OSPFv3 instance with the configured instance tag.
Step 3	switch(config-router)# address-family ipv6 unicast	Enters IPv6 unicast address family mode.
Step 4	switch(config-router-af)# area area-id range ipv6-prefix/length [no-advertise] [cost cost]	Creates a summary address on an ABR for a range of addresses and optionally advertises this summary address in a Inter-Area Prefix (type 3) LSA. The cost range is from 0 to 16777215.
Step 5	switch(config-router-af)# summary-address ipv6-prefix/length [no-advertise] [tag tag]	Creates a summary address on an ASBR for a range of addresses and optionally assigns a tag for this summary address that can be used for redistribution with route maps.
Step 6	(Optional) switch(config-router)# show ipv6 ospfv3 summary-address	Displays information about OSPFv3 summary addresses.
Step 7	(Optional) switch(config)# copy running-config startup-config	Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

Example

This shows how to create a stub area that blocks all summary route updates:

```
switch# configure terminal
switch(config)# router ospfv3 201
switch(config-router)# address-family ipv6 unicast
switch(config-router)# area 0.0.0.10 range 2001:0DB8::/48
switch(config-router)# copy running-config startup-config
```

This example shows how to create summary addresses on an ASBR:

```
switch# configure terminal
switch(config)# router ospfv3 201
switch(config-router)# address-family ipv6 unicast
switch(config-router)# summary-address 2001:0DB8::/48
switch(config-router)# copy running-config startup-config
```

Configuring the Administrative Distance of Routes

Beginning with Cisco NX-OS Release 6.1, you can set the administrative distance of routes added by OSPFv3 into the RIB.

The administrative distance is a rating of the trustworthiness of a routing information source. A higher value indicates a lower trust rating. Typically, a route can be learned through more than one routing protocol. The administrative distance is used to discriminate between routes learned from more than one routing protocol. The route with the lowest administrative distance is installed in the IP routing table.

Before you begin

Ensure that you have enabled OSPFv3.

Ensure that you are in the correct VDC (or use the **switchto vdc** command).

See the guidelines and limitations for this feature.

Procedure

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# router ospf instance-tag	Creates a new OSPFv3 instance with the configured instance tag.
Step 3	switch(config-router)# address-family ipv6 unicast	Enters IPv6 unicast address family mode.
Step 4	switch(config-router-af)# [no] table-map map-name [filter]	Configures the policy for filtering or modifying OSPFv2 routes before sending them to the RIB. You can enter up to 63 alphanumeric characters for the map name. The filter keyword specifies that only routes that are permitted by the route map(<i>map-name</i>) configuration are downloaded to the routing information base (RIB).
Step 5	switch(config-router-af)# exit	Exits router address-family configuration mode.
Step 6	switch(config-router)# exit	Exits router configuration mode.
Step 7	switch(config)# route-map map-name [permit deny] [seq]	Creates a route map or enters route-map configuration mode for an existing route map. Use <i>seq</i> to order the entries in a route map. Note The permit option enables you to set the distance. If you use the deny option, the default distance is applied.
Step 8	switch(config-route-map)# match route-type route-type	Matches against one of the following route types: <ul style="list-style-type: none"> external—The external route (BGP, EIGRP, and OSPF type 1 or 2) inter-area—OSPF inter-area route

	Command or Action	Purpose
		<ul style="list-style-type: none"> • internal—The internal route (including the OSPF intra- or inter-area) • intra-area—OSPF intra-area route • nssa-external—The NSSA external route (OSPF type 1 or 2) • type-1—The OSPF external type 1 route • type-2—The OSPF external type 2 route
Step 9	switch(config-route-map)# match ip route-source prefix-list name	Matches the IPv6 route source address or router ID of a route to one or more IP prefix lists. Use the ip prefix-list command to create the prefix list. Note For OSPFv3, the router ID is 4 bytes.
Step 10	switch(config-route-map)# match ipv6 address prefix-list name	Matches against one or more IPv6 prefix lists. Use the ip prefix-list command to create the prefix list.
Step 11	switch(config-route-map)# set distance value	Sets the administrative distance of routes for OSPFv3. The range is from 1 to 255.
Step 12	(Optional) switch(config)# copy running-config startup-config	Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

Example

This example shows how to configure the OSPFv3 administrative distance for inter-area routes to 150, for external routes to 200, and for all prefixes in prefix list p1 to 190:

```
switch# configure terminal
switch(config)# router ospfv3 201
switch(config-router)# address-family ipv6 unicast
switch(config-router-af)# table-map foo
switch(config-router)# exit
switch(config)# exit
switch(config)# route-map foo permit 10
switch(config-route-map)# match route-type inter-area
switch(config-route-map)# set distance 150
switch(config)# route-map foo permit 20
switch(config-route-map)# match route-type external
switch(config-route-map)# set distance 200
switch(config)# route-map foo permit 30
switch(config-route-map)# match ip route-source prefix-list p1
switch(config-route-map)# match ipv6 address prefix-list p1
switch(config-route-map)# set distance 190
```

The following example shows how to configure a route map for blocking the next hops that are learned through VLAN 10:

```
switch(config)# route-map Filter-OSPF 10 deny
switch(config-route-map)# match interface VLAN 10
switch(config-route-map)# exit
switch(config)# route-map Filter-OSPF 20 permit
```

The following example shows how to configure the **table-map** command with the **filter** keyword to use a route map (Filter-OSPF) to remove the next-hop path that is learned through VLAN 10 but not the next-hop path that is learned through VLAN 20:

```
switch(config)# route ospfv3 p1
switch(config-router)# table-map Filter-OSPF filter
```

Modifying the Default Timers

OSPFv3 includes a number of timers that control the behavior of protocol messages and shortest path first (SPF) calculations. OSPFv3 includes the following optional timer parameters:

- LSA arrival time—Sets the minimum interval allowed between LSAs arriving from a neighbor. LSAs that arrive faster than this time are dropped.
- Pacing LSAs—Sets the interval at which LSAs are collected into a group and refreshed, checksummed, or aged. This timer controls how frequently LSA updates occur and optimizes how many are sent in an LSA update message.
- Throttle LSAs—Sets rate limits for generating LSAs. This timer controls how frequently LSAs are generated after a topology change occurs.
- Throttle SPF calculation—Controls how frequently the SPF calculation is run.

At the interface level, you can also control the following timers:

- Retransmit interval—Sets the estimated time between successive LSAs
- Transmit delay—Sets the estimated time to transmit an LSA to a neighbor.

Before you begin

Ensure that you are in the correct VDC (or use the **switchto vdc** command).

Procedure

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# router ospfv3 instance-tag	Creates a new OSPFv3 instance with the configured instance tag.
Step 3	switch(config-router)# timers lsa-arrival msec	Sets the LSA arrival time in milliseconds. The range is from 10 to 600000. The default is 1000 milliseconds.

	Command or Action	Purpose
Step 4	switch(config-router)# timers lsa-group-pacing <i>seconds</i>	Sets the interval in seconds for grouping LSAs. The range is from 1 to 1800. The default is 10 seconds.
Step 5	switch(config-router)# timers throttle lsa <i>start-time hold-interval max-time</i>	Sets the rate limit in milliseconds for generating LSAs. You can configure the following timers: <ul style="list-style-type: none"> • <i>start-time</i>—The range is from 50 to 5000 milliseconds. The default value is 50 milliseconds. • <i>hold-interval</i>—The range is from 50 to 30,000 milliseconds. The default value is 5000 milliseconds. • <i>max-time</i>—The range is from 50 to 30,000 milliseconds. The default value is 5000 milliseconds.
Step 6	switch(config-router)# address-family ipv6 unicast	Enters IPv6 unicast address family mode.
Step 7	switch(config-router)# timers throttle spf <i>delay-time hold-time</i>	Sets the SPF best path schedule initial delay time and the minimum hold time in seconds between SPF best path calculations. The range is from 1 to 600000. The default is no delay time and 5000 millisecond hold time.
Step 8	switch(config)# interface <i>type slot/port</i>	Enters interface configuration mode.
Step 9	switch(config-if)# ospfv3 retransmit-interval <i>seconds</i>	Sets the estimated time in seconds between LSAs transmitted from this interface. The range is from 1 to 65535. The default is 5.
Step 10	switch(config-if)# ospfv3 transmit-delay <i>seconds</i>	Sets the estimated time in seconds to transmit an LSA to a neighbor. The range is from 1 to 450. The default is 1.
Step 11	(Optional) switch(config)# copy running-config startup-config	Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

Example

This shows how to create a stub area that blocks all summary route updates:

```
switch# configure terminal
switch(config)# router ospfv3 201
switch(config-router)# timers lsa-group-pacing 300
switch(config-router)# copy running-config startup-config
```

Configuring the OSPFv3 Max-Metric Router LSA

You can configure OSPFv3 to advertise its locally generated router LSAs with the maximum metric value possible (the infinity metric 0xFFFF). This feature allows OSPFv3 processes to converge but not attract transit traffic through the device if there are better alternate paths. After a specified timeout or a notification from BGP, OSPFv3 advertises the LSAs with normal metrics.

Before you begin

Ensure that you are in the correct VDC (or use the **switchto vdc** command).

Procedure

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# router ospfv3 instance-tag	Creates a new OSPFv3 instance with the configured instance tag.
Step 3	switch(config-router)# max-metric router-lsa [external-lsa <i>max-metric-value</i>] [stub-prefix-lsa] [on-startup <i>seconds</i>] [wait-for bgp tag] [inter-area-prefix-lsa <i>max-metric-value</i>]	Configures a device that is running the OSPFv3 protocol to advertise a maximum metric so that other devices do not prefer the device as an intermediate hop in their SPF calculations.
Step 4	(Optional) switch(config)# copy running-config startup-config	Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

Example

This example shows how to configure a router to advertise a maximum metric for the stub links:

```
switch(config)# router ospfv3 200
switch(config-router)# max-metric router-lsa stub-prefix-lsa
```

Configuring Graceful Restart

Graceful restart is enabled by default. You can configure the following optional parameters for graceful restart in an OSPFv3 instance:

- Grace period—Configures how long neighbors should wait after a graceful restart has started before tearing down adjacencies.
- Helper mode disabled—Disables helper mode on the local OSPFv3 instance. OSPFv3 does not participate in the graceful restart of a neighbor.
- Planned graceful restart only—Configures OSPFv3 to support graceful restart only in the event of a planned restart.

Before you begin

You must enable OSPF.

Ensure that all neighbors are configured for graceful restart with matching optional parameters set.

Ensure that you are in the correct VDC (or use the **switchto vdc** command).

Procedure

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# router ospfv3 <i>instance-tag</i>	Creates a new OSPFv3 instance with the configured instance tag.
Step 3	switch(config-router)# graceful restart	Enables graceful restart. A graceful restart is enabled by default.
Step 4	switch(config-router)# graceful-restart grace-period <i>seconds</i>	Sets the grace period, in seconds. The range is from 5 to 1800. The default is 60 seconds.
Step 5	switch(config-router)# graceful-restart helper-disable	Disables helper mode. Enabled by default.
Step 6	switch(config-router)# graceful-restart planned-only	Configures graceful restart for planned restarts only.
Step 7	(Optional) switch(config-if)# show ipv6 ospfv3 <i>instance-tag</i>	Displays OSPFv3 information.
Step 8	(Optional) switch(config)# copy running-config startup-config	Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

Example

This shows how to create a stub area that blocks all summary route updates:

```
switch# configure terminal
switch(config)# router ospfv3 201
switch(config-router)# graceful-restart
switch(config-router)# graceful-restart grace-period 120
switch(config-router)# copy running-config startup-config
```

Restarting an OSPFv3 Instance

You can restart an OSPv3 instance. This action clears all neighbors for the instance.

To restart an OSPFv3 instance and remove all associated neighbors, use the following command:

Procedure

	Command or Action	Purpose
Step 1	switch(config)# restart ospfv3 <i>instance-tag</i>	Restarts the OSPFv3 instance and removes all neighbors.

Configuring OSPFv3 with Virtualization

You can configure multiple OSPFv3 instances in each VDC. You can also create multiple VRFs within each VDC and use the same or multiple OSPFv3 instances in each VRF. You assign an OSPFv3 interface to a VRF.



Note Configure all other parameters for an interface after you configure the VRF for an interface. Configuring a VRF for an interface deletes all the configuration for that interface.

Before you begin

Create the VDCs.

You must enable OSPF.

Ensure that you are in the correct VDC (or use the **switchto vdc** command).

Procedure

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# vrf context <i>vrf-name</i>	Creates a new VRF and enters VRF configuration mode.
Step 3	switch(config)# router ospfv3 <i>instance-tag</i>	Creates a new OSPFv3 instance with the configured instance tag.
Step 4	switch(config-router)# vrf <i>vrf-name</i>	Enters VRF configuration mode.
Step 5	(Optional) switch(config-router-vrf)# maximum-paths <i>paths</i>	Configures the maximum number of equal OSPFv3 paths to a destination in the route table for this VRF. Use this command for load balancing.
Step 6	switch(config)# interface <i>type slot/port</i>	Enters interface configuration mode.
Step 7	switch(config-if)# vrf member <i>vrf-name</i>	Adds this interface to a VRF.
Step 8	switch(config-if)# ipv6 address <i>ipv6-prefix/length</i>	Configures an IP address for this interface. You must do this step after you assign this interface to a VRF.

	Command or Action	Purpose
Step 9	switch(config-if)# ipv6 ospfv3 instance-tag area area-id	Assigns this interface to the OSPFv3 instance and area configured.
Step 10	(Optional) switch(config)# copy running-config startup-config	Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

Example

This example shows how to create a VRF and add an interface to the VRF:

```
switch# configure terminal
switch(config)# vrf context NewVRF
switch(config-vrf)# exit
switch(config)# router ospfv3 201
switch(config-router)# exit
switch(config)# interface ethernet 1/2
switch(config-if)# vrf member NewVRF
switch(config-if)# ipv6 address 2001:0DB8::1/48
switch(config-if)# ipv6 ospfv3 201 area 0
switch(config-if)# copy running-config startup-config
```

Configuring OSPFv3 Authentication at Router Level

You can enable authentication of OSPFv3 packets on a per-interface basis at the Router level using the following commands.

Before you begin

Ensure you have enabled OSPF.

Ensure that you are in the correct VDC(or use the **switchto vdc** command)

Enable the authentication package.

Procedure

-
- Step 1** Enter the global configuration mode:
- ```
switch# configure terminal
```
- Step 2** Enable the authentication package:
- ```
switch(config)# feature imp
```
- Step 3** Create a new OSPFv3 instance with the configured instance tag:
- ```
switch(config)# router ospfv3 instance-tag
```
- Step 4** Enable IPsec AH Authentication:
- ```
switch(config-router)# authentication ipsec spi spi auth [ 0 | 3 | 7 ] key
```

You can specify the security policy index through *spi* and define the authentication algorithm through *auth* which can be md5 or sha1. Numbers 0, 3 and 7 specify the format of *key*.

- Step 5** (Optional) Display OSPFv3 information:
switch(config)# **show running-config ospfv3**
-

Configuring OSPFv3 Authentication at Area Level

Authentication of OSPFv3 packets is enabled on a per-interface basis at the Area level using the following commands.

Before you begin

Ensure you have enabled OSPF.

Ensure that you are in the correct VDC(or use the **switchto vdc** command)

Enable the authentication package.

Procedure

- Step 1** Enter the global configuration mode:
switch# **configure terminal**
- Step 2** Enable the authentication package:
switch(config)# **feature imp**
- Step 3** Create a new OSPFv3 instance with the configured instance tag:
switch(config)#**router ospfv3** *instance-tag*
- Step 4** Enable IPsec AH Authentication:
switch(config-router)#**area** *area-num* **authentication ipsec spi** *spi* **auth** [**0** | **3** | **7**] *key*
You can specify the security policy index through *spi* and define the authentication algorithm through *auth* which can be md5 or sha1. Numbers 0, 3 and 7 specify the format of *key*.
- Step 5** (Optional) Display OSPFv3 information:
switch(config)# **show running-config ospfv3**
-

Configuring OSPFv3 Authentication at Interface Level

You can configure the authentication of OSPFv3 packets per interface using the following commands.

Before you begin

Ensure you have enabled OSPF.

Ensure that you are in the correct VDC(or use the **switchto vdc** command)

Enable the authentication package.

Procedure

-
- Step 1** Enter the global configuration mode:
switch# **configure terminal**
- Step 2** Enables the authentication mode:
switch(config)# **feature imp**
- Step 3** Enters the interface configuration mode:
switch(config)# **interface ethernet** *interface*
- Step 4** Change the port mode to Layer 3 interface:
switch(config-if)# **no switchport**
- Step 5** Specify the OSPFv3 instance and area for the interface:
switch(config-if)# **ipv6 router ospfv3** *instance-tag* **area** *area-id*
- Step 6** Enable IPsec AH Authentication:
switch(config-if)# **ospfv3 authentication ipsec spi** *spi* **auth** [0 | 3 | 7] *key*
- You can specify the security policy index through *spi* and define the authentication algorithm through *auth* which can be md5 or sha1. Numbers 0, 3 and 7 specify the format of *key*.
- Step 7** (Optional) Display the running configuration on the interface:
switch(config-if)#**show run interface** *interface*
-

Configuration Example

The following example shows how to enable security for Ethernet interface 2/1.

```
switch# configure terminal
switch(config)# interface ethernet 2/1
switch(config-if)# no switchport
switch(config-if)# ipv6 router ospfv3 1 area 0
switch(config-if)# ospfv3 authentication ipsec spi 256 md5 0 11111111111111111111111111111111
switch(config-if)# show run interface ethernet 2/1

!Command: show running-config interface Ethernet2/1
!Time: Mon Oct 26 09:19:30 2015

version 7.2(0)D1(1)

interface Ethernet2/1
```

```
shutdown
no switchport
medium p2p
ospfv3 authentication ipsec spi 256 md5 3 b54dc5a961fb42098f6902e512cb6e099d44
d3239f4e48e73668de6f52254f0e
ipv6 router ospfv3 1 area 0.0.0.0

switch(config-if)#
```

Configuring OSPFv3 Encryption at Router Level

You can configure OSPFv3 ESP to encrypt and authenticate OSPFv3 packets at the router level using the following commands.

Before you begin

You must enable OSPFv3.

Enable authentication package.

Procedure

-
- Step 1** Enter the global configuration mode:
- ```
switch# configure terminal
```
- Step 2** Enter the configuration of OSPFv3 mode:
- ```
switch# configure ospfv3
```
- Step 3** Enable authentication package:
- ```
switch(config)# feature imp
```
- Step 4** Create a new OSPFv3 instance with the configured instance tag:
- ```
switch(config)# router ospfv3 instance-tag
```
- Step 5** Enable IPsec ESP encryption:
- ```
switch(config-router)# encryption ipsec spi spi_id esp encrypt_algorithm [0 | 3 | 7] key authentication
auth_algorithm [0 | 3 | 7] key.
```
- You can specify the security policy index through *spi\_id* and define the encryption algorithm through *encrypt\_algorithm* which can be 3des, aes 128 or null. Numbers 0, 3, and 7 specify the format of the *key*. You can define the authentication algorithm through *auth\_algorithm* which can be sha1 or md5.
- Note** MD5 is not supported in FIPS mode.
- Step 6** (Optional) Display OSPFv3 information:
- ```
switch(config)# show running-config ospfv3
```
-

Configuring OSPFv3 Encryption at Area Level

You can configure OSPFv3 ESP to encrypt and authenticate OSPFv3 packets at the area level using the following commands.

Before you begin

You must enable OSPFv3.

Enable authentication package.

Procedure

-
- Step 1** Enter the global configuration mode:
switch# **configure terminal**
- Step 2** Enter the configuration of OSPFv3 mode:
switch# **configure ospfv3**
- Step 3** Enable the authentication package:
switch(config)# **feature imp**
- Step 4** Create a new OSPFv3 instance with the configured instance tag:
switch(config)# **router ospfv3** *instance-tag*
- Step 5** Enable IPsec ESP Encryption:
switch(config-router)#**area** *area-num* **encryption ipsec spi** *spi_val* **esp** *encrypt_algorithm* [**0** | **3** | **7**] *key*
authentication *auth_algorithm* [**0** | **3** | **7**] *key*
- You can specify the security policy index through *spi_id* and define the encryption algorithm through *encrypt_algorithm* which can be 3des, aes 128 or null. Numbers 0, 3, and 7 specify the format of the *key*. You can define the authentication algorithm through *auth_algorithm* which can be sha1 or md5.
- Note** MD5 is not supported in FIPS mode.
- Step 6** (Optional) Display OSPFv3 information:
switch(config)# **show running-config ospfv3**
-

Configuring OSPFv3 Encryption at Interface Level

You can configure OSPFv3 ESP to encrypt and authenticate OSPFv3 packets at the interface level using the following commands.

Before you begin

You must enable OSPFv3.

Enable authentication package.

Procedure

-
- Step 1** Enter the global configuration mode:
- ```
switch# configure terminal
```
- Step 2** Enter the configuration of OSPFv3 mode:
- ```
switch# configure ospfv3
```
- Step 3** Enables the authentication mode:
- ```
switch(config)# feature imp
```
- Step 4** Enters the interface configuration mode:
- ```
switch(config)# interface ethernet interface
```
- Step 5** Specify the OSPFv3 instance and area for the interface:
- ```
switch(config-if)# ipv6 router ospfv3 instance-tag area area-id
```
- Step 6** Enable IPsec ESP Encryption:
- ```
switch(config-if)# ospfv3 encryption ipsec spi spi_id esp encrypt_algorithm [ 0 | 3 | 7 ] key authentication auth_algorithm [ 0 | 3 | 7 ] key
```
- You can specify the security policy index through *spi_id* and define the encryption algorithm through *encrypt_algorithm* which can be 3des, aes 128 or null. Numbers 0,3 and 7 specify the format of the *key*. You can define the authentication algorithm through *auth_algorithm* which can be sha1 or md5.
- Note** MD5 is not supported in FIPS mode.
- Step 7** (Optional) Display the running configuration on the interface:
- ```
switch(config-if)#show run interface interface
```
- 

**Configuration Example**

The following example shows how to enable security for Ethernet interface 3/2.

```
switch# configure terminal
switch(config)# feature ospfv3
switch(config)# feature imp
switch(config)# interface ethernet 3/2
switch(config-if)# ipv6 router ospfv3 1 area 0.0.0.0
switch(config-if)# ospfv3 encryption ipsec spi 444
 esp Specify encryption parameters
switch(config-if)# ospfv3 encryption ipsec spi 444 esp
 3des Use the triple DES algorithm
 aes Use the AES algorithm
```

```

 null Use NULL authentication
switch(config-if)# ospfv3 encryption ipsec spi 444 esp aes
 128 Use the 128-bit AES algorithm
switch(config-if)# ospfv3 encryption ipsec spi 444 esp aes 128
 0 Specifies an UNENCRYPTED encryption key will follow
 3 Specifies an 3DES ENCRYPTED encryption key will follow
 7 Specifies a Cisco type 7 ENCRYPTED encryption key will follow
 WORD The UNENCRYPTED (cleartext) encryption key
switch(config-if)# ospfv3 encryption ipsec spi 444 esp aes 128
12345678123456781234567812345678 authentication null
switch(config-if)# sh ospfv3 interface
Ethernet3/2 is up, line protocol is up
 IPv6 address 1:1:1:1::2/64
 Process ID 1 VRF default, Instance ID 0, area 0.0.0.0
 Enabled by interface configuration
 State DOWN, Network type BROADCAST, cost 40
 ESP Encryption AES, Authentication NULL, SPI 444, ConnId 444
switch(config-if)#

```

## Configuring OSPFv3 Encryption for Virtual Links

You can configure OSPFv3 ESP to encrypt and authenticate OSPFv3 packets for virtual links using the following commands.

### Before you begin

You must enable OSPFv3.

Enable authentication package.

### Procedure

- 
- Step 1** Enter the global configuration mode:
- ```
switch# configure terminal
```
- Step 2** Enter the configuration of OSPFv3 mode:
- ```
switch# configure ospfv3
```
- Step 3** Enable the authentication package:
- ```
switch(config)# feature imp
```
- Step 4** Create a new OSPFv3 instance with the configured instance tag:
- ```
switch(config)#router ospfv3 instance-tag
```
- Step 5** Enable IPsec ESP Encryption:
- ```
switch(config-router)# encryption ipsec spi spi_id esp encrypt_algorithm [ 0 | 3 | 7 ] key authentication
auth_algorithm [ 0 | 3 | 7 ] key
```

You can specify the security policy index through *spi_id* and define the encryption algorithm through *encrypt_algorithm* which can be 3des, aes 128 or null. Numbers 0,3 and 7 specify the format of the *key*. You can define the authentication algorithm through *auth_algorithm* which can be sha1 or md5.

Note MD5 is not supported in FIPS mode.

Step 6 (Optional) Display OSPFv3 information:
 switch(config)# **show running-config ospfv3**

Configuration Example

The following example shows how to encrypt Virtual links.

```
switch(config)# feature ospfv3
switch(config)# feature imp
switch(config-if)# router ospfv3 1
switch(config-router)# area 0.0.0.1 virtual-link 3.3.3.3
switch(config-router-vlink)# encryption ipsec spi ?
<256-4294967295> SPI Value
switch(config-router-vlink)# encryption ipsec spi 256 esp ?
3des Use the triple DES algorithm
aes Use the AES algorithm
null Use NULL authentication
switch(config-router-vlink)# encryption ipsec spi 256 esp aes 128
123456789A123456789B123456789C12 authentication ?
null Use NULL authentication
sha1 Use the SHA1 algorithm
switch(config-router-vlink)# encryption ipsec spi 256 esp aes 128
123456789A123456789B123456789C12 authentication null
```

Configuration Examples for OSPFv3

This example shows how to configure OSPFv3:

```
This example shows how to configure OSPFv3:
feature ospfv3
router ospfv3 201
  router-id 290.0.2.1

interface ethernet 1/2
  ipv6 address 2001:0DB8::1/48
  ipv6 ospfv3 201 area 0.0.0.10
```

Related Documents for OSPFv3

Related Topic	Document Title
OSPFv3 CLI commands	<i>Cisco Nexus 7000 Series NX-OS Unicast Routing Command Reference</i>
VDCs	<i>Cisco Nexus 7000 Series NX-OS Virtual Device Context Command Reference</i>

Feature History for OSPFv3

This table includes only the updates for those releases that have resulted in additions or changes to the feature.

Table 18: Feature History for OSPFv3

Feature Name	Release	Feature Information
OSPFv3 ESP Encryption	8.4(4)	Added ESP encryption for OSPFv3 packets.
OSPF—Distribute List to Filter Paths	6.2(6a)	Added support for filtering next-hop paths for an OSPF route to prevent the path from being added to the RIB.
Administrative distance of routes	6.2(2)	Added the filter keyword to the table-map command to specify that only routes permitted by the route map are downloaded to the RIB.
Route summarization	6.2(2)	Added the ability to prevent discard routes from being created.
OSPFv3	6.2(2)	<ul style="list-style-type: none"> • Bidirectional Forwarding Detection (BFD) was enhanced to add a client for OSPFv3 • Added the ability to advertise locally generated router LSAs with the maximum metric value possible. • Added the optional name-lookup parameter for OSPFv3 instances.
MIBs	6.2(2)	Added OSPFv3 SNMP/trap support.
OSPFv3	6.1(1)	Added support for configuring the administrative distance of routes for OSPFv3.
Passive interface	5.2(1)	Added support for setting the passive interface mode on all interfaces in the router or VRF.
OSPFv3	4.0(1)	This feature was introduced.



CHAPTER 9

Configuring EIGRP

This chapter contains the following sections:

- [Finding Feature Information, on page 193](#)
- [Information About EIGRP, on page 193](#)
- [Prerequisites for EIGRP, on page 201](#)
- [Guidelines and Limitations for EIGRP, on page 201](#)
- [Default Settings for EIGRP Parameters, on page 202](#)
- [Configuring Basic EIGRP, on page 203](#)
- [Configuring Advanced EIGRP, on page 207](#)
- [Configuring Virtualization for EIGRP, on page 222](#)
- [Verifying the EIGRP Configuration, on page 224](#)
- [Displaying EIGRP Statistics, on page 224](#)
- [Configuration Example for EIGRP, on page 224](#)
- [Related Documents for EIGRP, on page 225](#)
- [MIBs, on page 226](#)
- [Feature History for EIGRP, on page 226](#)

Finding Feature Information

Your software release might not support all the features documented in this module. For the latest caveats and feature information, see the Bug Search Tool at <https://tools.cisco.com/bugsearch/> and the release notes for your software release. To find information about the features documented in this module, and to see a list of the releases in which each feature is supported, see the "New and Changed Information" chapter or the Feature History table in this chapter.

Information About EIGRP

EIGRP combines the benefits of distance vector protocols with the features of link-state protocols. EIGRP sends out periodic Hello messages for neighbor discovery. Once EIGRP learns a new neighbor, it sends a one-time update of all the local EIGRP routes and route metrics. The receiving EIGRP router calculates the route distance based on the received metrics and the locally assigned cost of the link to that neighbor. After this initial full route table update, EIGRP sends incremental updates to only those neighbors affected by the route change. This process speeds convergence and minimizes the bandwidth used by EIGRP.

EIGRP Components

EIGRP has the following basic components:

- Reliable Transport Protocol
- Neighbor Discovery and Recovery
- Diffusing Update Algorithm

Reliable Transport Protocol

The Reliable Transport Protocol guarantees ordered delivery of EIGRP packets to all neighbors. The Reliable Transport Protocol supports an intermixed transmission of multicast and unicast packets. The reliable transport can send multicast packets quickly when unacknowledged packets are pending. This provision helps to ensure that the convergence time remains low for various speed links.

The Reliable Transport Protocol includes the following message types:

- Hello—Used for neighbor discovery and recovery. By default, EIGRP sends a periodic multicast Hello message on the local network at the configured hello interval. By default, the hello interval is 5 seconds.
- Acknowledgement—Verify reliable reception of Updates, Queries, and Replies.
- Updates—Send to affected neighbors when routing information changes. Updates include the route destination, address mask, and route metrics such as delay and bandwidth. The update information is stored in the EIGRP topology table.
- Queries and Replies—Sent as part of the Diffusing Update Algorithm used by EIGRP.

Neighbor Discovery and Recovery

EIGRP uses the Hello messages from the Reliable Transport Protocol to discover neighboring EIGRP routers on directly attached networks. EIGRP adds neighbors to the neighbor table. The information in the neighbor table includes the neighbor address, the interface it was learned on, and the hold time, which indicates how long EIGRP should wait before declaring a neighbor unreachable. By default, the hold time is three times the hello interval or 15 seconds.

EIGRP sends a series of Update messages to new neighbors to share the local EIGRP routing information. This route information is stored in the EIGRP topology table. After this initial transmission of the full EIGRP route information, EIGRP sends Update messages only when a routing change occurs. These Update messages contain only the new or changed information and are sent only to the neighbors affected by the change.

EIGRP also uses the Hello messages as a keepalive to its neighbors. As long as Hello messages are received, Cisco NX-OS can determine that a neighbor is alive and functioning.

Diffusing Update Algorithm

The Diffusing Update Algorithm (DUAL) calculates the routing information based on the destination networks in the topology table. The topology table includes the following information:

- IPv4 or IPv6 address/mask—The network address and network mask for this destination.
- Successors—The IP address and local interface connection for all feasible successors or neighbors that advertise a shorter distance to the destination than the current feasible distance.

- Feasibility distance (FD)—The lowest calculated distance to the destination.

DUAL uses the distance metric to select efficient, loop-free paths. DUAL selects routes to insert into the unicast Routing Information Base (RIB) based on feasible successors. When a topology change occurs, DUAL looks for feasible successors in the topology table. If there are feasible successors, DUAL selects the feasible successor with the lowest feasible distance and inserts that into the unicast RIB, avoiding unnecessary recomputation.

When there are no feasible successors but there are neighbors advertising the destination, DUAL transitions from the passive state to the active state and triggers a recomputation to determine a new successor or next-hop router to the destination. The amount of time required to recompute the route affects the convergence time. EIGRP sends Query messages to all neighbors, searching for feasible successors. Neighbors that have a feasible successor send a Reply message with that information. Neighbors that do not have feasible successors trigger a DUAL recomputation.

EIGRP Route Updates

When a topology change occurs, EIGRP sends an Update message with only the changed routing information to affected neighbors. This Update message includes the distance information to the new or updated network destination.

The distance information in EIGRP is represented as a composite of available route metrics, including bandwidth, delay, load utilization, and link reliability. Each metric has an associated weight that determines if the metric is included in the distance calculation. You can configure these metric weights. You can fine-tune link characteristics to achieve optimal paths, but we recommend that you use the default settings for most configurable metrics.

Internal Route Metrics

Internal routes are routes that occur between neighbors within the same EIGRP autonomous system. These routes have the following metrics:

- Next hop—The IP address of the next-hop router.
- Delay—The sum of the delays configured on the interfaces that make up the route to the destination network. The delay is configured in tens of microseconds.
- Bandwidth—The calculation from the lowest configured bandwidth on an interface that is part of the route to the destination.



Note We recommend that you use the default bandwidth value. This bandwidth parameter is also used by EIGRP.

- MTU—The smallest maximum transmission unit value along the route to the destination.
- Hop count—The number of hops or routers that the route passes through to the destination. This metric is not directly used in the DUAL computation.
- Reliability—An indication of the reliability of the links to the destination.
- Load—An indication of how much traffic is on the links to the destination.

By default, EIGRP uses the bandwidth and delay metrics to calculate the distance to the destination. You can modify the metric weights to include the other metrics in the calculation.

Wide Metrics

EIGRP supports wide (64-bit) metrics to improve route selection on higher-speed interfaces or bundled interfaces. Routers supporting wide metrics can interoperate with routers that do not support wide metrics as follows:

- A router that supports wide metrics—Adds local wide metrics values to the received values and sends the information on.
- A router that does not support wide metrics—Sends any received metrics on without changing the values.

EIGRP uses the following equation to calculate path cost with wide metrics:

$$\text{metric} = [k1 \times \text{bandwidth} + (k2 \times \text{bandwidth}) / (256 - \text{load}) + k3 \times \text{delay} + k6 \times \text{extended attributes}] \times [k5 / (\text{reliability} + k4)]$$

Because the unicast RIB cannot support 64-bit metric values, EIGRP wide metrics use the following equation with a RIB scaling factor to convert the 64-bit metric value to a 32-bit value:

$$\text{RIB Metric} = (\text{Wide Metric} / \text{RIB scale value})$$

where the RIB scale value is a configurable parameter.

EIGRP wide metrics introduce the following two new metric values represented as k6 in the EIGRP metrics configuration:

- Jitter—(Measured in microseconds) accumulated across all links in the route path. Routes lower jitter values are preferred for EIGRP path selection.
- Energy—(Measured in watts per kilobit) accumulated across all links in the route path. Routes lower energy values are preferred for EIGRP path selection.

EIGRP prefers a path with no jitter or energy metric values or lower jitter or metric values over a path with higher values.



Note EIGRP wide metrics are sent with a TLV version of 2.

External Route Metrics

External routes are routes that occur between neighbors in different EIGRP autonomous systems. These routes have the following metrics:

- Next hop—The IP address of the next-hop router.
- Router ID—The router ID of the router that redistributed this route into EIGRP.
- AS number—The autonomous system number of the destination.
- Protocol ID—A code that represents the routing protocol that learned the destination route.
- Tag—An arbitrary tag that can be used for route maps.
- Metric—The route metric for this route from the external routing protocol.

EIGRP and the Unicast RIB

EIGRP adds all learned routes to the EIGRP topology table and the unicast RIB. When a topology change occurs, EIGRP uses these routes to search for a feasible successor. EIGRP also listens for notifications from the unicast RIB for changes in any routes redistributed to EIGRP from another routing protocol.

Advanced EIGRP

You can use the advanced features of EIGRP to optimize your EIGRP configuration.

Address Families

EIGRP supports both IPv4 and IPv6 address families. For backward compatibility, you can configure EIGRPv4 in route configuration mode or in IPv4 address family mode. You must configure EIGRP for IPv6 in address family mode.

Address family configuration mode includes the following EIGRP features:

- Authentication
- AS number
- Default route
- Metrics
- Distance
- Graceful restart
- Logging
- Load balancing
- Redistribution
- Router ID
- Stub router
- Timers

You cannot configure the same feature in more than one configuration mode. For example, if you configure the default metric in router configuration mode, you cannot configure the default metric in address family mode.

Authentication

You can configure authentication on EIGRP messages to prevent unauthorized or invalid routing updates in your network. EIGRP authentication supports MD5 authentication digest.

You can configure the EIGRP authentication per virtual routing and forwarding (VRF) instance or interface using key-chain management for the authentication keys. Key-chain management allows you to control changes to the authentication keys used by MD5 authentication digest. See the *Cisco Nexus 7000 Series NX-OS Security Configuration Guide*, for more details about creating key chains.

For MD5 authentication, you configure a password that is shared at the local router and all remote EIGRP neighbors. When an EIGRP message is created, Cisco NX-OS creates an MD5 one-way message digest based on the message itself and the encrypted password and sends this digest along with the EIGRP message. The receiving EIGRP neighbor validates the digest using the same encrypted password. If the message has not changed, the calculation is identical and the EIGRP message is considered valid.

MD5 authentication also includes a sequence number with each EIGRP message that is used to ensure that no message is replayed in the network.

Stub Routers

You can use the EIGRP stub routing feature to improve network stability, reduce resource usage, and simplify stub router configuration. Stub routers connect to the EIGRP network through a remote router.

When using EIGRP stub routing, you need to configure the distribution and remote routers to use EIGRP and configure only the remote router as a stub. EIGRP stub routing does not automatically enable summarization on the distribution router. In most cases, you need to configure summarization on the distribution routers.

Without EIGRP stub routing, even after the routes that are sent from the distribution router to the remote router have been filtered or summarized, a problem might occur. For example, if a route is lost somewhere in the corporate network, EIGRP could send a query to the distribution router. The distribution router could then send a query to the remote router even if routes are summarized. If a problem communicating over the WAN link between the distribution router and the remote router occurs, EIGRP could get stuck in an active condition and cause instability elsewhere in the network. EIGRP stub routing allows you to prevent queries to the remote router.

Route Summarization

You can configure a summary aggregate address for a specified interface. Route summarization simplifies route tables by replacing a number of more-specific addresses with an address that represents all the specific addresses. For example, you can replace 10.1.1.0/24, 10.1.2.0/24, and 10.1.3.0/24 with one summary address, 10.1.0.0/16.

If more specific routes are in the routing table, EIGRP advertises the summary address from the interface with a metric equal to the minimum metric of the more specific routes.

In case of process restart or system switchover, the summary address can cause traffic loss. The traffic loss will be seen on the PEER where traffic is routed using the summary address.



Note EIGRP does not support automatic route summarization.

Route Redistribution

You can use EIGRP to redistribute static routes, routes learned by other EIGRP autonomous systems, or routes from other protocols. You must configure a route map with the redistribution to control which routes are passed into EIGRP. A route map allows you to filter routes based on attributes such as the destination, origination protocol, route type, route tag, and so on.

You also configure the default metric that is used for all imported routes into EIGRP.

You use distribute lists to filter routes from routing updates. These filtered routes are applied to each interface with the **ip distribute-list eigrp** command.

Load Balancing

You can use load balancing to allow a router to distribute traffic over all the router network ports that are the same distance from the destination address. Load balancing increases the usage of network segments, which increases effective network bandwidth.

Cisco NX-OS supports the Equal Cost Multiple Paths (ECMP) feature with up to 16 equal-cost paths in the EIGRP route table and the unicast RIB. You can configure EIGRP to load balance traffic across some or all of those paths.



Note EIGRP in Cisco NX-OS does not support unequal cost load balancing.

Split Horizon

You can use split horizon to ensure that EIGRP never advertises a route out of the interface where it was learned.

Split horizon is a method that controls the sending of EIGRP update and query packets. When you enable split horizon on an interface, Cisco NX-OS does not send update and query packets for destinations that were learned from this interface. Controlling update and query packets in this manner reduces the possibility of routing loops.

Split horizon with poison reverse configures EIGRP to advertise a learned route as unreachable back through that the interface that EIGRP learned the route from.

EIGRP uses split horizon or split horizon with poison reverse in the following scenarios:

- Exchanging topology tables for the first time between two routers in startup mode.
- Advertising a topology table change.
- Sending a Query message.

By default, the split horizon feature is enabled on all interfaces.

BFD

This feature supports bidirectional forwarding detection (BFD). BFD is a detection protocol designed to provide fast forwarding-path failure detection times. BFD provides subsecond failure detection between two adjacent devices and can be less CPU-intensive than protocol hello messages because some of the BFD load can be distributed onto the data plane on supported modules. See the *Cisco Nexus 7000 Series NX-OS Interfaces Configuration Guide*, for more information.

Virtualization Support for EIGRP

Cisco NX-OS supports multiple instances of EIGRP that runs on the same system. EIGRP supports Virtual Routing and Forwarding instances (VRFs). VRFs exist within virtual device contexts (VDCs). By default, Cisco NX-OS places you in the default VDC and default VRF unless you specifically configure another VDC and VRF. See the *Cisco Nexus 7000 Series NX-OS Virtual Device Context Configuration Guide*, for more information.

Graceful Restart and High Availability

Cisco NX-OS supports nonstop forwarding and graceful restart for EIGRP.

You can use nonstop forwarding for EIGRP to forward data packets along known routes in the FIB while the EIGRP routing protocol information is being restored following a failover. With nonstop forwarding (NSF), peer networking devices do not experience routing flaps. During failover, data traffic is forwarded through intelligent modules while the standby supervisor becomes active.

If a Cisco NX-OS system experiences a cold reboot, the device does not forward traffic to the system and removes the system from the network topology. In this scenario, EIGRP experiences a stateless restart, and all neighbors are removed. Cisco NX-OS applies the startup configuration, and EIGRP rediscovers the neighbors and shares the full EIGRP routing information again.

A dual supervisor platform that runs Cisco NX-OS can experience a stateful supervisor switchover. Before the switchover occurs, EIGRP uses a graceful restart to announce that EIGRP will be unavailable for some time. During a switchover, EIGRP uses nonstop forwarding to continue forwarding traffic based on the information in the FIB, and the system is not taken out of the network topology.

The graceful restart-capable router uses Hello messages to notify its neighbors that a graceful restart operation has started. When a graceful restart-aware router receives a notification from a graceful restart-capable neighbor that a graceful restart operation is in progress, both routers immediately exchange their topology tables. The graceful restart-aware router performs the following actions to assist the restarting router as follows:

- The router expires the EIGRP Hello hold timer to reduce the time interval set for Hello messages. This process allows the graceful restart-aware router to reply to the restarting router more quickly and reduces the amount of time required for the restarting router to rediscover neighbors and rebuild the topology table.
- The router starts the route-hold timer. This timer sets the period of time that the graceful restart-aware router will hold known routes for the restarting neighbor. The default time period is 240 seconds.
- The router notes in the peer list that the neighbor is restarting, maintains adjacency, and holds known routes for the restarting neighbor until the neighbor signals that it is ready for the graceful restart-aware router to send its topology table or the route-hold timer expires. If the route-hold timer expires on the graceful restart-aware router, the graceful restart-aware router discards held routes and treats the restarting router as a new router that joins the network and reestablishes adjacency.

After the switchover, Cisco NX-OS applies the running configuration, and EIGRP informs the neighbors that it is operational again.



Note You must enable graceful restart to support in-service software upgrades (ISSU) for EIGRP. If you disable graceful restart, Cisco NX-OS issues a warning that an ISSU cannot be supported with this configuration.

Multiple EIGRP Instances

Cisco NX-OS supports multiple instances of the EIGRP protocol that run on the same system. Every instance uses the same system router ID. You can optionally configure a unique router ID for each instance. For the number of supported EIGRP instances, see the *Cisco Nexus 7000 Series NX-OS Verified Scalability Guide*.

Prerequisites for EIGRP

You must enable EIGRP.

If you configure VDCs, you must install the Advanced Services license and enter the desired VDC (see the *Cisco Nexus 7000 Series NX-OS Virtual Device Context Configuration Guide*).

Guidelines and Limitations for EIGRP

- A system configured with EIGRP 64bit metric version and having 32bit metric version neighbors precision in delay conversion from 64bit metric to 32bit metric and from 32bit metric to 64bit metric is improved. It implies that the 32bit metric of prefixes in the 64bit EIGRP system and 64bit metric of prefixed in 32bit EIGRP system changes from previous releases. This is not applicable for the 32bit metric EIGRP system or if all neighbors are the 64bit metric version.
- When you configure a table map, administrative distance of the routes and the metric, the configuration commands make the EIGRP neighbours to flap. This is an expected behavior.
- If the filtered list is modified when redistributing routes into EIGRP and filtering prefixes with a route map or prefix list, all prefixes permitted by the filter, even those not touched, are refreshed in the EIGRP topology table. This refresh is signaled to all EIGRP routers in the query domain for this set of prefixes.
- A metric configuration (either through the default-metric configuration option or through a route map) is required for redistribution from any other protocol, connected routes, or static routes.
- For graceful restart, an NSF-aware router must be up and completely converged with the network before it can assist an NSF-capable router in a graceful restart operation.
- For graceful restart, neighboring devices participating in the graceful restart must be NSF-aware or NSF-capable.
- Cisco NX-OS EIGRP is compatible with EIGRP in the Cisco IOS software.
- Do not change the metric weights without a good reason. If you change the metric weights, you must apply the change to all EIGRP routers in the same autonomous system.
- A mix of standard metrics and wide metrics in an EIGRP network with interface speeds of 1 Gigabit or greater may result in suboptimal routing.
- Consider using stubs for larger networks.
- Avoid redistribution between different EIGRP autonomous systems because the EIGRP vector metric will not be preserved.
- The **no {ip | ipv6} next-hop-self** command does not guarantee reachability of the next hop.
- The **{ip | ipv6} passive-interface eigrp** command suppresses neighbors from forming.
- Cisco NX-OS does not support IGRP or connecting IGRP and EIGRP clouds.
- Autosummarization is disabled by default and cannot be enabled.
- Cisco NX-OS supports only IP.
- EIGRPv6 adjacency cannot be formed over an interface that only has IPv6 link local address.

Global IPv6 address is required on the interface for EIGRPv6 neighbour adjacency to be formed over such interface.

- High availability is not supported with EIGRP aggressive timers.
- If you are familiar with the Cisco IOS CLI, be aware that the Cisco NX-OS commands for this feature might differ from the Cisco IOS commands that you would use.

Default Settings for EIGRP Parameters

Table 19: Default Settings for EIGRP Parameters

Parameters	Default
Administrative distance	Internal routes—90 External routes—170
Bandwidth percent	50 percent
Default metric for redistributed routes	Bandwidth—100000 Kb/s Delay—100 (10 microsecond units) Reliability—255 Loading—1 MTU—1500
EIGRP feature	Disabled
Hello interval	5 seconds
Hold time	15 seconds
Equal-cost paths	8
Metric weights	1 0 1 0 0 0
Next-hop address advertised	IP address of local interface
NSF convergence time	120
NSF route-hold time	240
NSF signal time	20
Redistribution	Disabled
Split horizon	Enabled

Configuring Basic EIGRP

Enabling or Disabling the EIGRP Feature

You must enable the EIGRP feature before you can configure EIGRP.

Before you begin

Ensure that you are in the correct VDC (or use the `switchto vdc` command).

Procedure

	Command or Action	Purpose
Step 1	<code>switch# configure terminal</code>	Enters global configuration mode.
Step 2	<code>switch(config)# [no] feature eigrp</code>	Enables the EIGRP feature. The <code>no</code> option disables the EIGRP feature and removes all associated configuration.
Step 3	(Optional) <code>switch(config)# show feature</code>	Displays information about enabled features.
Step 4	(Optional) <code>switch(config)# copy running-config startup-config</code>	Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

Example

The following example enables EIGRP:

```
switch# configure terminal
switch(config)# feature eigrp
switch(config)# copy running-config startup-config
```

Creating an EIGRP Instance

You can create an EIGRP instance and associate an interface with that instance. You assign a unique autonomous system number for this EIGRP process. Routes are not advertised or accepted from other autonomous systems unless you enable route redistribution.

Before you begin

- Ensure that you have enabled the EIGRP feature.
- EIGRP must be able to obtain a router ID (for example, a configured loopback address) or you must configure the router ID option.
- If you configure an instance tag that does not qualify as an AS number, you must configure the AS number explicitly or this EIGRP instance remains in the shutdown state. For IPv6, this number must be configured under address family.

- Confirm that you are in the correct VDC. To change the VDC, use the **switchto vdc** command.

Procedure

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# [no] router eigrp <i>instance-tag</i>	<p>Creates a new EIGRP process with the configured instance tag. The instance tag can be any case-sensitive, alphanumeric string up to 20 characters.</p> <p>If you configure an instance tag that does not qualify as an AS number, you must use the autonomous-system command to configure the AS number explicitly or this EIGRP instance will remain in the shutdown state.</p> <p>The no option deletes the EIGRP process and all associated configuration. You should also remove any EIGRP commands configured in interface mode if you remove the EIGRP process.</p>
Step 3	(Optional) switch(config-router)# autonomous-system <i>as-number</i>	<p>Configures a unique AS number for this EIGRP instance.</p> <p>The range is from 1 to 65535.</p>
Step 4	(Optional) switch(config-router)# log-adjacency-changes	Generates a system message whenever an adjacency changes state. This command is enabled by default.
Step 5	(Optional) switch(config-router)# log-neighbor-warnings [<i>seconds</i>]	<p>Generates a system message whenever a neighbor warning occurs.</p> <p>You can configure the time between warning messages, from 1 to 65535, in seconds. The default is 10 seconds. This command is enabled by default.</p>
Step 6	switch(config-router)# interface <i>interface-type slot/port</i>	Enters interface configuration mode. Use ? to determine the slot and port ranges.
Step 7	switch(config-if)# {ip ipv6} router eigrp <i>instance-tag</i>	Associates this interface with the configured EIGRP process. The instance tag can be any case-sensitive, alphanumeric string up to 20 characters.
Step 8	(Optional) switch(config-if)# show {ip ipv6} eigrp interfaces	Displays information about EIGRP interfaces.

	Command or Action	Purpose
Step 9	(Optional) switch(config-if)# copy running-config startup-config	Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

Example

This example shows how to create an EIGRP process and configure an interface for EIGRP:

```
switch# configure terminal
switch(config)# router eigrp Test1
switch(config)# interface ethernet 1/2

switch(config-if)# ip router eigrp Test1
switch(config-if)# no shutdown
switch(config-if)# copy running-config startup-config
```

Restarting an EIGRP Instance

You can restart an EIGRP instance. This action clears all neighbors for the instance.

Procedure

	Command or Action	Purpose
Step 1	(Optional) switch(config)# flush-routes	Flushes all EIGRP routes in the unicast RIB when this EIGRP instance restarts.
Step 2	Required: switch(config)# restart eigrp instance-tag	Restarts the EIGRP instance and removes all neighbors. The instance tag can be any case-sensitive, alphanumeric string up to 20 characters.
Step 3	(Optional) switch(config)# copy running-config startup-config	Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

Example

The following example shows how to restart an EIGRP instance:

```
switch(config)# flush-routes
switch(config)# restart eigrp Test1
switch(config)# copy running-config startup-config
```

Shutting Down an EIGRP Instance

You can gracefully shut down an EIGRP instance. This action removes all routes and adjacencies but preserves the EIGRP configuration.

Procedure

	Command or Action	Purpose
Step 1	switch(config-router)# shutdown	Disables this instance of EIGRP. The EIGRP router configuration remains.
Step 2	(Optional) switch(config-router)# copy running-config startup-config	Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

Example

The following example shows how to shut down an EIGRP instance:

```
switch(config-router)# shutdown
switch(config-router)# copy running-config startup-config
```

Configuring a Passive Interface for EIGRP

You can configure a passive interface for EIGRP. A passive interface does not participate in EIGRP adjacency, but the network address for the interface remains in the EIGRP topology table.

Procedure

	Command or Action	Purpose
Step 1	switch(config-if)# [default no]{ ip ipv6 } passive-interface eigrp <i>instance-tag</i>	Suppresses EIGRP hellos, which prevents neighbors from forming and sending routing updates on an EIGRP interface. The instance tag can be any case-sensitive, alphanumeric string up to 20 characters. Note <ul style="list-style-type: none"> To configure all EIGRP interfaces as passive by default, use the passive-interface default command. To configure disabled passive interface on EIGRP, you must use no command with this command. To unconfigure passive interface on EIGRP, you must use default with this command.
Step 2	(Optional) switch(config-if)# copy running-config startup-config	Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

Example

The following example shows how to configure a passive interface for EIGRP:

```
switch(config-if) # ip passive-interface eigrp tag10
switch(config-if) # copy running-config startup-config
```

The following example shows how to configure disabled passive-interface for EIGRP:

```
switch(config-if) # no ip passive-interface eigrp tag10
switch(config-if) # copy running-config startup-config
```

The following example shows how to unconfigure passive-interface for EIGRP:

```
switch(config-if) # default ip passive-interface eigrp tag10
switch(config-if) # copy running-config startup-config
```

Shutting Down EIGRP on an Interface

You can gracefully shut down EIGRP on an interface. This action removes all adjacencies and stops EIGRP traffic on this interface but preserves the EIGRP configuration.

Procedure

	Command or Action	Purpose
Step 1	switch(config-if)# {ip ipv6} eigrp instance-tag shutdown	Disables EIGRP on this interface. The EIGRP interface configuration remains. The instance tag can be any case-sensitive, alphanumeric string up to 20 characters.
Step 2	(Optional) switch(config-if)# copy running-config startup-config	Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

Example

The following example shows how to shut down EIGRP on an interface:

```
switch(config-if) # ip eigrp Test1 shutdown
switch(config-if) # copy running-config startup-config
```

Configuring Advanced EIGRP

Configuring Authentication in EIGRP

You can configure authentication between neighbors for EIGRP. You can configure EIGRP authentication for the EIGRP process or for individual interfaces. The interface EIGRP authentication configuration overrides the EIGRP process-level authentication configuration.

Before you begin

- Ensure that you have enabled the EIGRP feature.
- Ensure that all neighbors for an EIGRP process share the same authentication configuration, including the shared authentication key.
- Create the key chain for this authentication configuration. For more information, see the *Cisco Nexus 7000 Series NX-OS Security Configuration Guide*.
- Confirm that you are in the correct VDC. To change the VDC, use the **switchto vdc** command.

Procedure

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# router eigrp <i>instance-tag</i>	Creates a new EIGRP process with the configured instance tag. The instance tag can be any case-sensitive, alphanumeric string up to 20 characters. If you configure an instance tag that does not qualify as an AS number, you must use the autonomous-system command to configure the AS number explicitly or this EIGRP instance will remain in the shutdown state.
Step 3	switch(config-router)# address-family { ipv4 ipv6 } unicast	Enters the address-family configuration mode. This command is optional for IPv4.
Step 4	switch(config-router-af)# authentication key-chain <i>key-chain</i>	Associates a key chain with this EIGRP process for this VRF. The key chain can be any case-sensitive, alphanumeric string up to 20 characters.
Step 5	switch(config-router-af)# authentication mode md5	Configures MD5 message digest authentication mode for this VRF.
Step 6	switch(config-router-af)# interface <i>interface-type slot/port</i>	Enters interface configuration mode. Use ? to find the supported interfaces.
Step 7	switch(config-if)# { ip ipv6 } router eigrp <i>instance-tag</i>	Associates this interface with the configured EIGRP process. The instance tag can be any case-sensitive, alphanumeric string up to 20 characters.
Step 8	switch(config-if)# { ip ipv6 } authentication key-chain eigrp <i>instance-tag key-chain</i>	Associates a key chain with this EIGRP process for this interface. This configuration overrides the authentication configuration set in the router VRF mode. The instance tag can be any case-sensitive, alphanumeric string up to 20 characters.

	Command or Action	Purpose
Step 9	switch(config-if)# {ip ipv6} authentication mode eigrp instance-tag md5	Configures the MD5 message digest authentication mode for this interface. This configuration overrides the authentication configuration set in the router VRF mode. The instance tag can be any case-sensitive, alphanumeric string up to 20 characters.
Step 10	(Optional) switch(config-if)# copy running-config startup-config	Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

Example

The following example shows how to configure MD5 message digest authentication for EIGRP over Ethernet interface 1/2:

```
switch# configure terminal
switch(config)# router eigrp Test1
switch(config-router)# exit
switch(config)# interface ethernet 1/2

switch(config-if)# ip router eigrp Test1
switch(config-if)# ip authentication key-chain eigrp Test1 routeKeys
switch(config-if)# ip authentication mode eigrp Test1 md5
switch(config-if)# copy running-config startup-config
```

Configuring EIGRP Stub Routing

To configure a router for EIGRP stub routing, use these commands in address-family configuration mode:

Procedure

	Command or Action	Purpose
Step 1	switch(config-router-af)# stub [direct receive-only redistributed [direct] leak-map map-name]	Configures a remote router as an EIGRP stub router. The map name can be any case-sensitive, alphanumeric string up to 20 characters.
Step 2	(Optional) switch(config-router-af)# show ip eigrp neighbor detail	Verifies that the router has been configured as a stub router.
Step 3	(Optional) switch(config-router-af)# copy running-config startup-config	Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

Example

The following example shows how to configure a stub router to advertise directly connected and redistributed routes. The last line of the output for the **show ip eigrp neighbor detail** command shows the stub status of the remote or spoke router.

```
switch# configure terminal
switch(config)# router eigrp Test1
switch(config-router)# address-family ipv4 unicast
switch(config-router-af)# stub direct redistributed
switch(config-router-af)# show ip eigrp neighbor detail

IP-EIGRP neighbors for process 201
H   Address                Interface    Hold Uptime    SRTT   RTO  Q  Seq Type
      (sec)                (ms)                Cnt Num
0   10.1.1.2                Se3/1       11 00:00:59    1  4500  0  7
      Version 12.1/1.2, Retrans: 2, Retries: 0
      Stub Peer Advertising ( CONNECTED SUMMARY ) Routes

switch(config-router-af)# copy running-config startup-config
```

Configuring a Summary Address for EIGRP

You can configure a summary aggregate address for a specified interface. If any more specific routes are in the routing table, EIGRP advertises the summary address out the interface with a metric equal to the minimum of all more specific routes.

Procedure

	Command or Action	Purpose
Step 1	switch(config-if)# {ip ipv6} summary-address eigrp instance-tag ip-prefix/length [distance leak-map map-name]	Configures a summary aggregate address as either an IP address and network mask or an IP prefix/length. The instance tag and map name can be any case-sensitive, alphanumeric string up to 20 characters. You can optionally configure the administrative distance for this aggregate address. The default administrative distance is 5 for aggregate addresses.
Step 2	(Optional) switch(config-if)# copy running-config startup-config	Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

Example

The following example shows how to cause EIGRP to summarize network 192.0.2.0 out Ethernet 1/2 only:

```
switch(config)# interface ethernet 1/2
```

```
switch(config-if)# ip summary-address eigrp Test1 192.0.2.0 255.255.255.0
switch(config-if)# copy running-config startup-config
```

Redistributing Routes into EIGRP

You can redistribute routes in EIGRP from other routing protocols.

Before you begin

- Ensure that you have enabled the EIGRP feature.
- You must configure the metric (either through the default-metric configuration option or through a route map) for routes redistributed from any other protocol.
- You must create a route map to control the types of routes that are redistributed into EIGRP.
- Confirm that you are in the correct VDC. To change the VDC, use the **switchto vdc** command.

Procedure

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# router eigrp <i>instance-tag</i>	Creates a new EIGRP process with the configured instance tag. The instance tag can be any case-sensitive, alphanumeric string up to 20 characters. If you configure an instance tag that does not qualify as an AS number, you must use the autonomous-system command to configure the AS number explicitly or this EIGRP instance will remain in the shutdown state.
Step 3	switch(config-router)# address-family { ipv4 ipv6 } unicast	Enters the address-family configuration mode. This command is optional for IPv4.
Step 4	switch(config-router-af)# redistribute { bgp <i>as</i> { eigrp isis ospf ospfv3 rip } <i>instance-tag</i> direct static } route-map <i>map-name</i>	Injects routes from one routing domain into EIGRP. The instance tag and map name can be any case-sensitive, alphanumeric string up to 20 characters.
Step 5	switch(config-router-af)# default-metric <i>bandwidth delay reliability loading mtu</i>	Sets the metrics assigned to routes learned through route redistribution. The default values are as follows: <ul style="list-style-type: none"> • bandwidth—100000 Kb/s • delay—100 (10 microsecond units) • reliability—255 • loading—1

	Command or Action	Purpose
		<ul style="list-style-type: none"> • MTU—1492
Step 6	(Optional) switch(config-router-af)# show {ip ipv6} eigrp route-map statistics redistribute	Displays information about EIGRP route map statistics.
Step 7	(Optional) switch(config-router-af)# copy running-config startup-config	Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

Example

The following example shows how to redistribute BGP into EIGRP for IPv4:

```
switch# configure terminal
switch(config)# router eigrp Test1
switch(config-router)# redistribute bgp 100 route-map BGPFilter
switch(config-router)# default-metric 500000 30 200 1 1500
switch(config-router)# copy running-config startup-config
```

Limiting the Number of Redistributed Routes

Route redistribution can add many routes to the EIGRP route table. You can configure a maximum limit to the number of routes accepted from external protocols. EIGRP provides the following options to configure redistributed route limits:

- Fixed limit—Logs a message when EIGRP reaches the configured maximum. EIGRP does not accept any more redistributed routes. You can optionally configure a threshold percentage of the maximum where EIGRP logs a warning when that threshold is passed.
- Warning only—Logs a warning only when EIGRP reaches the maximum. EIGRP continues to accept redistributed routes.
- Withdraw—Starts the timeout period when EIGRP reaches the maximum. After the timeout period, EIGRP requests all redistributed routes if the current number of redistributed routes is less than the maximum limit. If the current number of redistributed routes is at the maximum limit, EIGRP withdraws all redistributed routes. You must clear this condition before EIGRP accepts more redistributed routes. You can optionally configure the timeout period.

Before you begin

- Ensure that you have enabled the EIGRP feature.
- Confirm that you are in the correct VDC. To change the VDC, use the **switchto vdc** command.

Procedure

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.

	Command or Action	Purpose
Step 2	switch(config)# router eigrp <i>instance-tag</i>	Creates a new EIGRP process with the configured instance tag.
Step 3	switch(config-router)# redistribute { bgp id direct eigrp id isis id ospf id rip id static } route-map <i>map-name</i>	Redistributes the selected protocol into EIGRP through the configured route map.
Step 4	switch(config-router)# redistribute maximum-prefix <i>max</i> [<i>threshold</i>] [warning-only withdraw [<i>num-retries</i> <i>timeout</i>]]	Specifies a maximum number of prefixes that EIGRP distributes. The range is from 0 to 65536. Optionally specifies the following: <ul style="list-style-type: none"> • threshold—Percent of maximum prefixes that triggers a warning message. • warning-only—Logs a warning message when the maximum number of prefixes is exceeded. • withdraw—Withdraws all redistributed routes. Optionally tries to retrieve the redistributed routes. The <i>num-retries</i> range is from 1 to 12. The timeout is from 60 to 600 seconds. The default is 300 seconds. Use the clear ip eigrp redistribution command if all routes are withdrawn.
Step 5	(Optional) switch(config-router)# show running-config eigrp	Displays the EIGRP configuration.
Step 6	(Optional) switch(config-router)# copy running-config startup-config	Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

Example

The following example shows how to limit the number of redistributed routes into EIGRP:

```
switch# configure terminal
switch(config)# router eigrp Test1
switch(config-router)# redistribute bgp route-map FilterExternalBGP
switch(config-router)# redistribute maximum-prefix 1000 75
```

Configuring the Administrative Distance of Routes

You can set the administrative distance of routes added by EIGRP into the RIB.

Before you begin

You must enable EIGRP.

Ensure that you are in the correct VDC (or use the **switchto vdc** command).

Procedure

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# router eigrp <i>instance-tag</i>	Creates a new EIGRP instance and enters router configuration mode.
Step 3	switch(config-router)# table-map <i>route-map-name</i> [filter]	Configures a table map with route map information. You can enter up to 63 alphanumeric characters for the map name. The filter keyword filters routes rejected by the route map and does not download them to the RIB. Note When you configure a table map, administrative distance of the routes and the metric, the configuration commands make the EIGRP neighbours to flap. This is an expected behavior.
Step 4	(Optional) switch(config)# copy running-config startup-config	Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

Configuring Route-Map Filtering

You can enable EIGRP to interoperate with other protocols to leverage additional routing functionality by filtering inbound and outbound traffic based on route-map options.

Before you begin

You must enable EIGRP.

Ensure that you are in the correct VDC (or use the **switchto vdc** command).

Procedure

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# route-map <i>map-tag</i> [permit deny] [<i>sequence-number</i>]	Enters route-map configuration mode.
Step 3	switch(config-route-map)# match metric <i>metric-value</i> [+ <i>deviation-number</i>] [... <i>metric-value</i> [+ <i>deviation-number</i>]]	Specifies a match clause that filters inbound updates that match an internal or external protocol metric. The <i>metric-value</i> argument is an internal protocol metric that can be an EIGRP five-part metric. The range is from 1 to 4294967295.

	Command or Action	Purpose
		The +- <i>deviation-number</i> argument represents a standard deviation, which can be any number. When you specify a metric deviation with the + and - keywords, the router matches any metric that falls inclusively in that range.
Step 4	switch(config-route-map)# match source-protocol <i>source-protocol</i> [<i>as-number</i>]	Specifies a match clause that matches external routes from sources that match the source protocol. The <i>source-protocol</i> argument is the protocol to match. The valid options are bgp , connected , eigrp , isis , ospf , rip , and static . The <i>as-number</i> argument does not apply to the connected , rip , and static options. The range is from 1 to 65535.
Step 5	switch(config-route-map)# set tag <i>tag-value</i>	Sets a tag value on the route in the destination routing protocol when all the match criteria of a route map are met.
Step 6	switch(config-route-map)# exit	Exits route-map configuration mode.
Step 7	switch(config)# router eigrp <i>instance-tag</i>	Creates a new EIGRP instance and enters router configuration mode.
Step 8	switch(config-router)# exit	Exits router configuration mode.
Step 9	switch(config)# interface <i>interface-type</i> <i>slot/port</i>	Enters interface configuration mode. Use ? to determine the slot and port ranges.
Step 10	switch(config-if)# ip address <i>ip-address</i>	Specifies an IP address for the EIGRP routing process.
Step 11	switch(config-if)# ip router eigrp <i>as-number</i>	Configures the EIGRP routing process and enters the router configuration mode.
Step 12	switch(config-if)# ip distribute-list eigrp <i>as-number</i> route-map <i>map-tag</i> in	Filters networks received in updates.
Step 13	(Optional) switch(config)# copy running-config startup-config	Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

Configuring Load Balancing in EIGRP

You can configure the number of Equal Cost Multiple Path (ECMP) routes using the **maximum-paths** option.

Before you begin

- Ensure that you have enabled the EIGRP feature.

- Confirm that you are in the correct VDC. To change the VDC, use the **switchto vdc** command.

Procedure

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# router eigrp <i>instance-tag</i>	Creates a new EIGRP process with the configured instance tag. The instance tag can be any case-sensitive, alphanumeric string up to 20 characters. If you configure an instance tag that does not qualify as an AS number, you must use the autonomous-system command to configure the AS number explicitly or this EIGRP instance remains in the shutdown state.
Step 3	switch(config-router)# address-family { ipv4 ipv6 } unicast	Enters the address-family configuration mode. This command is optional for IPv4.
Step 4	switch(config-router-af)# maximum-paths <i>num-paths</i>	Sets the number of equal cost paths that EIGRP accepts in the route table. The range is from 1 to 16. The default is 8.
Step 5	(Optional) switch(config-router-af)# copy running-config startup-config	Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

Example

The following example shows how to configure equal cost load balancing for EIGRP over IPv4 with a maximum of six equal cost paths:

```
switch# configure terminal
switch(config)# router eigrp Test1
switch(config-router)# address-family ipv4 unicast
switch(config-router-af)# maximum-paths 6
switch(config-router-af)# copy running-config startup-config
```

Configuring Graceful Restart for EIGRP

You can configure graceful restart or nonstop forwarding for EIGRP.



Note Graceful restart is enabled by default.

Before you begin

- Ensure that you have enabled the EIGRP feature.
- An NSF-aware router must be up and completely converged with the network before it can assist an NSF-capable router in a graceful restart operation.
- Neighboring devices participating in the graceful restart must be NSF aware or NSF capable.
- Confirm that you are in the correct VDC. To change the VDC, use the **switchto vdc** command.

Procedure

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# router eigrp <i>instance-tag</i>	Creates a new EIGRP process with the configured instance tag. The instance tag can be any case-sensitive, alphanumeric string up to 20 characters. If you configure an instance tag that does not qualify as an AS number, you must use the autonomous-system command to configure the AS number explicitly or this EIGRP instance remains in the shutdown state.
Step 3	switch(config-router)# address-family { ipv4 ipv6 } unicast	Enters the address-family configuration mode. This command is optional for IPv4.
Step 4	switch(config-router-af)# graceful-restart	Enables graceful restart. This feature is enabled by default.
Step 5	switch(config-router-af)# timers nsf converge <i>seconds</i>	Sets the time limit for the convergence after a switchover. The range is from 60 to 180 seconds. The default is 120.
Step 6	switch(config-router-af)# timers nsf route-hold <i>seconds</i>	Sets the hold time for routes learned from the graceful restart-aware peer. The range is from 20 to 300 seconds. The default is 240.
Step 7	switch(config-router-af)# timers nsf signal <i>seconds</i>	Sets the time limit for signaling a graceful restart. The range is from 10 to 30 seconds. The default is 20.
Step 8	(Optional) switch(config-router-af)# copy running-config startup-config	Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

Example

This example shows how to configure graceful restart for EIGRP over IPv6 using the default timer values:

```
switch# configure terminal
switch(config)# router eigrp Test1
switch(config-router)# address-family ipv6 unicast
switch(config-router-af)# graceful-restart
switch(config-router-af)# copy running-config startup-config
```

Adjusting the Interval Between Hello Packets and the Hold Time

You can adjust the interval between Hello messages and the hold time.

By default, Hello messages are sent every 5 seconds. The hold time is advertised in Hello messages and indicates to neighbors the length of time that they should consider the sender valid. The default hold time is three times the hello interval, or 15 seconds.

On very congested and large networks, the default hold time might not be sufficient time for all routers to receive hello packets from their neighbors. In this case, you might want to increase the hold time.

Procedure

	Command or Action	Purpose
Step 1	switch(config-if)# {ip ipv6} hello-interval eigrp instance-tag seconds	Configures the hello interval for an EIGRP routing process. The instance tag can be any case-sensitive, alphanumeric string up to 20 characters. The range is from 1 to 65535 seconds. The default is 5.
Step 2	switch(config-if)# {ip ipv6} hold-time eigrp instance-tag seconds	Configures the hold time for an EIGRP routing process. The instance tag can be any case-sensitive, alphanumeric string up to 20 characters. The range is from 1 to 65535 seconds.
Step 3	(Optional) switch(config-if)# show ip eigrp interface detail	Verifies the timer configuration.
Step 4	(Optional) switch(config-if)# copy running-config startup-config	Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

Example

The following example shows how to change the interval between Hello packets and the hold time:

```
switch(config)# interface ethernet 1/2
switch(config-if)# ip hello-interval eigrp Test1 30
switch(config-if)# ip hold-time eigrp Test1 30
switch(config-if)# show ip eigrp interface detail
switch(config-if)# copy running-config startup-config
```

Disabling Split Horizon

You can use split horizon to block route information from being advertised by a router out of any interface from which that information originated. Split horizon usually optimizes communications among multiple routing devices, particularly when links are broken.

By default, split horizon is enabled on all interfaces.

Procedure

	Command or Action	Purpose
Step 1	switch(config-if)# no {ip ipv6} split-horizon eigrp instance-tag	Disables split horizon.
Step 2	(Optional) switch(config-if)# copy running-config startup-config	Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

Example

The following example shows how to disable split horizon on a particular interface:

```
switch(config)# interface ethernet 1/2
switch(config-if)# no ip split-horizon eigrp Test1
switch(config-if)# copy running-config startup-config
```

Enabling Wide Metrics

You can enable wide metrics in router or address family configuration mode.

Procedure

	Command or Action	Purpose
Step 1	switch(config-router)# metrics version 64bit	Enables 64-bit metric values.
Step 2	switch(config-router)# metrics rib-scale value	(Optional) Configures the scaling factor used to convert the 64-bit metric values to 32 bit in the RIB. The range is from 1 to 255. The default value is 128.
Step 3	(Optional) switch(config-router)# copy running-config startup-config	Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

Example

This example shows how to enable wide metrics in router configuration mode:

```
switch# configure terminal
switch(config)# router eigrp Test1
```

```
switch(config-router)# metrics version 64bit
switch(config-router)# metrics rib-scale 128
switch(config-router)# copy running-config startup-config
```

Tuning EIGRP

You can configure optional parameters to tune EIGRP for your network. Some of the parameters can be configured in address-family configuration mode, and others can be configured in interface configuration mode.

Procedure

	Command or Action	Purpose
Step 1	(Optional) switch(config-router-af)# default-information originate [always route-map <i>map-name</i>]	Originates or accepts the default route with prefix 0.0.0.0/0. When a route-map is supplied, the default route is originated only when the route map yields a true condition. The route-map name can be any case-sensitive, alphanumeric string up to 20 characters.
Step 2	(Optional) switch(config-router-af)# distance <i>internal external</i>	Configures the administrative distance for this EIGRP process. The range is from 1 to 255. The internal value sets the distance for routes learned from within the same autonomous system (the default value is 90). The external value sets the distance for routes learned from an external autonomous system (the default value is 170).
Step 3	(Optional) switch(config-router-af)# metric max-hops <i>hop-count</i>	Sets the maximum allowed hops for an advertised route. Routes over this maximum are advertised as unreachable. The range is from 1 to 255. The default is 100.
Step 4	(Optional) switch(config-router-af)# metric weights <i>tos k1 k2 k3 k4 k5 k6</i>	Adjusts the EIGRP metric or K value. EIGRP uses the following formula to determine the total metric to the network: $\text{metric} = [k1 \times \text{bandwidth} + (k2 \times \text{bandwidth}) / (256 - \text{load}) + k3 \times \text{delay} + k6 \times \text{extended attributes}] \times [k5 / (\text{reliability} + k4)]$ Default values and ranges are as follows: <ul style="list-style-type: none"> • TOS—0. The range is from 0 to 8. • k1—1. The range is from 0 to 255. • k2—0. The range is from 0 to 255. • k3—1. The range is from 0 to 255. • k4—0. The range is from 0 to 255.

	Command or Action	Purpose
		<ul style="list-style-type: none"> • k5—0. The range is from 0 to 255. • k6—0. The range is from 0 to 255.
Step 5	(Optional) switch(config-router-af)# timers active-time { <i>time-limit</i> disabled }	Sets the time the router waits in minutes (after sending a query) before declaring the route to be stuck in the active (SIA) state. The range is from 1 to 65535. The default is 3.
Step 6	switch(config-router-af)# exit	Exits address-family configuration mode.
Step 7	switch(config-router)# exit	Exits router configuration mode.
Step 8	switch(config)# interface ethernet <i>slot/port</i>	Enters interface configuration mode.
Step 9	(Optional) switch(config-if)# { ip ipv6 } bandwidth eigrp <i>instance-tag bandwidth</i>	Configures the bandwidth metric for EIGRP on an interface. The instance tag can be any case-sensitive, alphanumeric string up to 20 characters. The bandwidth range is from 1 to 2,560,000,000 Kb/s.
Step 10	(Optional) switch(config-if)# { ip ipv6 } bandwidth-percent eigrp <i>instance-tag percent</i>	Configures the percentage of bandwidth that EIGRP might use on an interface. The instance tag can be any case-sensitive, alphanumeric string up to 20 characters. The percent range is from 0 to 100. The default is 50.
Step 11	(Optional) switch(config-if)# no { ip ipv6 } delay eigrp <i>instance-tag delay</i>	Configures the delay metric for EIGRP on an interface. The instance tag can be any case-sensitive, alphanumeric string up to 20 characters. The delay range is from 1 to 16777215 (in tens of microseconds).
Step 12	(Optional) switch(config-if)# { ip ipv6 } distribute-list eigrp <i>instance-tag</i> { prefix-list <i>name</i> route-map <i>map-name</i> } { in out }	Configures the route filtering policy for EIGRP on this interface. The instance tag, prefix list name, and route-map name can be any case-sensitive, alphanumeric string up to 20 characters.
Step 13	(Optional) switch(config-if)# no { ip ipv6 } next-hop-self eigrp <i>instance-tag</i>	Configures EIGRP to use the received next-hop address rather than the address for this interface. The default is to use the IP address of this interface for the next-hop address. The instance tag can be any case-sensitive, alphanumeric string up to 20 characters.
Step 14	(Optional) switch(config-if)# { ip ipv6 } offset-list eigrp <i>instance-tag</i> { prefix-list <i>name</i> route-map <i>map-name</i> } { in out } <i>offset</i>	Adds an offset to incoming and outgoing metrics to routes learned by EIGRP. The instance tag, prefix list name, and route-map name can be any case-sensitive, alphanumeric string up to 20 characters.

	Command or Action	Purpose
Step 15	(Optional) switch(config-if)# {ip ipv6} passive-interface eigrp instance-tag	Suppresses EIGRP hellos, which prevents neighbors from forming and sending routing updates on an EIGRP interface. The instance tag can be any case-sensitive, alphanumeric string up to 20 characters.
Step 16	(Optional) switch(config-if)# copy running-config startup-config	Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

Example

The following example shows how to configure optional parameters in address-family configuration mode to tune EIGRP for your network:

```
switch# configure terminal
switch(config)# router eigrp Test1
switch(config-router)# address-family ipv4 unicast
switch(config-router-af)# default-information originate always
switch(config-router-af)# distance 25 100
switch(config-router-af)# metric max-hops 70
switch(config-router-af)# metric weights 0 1 3 2 1 0
switch(config-router-af)# timers active-time 200
switch(config-router-af)# copy running-config startup-config
```

The following example shows how to configure optional parameters in interface configuration mode to tune EIGRP for your network:

```
switch# configure terminal
switch(config)# interface ethernet 1/2
switch(config-if)# ip bandwidth eigrp Test1 30000
switch(config-if)# ip bandwidth-percent eigrp Test1 30
switch(config-if)# ip delay eigrp Test1 100
switch(config-if)# ip distribute-list eigrp Test1 route-map EigrpTest in
switch(config-if)# ip next-hop-self eigrp Test1
switch(config-if)# ip offset-list eigrp Test1 prefix-list EigrpList in
switch(config-if)# ip passive-interface eigrp Test1
switch(config-if)# copy running-config startup-config
```

Configuring Virtualization for EIGRP

You can configure multiple EIGRP processes in each VDC. You can also create multiple VRFs within each VDC and use the same or multiple EIGRP processes in each VRF. You assign an interface to a VRF.



Note Configure all other parameters for an interface after you configure the VRF for an interface. Configuring a VRF for an interface deletes all other configuration for that interface.

Before you begin

- Ensure that you have enabled the EIGRP feature.
- Create the VDCs and VRFs.
- Confirm that you are in the correct VDC. To change the VDC, use the **switchto vdc** command.

Procedure

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# vrf context <i>vrf-name</i>	Creates a new VRF and enters VRF configuration mode. The VRF name can be any case-sensitive, alphanumeric string up to 20 characters.
Step 3	switch(config-vrf)# router eigrp <i>instance-tag</i>	Creates a new EIGRP process with the configured instance tag. The instance tag can be any case-sensitive, alphanumeric string up to 20 characters. If you configure an instance tag that does not qualify as an AS number, you must use the autonomous-system command to configure the AS number explicitly or this EIGRP instance remains in the shutdown state.
Step 4	switch(config-router)# interface ethernet <i>slot/port</i>	Enters interface configuration mode. Use ? to find the slot and port ranges.
Step 5	switch(config-if)# vrf member <i>vrf-name</i>	Adds this interface to a VRF. The VRF name can be any case-sensitive, alphanumeric string up to 20 characters.
Step 6	switch(config-if)# { ip ipv6 } router eigrp <i>instance-tag</i>	Adds this interface to the EIGRP process. The instance tag can be any case-sensitive, alphanumeric string up to 20 characters.
Step 7	(Optional) switch(config-if)# copy running-config startup-config	Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

Example

This example shows how to create a VRF and add an interface to the VRF:

```
switch# configure terminal
switch(config)# vrf context NewVRF
switch(config-vrf)# router eigrp Test1
switch(config-router)# interface ethernet 1/2

switch(config-if)# ip router eigrp Test1
```

```
switch(config-if)# vrf member NewVRF
switch(config-if)# copy running-config startup-config
```

Verifying the EIGRP Configuration

Use one of the following commands to verify the configuration:

Command	Purpose
show {ip ipv6} eigrp [<i>instance-tag</i>]	Displays a summary of the configured EIGRP processes.
show {ip ipv6} eigrp [<i>instance-tag</i>] interfaces [<i>type number</i>] [brief] [detail]	Displays information about all configured EIGRP interfaces.
show {ip ipv6} eigrp <i>instance-tag</i> neighbors [<i>type number</i>] [detail]	Displays information about all the EIGRP neighbors. Use this command to verify the EIGRP neighbor configuration.
show {ip ipv6} eigrp [<i>instance-tag</i>] route [<i>ip-prefix/length</i>] [active] [all-links] [detail-links] [pending] [summary] [zero-successors] [vrf vrf-name]	Displays information about all the EIGRP routes.
show {ip ipv6} eigrp [<i>instance-tag</i>] topology [<i>ip-prefix/length</i>] [active] [all-links] [detail-links] [pending] [summary] [zero-successors] [vrf vrf-name]	Displays information about the EIGRP topology table.
show running-configuration eigrp	Displays the current running EIGRP configuration.

Displaying EIGRP Statistics

Use one of the following commands to display EIGRP statistics:

Command	Purpose
show {ip ipv6} eigrp [<i>instance-tag</i>] accounting [vrf vrf-name]	Displays accounting statistics for EIGRP.
show {ip ipv6} eigrp [<i>instance-tag</i>] route-map statistics redistribute	Displays redistribution statistics for EIGRP.
show {ip ipv6} eigrp [<i>instance-tag</i>] traffic [vrf vrf-name]	Displays traffic statistics for EIGRP.

Configuration Example for EIGRP

```
switch# configure terminal
switch(config)# feature eigrp
switch(config)# interface ethernet 1/2
```

```

switch(config-if)# ip address 192.0.2.55/24
switch(config-if)# ip router eigrp Test1
switch(config)# exit
switch(config)# no shutdown
switch(config)# router eigrp Test1
switch(config-router)# router-id 192.0.2.1

```

The following example shows how to use a route map with the distribute-list command to filter routes that are dynamically received from (or advertised to) EIGRP peers. The example configures a route table with a metric of 50, a source protocol of BGP, and an autonomous system number of 45000. When the match clauses is true, the tag value of the destination routing protocol is set to 5. The route map is used to distribute incoming packets for an EIGRP process

```

switch(config)# route-map metric-range
switch(config-route-map)# match metric 50
switch(config-route-map)# match source-protocol bgp 45000
switch(config-route-map)# set tag 5
switch(config-route-map)# exit
switch(config)# router eigrp 1
switch(config-router)# exit
switch(config)# interface ethernet 1/2
switch(config-if)# ip address 172.16.0.0
switch(config-if)# ip router eigrp 1
switch(config-if)# ip distribute-list eigrp 1 route-map metric-range in

```

The following example shows how to use a route map with the redistribute command to allow routes that are redistributed from the routing table to be filtered with a route map before being admitted into an EIGRP topology table. The example shows how to configure a route map to match EIGRP routes with a metric of 110, 200, or an inclusive range of 700 to 800. When the match clause is true, the tag value of the destination routing protocol is set to 10. The route map is used to redistribute EIGRP packets.

```

switch(config)# route-map metric-eigrp
switch(config-route-map)# match metric 110 200 750 +- 50
switch(config-route-map)# set tag 10
switch(config-route-map)# exit
switch(config)# router eigrp 1
switch(config-router)# redistribute eigrp route-map metric-eigrp
switch(config-router)# exit
switch(config)# interface ethernet 1/2
switch(config-if)# ip address 172.16.0.0
switch(config-if)# ip router eigrp 1

```

Related Documents for EIGRP

Related Topic	Document Title
EIGRP CLI commands	<i>Cisco Nexus 7000 Series NX-OS Unicast Routing Command Reference</i>
VDCs and VRFs	<i>Cisco Nexus 7000 Series NX-OS Virtual Device Context Configuration Guide</i>
EIGRP overview	Introduction to EIGRP Tech Note
EIGRP FAQs	EIGRP Frequently Asked Questions

MIBs

MIBs	MIBs Link
CISCO-EIGRP-MIB	To locate and download MIBs, go to the following URL: https://cfng.cisco.com/mibs .

Feature History for EIGRP

This table includes only the updates for those releases that have resulted in additions or changes to the feature.

Table 20: Feature History for EIGRP

Feature Name	Release	Feature Information
EIGRP	6.2(2)	Added support for route-map filtering.
EIGRP	6.2(2)	Added support for configuring the administrative distance of routes.
EIGRP	6.2(2)	Added the ability to configure all EIGRP interfaces as passive by default.
Wide metrics	5.2(1)	Added support for EIGRP wide metrics.
BFD	5.0(2)	Added support for BFD. See the <i>Cisco Nexus 7000 Series NX-OS Interfaces Configuration Guide</i> for more information.
Graceful shutdown	4.2(1)	Added support to gracefully shut down an EIGRP instance or EIGRP on an interface but preserve the EIGRP configuration.
EIGRP instance tag	4.2(1)	Changed the length to 20 characters.
Limits on redistributed routes	4.2(1)	Added support for limiting the number of redistributed routes.
EIGRP IPv6 support	4.1(2)	Added support for IPv6.
Authentication	4.0(3)	Added the ability to configure authentication within a VRF for EIGRP.
EIGRP	4.0(1)	This feature was introduced.



CHAPTER 10

Configuring IS-IS

This chapter contains the following sections:

- [Finding Feature Information, on page 227](#)
- [Information About IS-IS, on page 227](#)
- [Prerequisites for IS-IS, on page 232](#)
- [Guidelines and Limitations for IS-IS, on page 232](#)
- [Default Settings for IS-IS, on page 232](#)
- [Configuring IS-IS, on page 233](#)
- [Monitoring IS-IS, on page 249](#)
- [Configuration Examples for IS-IS, on page 250](#)
- [Related Documents for IS-IS, on page 250](#)
- [Standards for IS-IS, on page 250](#)
- [Feature History for IS-IS, on page 250](#)

Finding Feature Information

Your software release might not support all the features documented in this module. For the latest caveats and feature information, see the Bug Search Tool at <https://tools.cisco.com/bugsearch/> and the release notes for your software release. To find information about the features documented in this module, and to see a list of the releases in which each feature is supported, see the "New and Changed Information" chapter or the Feature History table in this chapter.

Information About IS-IS

IS-IS is an Interior Gateway Protocol (IGP) based on Standardization (ISO)/International Engineering Consortium (IEC) 10589. Cisco NX-OS supports Internet Protocol version 4 (IPv4), and beginning with Cisco NX-OS Release 6.1, Cisco NX-OS supports IPv6. IS-IS is a dynamic link-state routing protocol that can detect changes in the network topology and calculate loop-free routes to other nodes in the network. Each router maintains a link-state database that describes the state of the network and sends packets on every configured link to discover neighbors. IS-IS floods the link-state information across the network to each neighbor. The router also sends advertisements and updates on the link-state database through all the existing neighbors.

IS-IS Overview

IS-IS sends a hello packet out every configured interface to discover IS-IS neighbor routers. The hello packet contains information, such as the authentication, area, and supported protocols, which the receiving interface uses to determine compatibility with the originating interface. The hello packets are also padded to ensure that IS-IS establishes adjacencies only with interfaces that have matching maximum transmission unit (MTU) settings. Compatible interfaces form adjacencies, which update routing information in the link-state database through link-state update messages (LSPs). By default, the router sends a periodic LSP refresh every 10 minutes and the LSPs remain in the link-state database for 20 minutes (the LSP lifetime). If the router does not receive an LSP refresh before the end of the LSP lifetime, the router deletes the LSP from the database.

The LSP interval must be less than the LSP lifetime or the LSPs time out before they are refreshed.

IS-IS sends periodic hello packets to adjacent routers. If you configure transient mode for hello packets, these hello packets do not include the excess padding used before IS-IS establishes adjacencies. If the MTU value on adjacent routers changes, IS-IS can detect this change and send padded hello packets for a period of time. IS-IS uses this feature to detect mismatched MTU values on adjacent routers.

IS-IS Areas


You can design IS-IS networks as a single area that includes all routers in the network or as multiple areas that connect into a backbone or Level 2 area. Routers in a nonbackbone area are Level 1 routers that establish adjacencies within a local area (intra-area routing). Level 2 area routers establish adjacencies to other Level 2 routers and perform routing between Level 1 areas (inter-area routing). A router can have both Level 1 and Level 2 areas configured. These Level 1/Level 2 routers act as area border routers that route information from the local area to the Level 2 backbone area.

Within a Level 1 area, routers know how to reach all other routers in that area. The Level 2 routers know how to reach other area border routers and other Level 2 routers. Level 1/Level 2 routers straddle the boundary between two areas, routing traffic to and from the Level 2 backbone area. Level 1/Level 2 routers use the attached (ATT) bit signal Level 1 routers to set a default route to this Level 1/Level 2 router to connect to the Level 2 area.

In some instances, such as when you have two or more Level 1/Level 2 routers in an area, you may want to control which Level 1/Level 2 router that the Level 1 routers use as the default route to the Level 2 area. You can configure which Level 1/Level 2 router sets the attached bit.

Each IS-IS instance in Cisco NX-OS supports either a single Level 1 or Level 2 area, or one of each. By default, all IS-IS instances automatically support Level 1 and Level 2 routing.

Figure 31: IS-IS Network Divided into Areas



An autonomous system boundary router (ASBR) advertises external destinations throughout the IS-IS autonomous system. External routes are the routes redistributed into IS-IS from any other protocol.

NET and System ID

Each IS-IS instance has an associated network entity title (NET). The NET is comprised of the IS-IS system ID, which uniquely identifies this IS-IS instance in the area and the area ID. For example, if the NET is 47.0004.004d.0001.0001.0c11.1111.00, the system ID is 0000.0c11.1111.00 and the area is ID 47.0004.004d.0001.

Designated Intermediate System

IS-IS uses a designated intermediate system (DIS) in broadcast networks to prevent each router from forming unnecessary links with every other router on the broadcast network. IS-IS routers send LSPs to the DIS, which manages all the link-state information for the broadcast network. You can configure the IS-IS priority that IS-IS uses to select the DIS in an area.



Note No DIS is required on a point-to-point network.

IS-IS Authentication

You can configure authentication to control adjacencies and the exchange of LSPs. Routers that want to become neighbors must exchange the same password for their configured level of authentication. IS-IS blocks a router that does not have the correct password. You can configure IS-IS authentication globally or for an individual interface for Level 1, Level 2, or both Level 1/Level 2 routing.

IS-IS supports the following authentication methods:

- Clear text—All packets exchanged carry a cleartext 128-bit password.
- MD5 digest—All packets exchanged carry a message digest that is based on a 128-bit key.

To provide protection against passive attacks, IS-IS never sends the MD5 secret key as cleartext through the network. In addition, IS-IS includes a sequence number in each packet to protect against replay attacks.

You can use also keychains for hello and LSP authentication. See the *Cisco Nexus 7000 Series NX-OS Security Configuration Guide*, for information on keychain management.

Mesh Groups

A mesh group is a set of interfaces in which all routers reachable over the interfaces have at least one link to every other router. Many links can fail without isolating one or more routers from the network.

In normal flooding, an interface receives a new LSP and floods the LSP out over all other interfaces on the router. With mesh groups, when an interface that is part of a mesh group receives a new LSP, the interface does not flood the new LSP over the other interfaces that are part of that mesh group.



Note You may want to limit LSPs in certain mesh network topologies to improve network scalability. Limiting LSP floods might also reduce the reliability of the network (in case of failures). For this reason, we recommend that you use mesh groups only if specifically required, and then only after you make a careful network design.

You can also configure mesh groups in block mode for parallel links between routers. In this mode, all LSPs are blocked on that interface in a mesh group after the routers initially exchange their link-state information.

Overload Bit

IS-IS uses the overload bit to tell other routers not to use the local router to forward traffic but to continue routing traffic destined for that local router.

You may want to use the overload bit in these situations:

- The router is in a critical condition.
- Graceful introduction and removal of the router to/from the network.
- Other (administrative or traffic engineering) reasons such as waiting for BGP convergence.

Route Summarization

You can configure a summary aggregate address. Route summarization simplifies route tables by replacing a number of more-specific addresses with an address that represents all the specific addresses. For example, you can replace 10.1.1.0/24, 10.1.2.0/24, and 10.1.3.0/24 with one summary address, 10.1.0.0/16.

If more specific routes are in the routing table, IS-IS advertises the summary address with a metric equal to the minimum metric of the more specific routes.



Note Cisco NX-OS does not support automatic route summarization.

Route Redistribution

You can use IS-IS to redistribute static routes, routes learned by other IS-IS autonomous systems, or routes from other protocols. You must configure a route map with the redistribution to control which routes are passed into IS-IS. A route map allows you to filter routes based on attributes such as the destination, origination protocol, route type, route tag, and so on.

Whenever you redistribute routes into an IS-IS routing domain, Cisco NX-OS does not, by default, redistribute the default route into the IS-IS routing domain. You can generate a default route into IS-IS, which can be controlled by a route policy.

You also configure the default metric that is used for all imported routes into IS-IS.

Administrative Distance

The administrative distance is a rating of the trustworthiness of a routing information source. A higher value indicates a lower trust rating. The administrative distance is used to discriminate between routes learned from more than one routing protocol. The route with the lowest administrative distance is installed in the IP routing table.

You can configure the administrative distance for internal and external routes based on various match criteria for a given prefix. Routing protocols such as IS-IS configure the prefix into the Routing Information Base (RIB), along with the next hops based on these metrics. If multiple paths are available for a prefix, the routing

protocol chooses the best path based on the cost to reach the next hop and the administrative distance. Beginning with Cisco NX-OS Release 6.2(2), you can specify that prefixes be considered based on specific routes. In prior releases, one administrative distance was sufficient for all internal routes.

Load Balancing

You can use load balancing to allow a router to distribute traffic over all the router network ports that are the same distance from the destination address. Load balancing increases the utilization of network segments and increases the effective network bandwidth.

Cisco NX-OS supports the Equal Cost Multiple Paths (ECMP) feature with up to 16 equal-cost paths in the IS-IS route table and the unicast RIB. You can configure IS-IS to load balance traffic across some or all of those paths.

BFD

This feature supports bidirectional forwarding detection (BFD). BFD is a detection protocol designed to provide fast forwarding-path failure detection times. BFD provides subsecond failure detection between two adjacent devices and can be less CPU-intensive than protocol hello messages because some of the BFD load can be distributed onto the data plane on supported modules. See the *Cisco Nexus 7000 Series NX-OS Interfaces Configuration Guide*, for more information.

Virtualization Support

Cisco NX-OS supports multiple instances of the IS-IS protocol that runs on the same system. IS-IS supports virtual routing and forwarding (VRF) instances. VRFs exist within virtual device contexts (VDCs). You can configure up to four IS-IS instances in a VDC.

By default, Cisco NX-OS places you in the default VDC and default VRF unless you specifically configure another VDC and VRF. See the *Cisco Nexus 7000 Series NX-OS Virtual Device Context Configuration Guide*.

High Availability and Graceful Restart

Cisco NX-OS provides a multilevel high-availability architecture. IS-IS supports stateful restart, which is also referred to as non-stop routing (NSR). If IS-IS experiences problems, it attempts to restart from its previous run-time state. The neighbors would not register any neighbor event in this case. If the first restart is not successful and another problem occurs, IS-IS attempts a graceful restart as per RFC 3847. A graceful restart, or non-stop forwarding (NSF), allows IS-IS to remain in the data forwarding path through a process restart. When the restarting IS-IS interface is operational again, it rediscovers its neighbors, establishes adjacency, and starts sending its updates again. At this point, the NSF helpers recognize that the graceful restart has finished.

A stateful restart is used in the following scenarios:

- First recovery attempt after process experiences problems
- ISSU
- User-initiated switchover using the **system switchover** command

A graceful restart is used in the following scenarios:

- Second recovery attempt after the process experiences problems within a 4-minute interval
- Manual restart of the process using the **restart isis** command
- Active supervisor removal
- Active supervisor reload using the **reload module active-sup** command



Note Graceful restart is on by default, and we strongly recommended that it not be disabled.

Multiple IS-IS Instances

Cisco NX-OS supports a maximum of four instances of the IS-IS protocol that run on the same node. You cannot configure multiple instances over the same interface. Every instance uses the same system router ID.

Prerequisites for IS-IS

IS-IS has the following prerequisites:

- You must enable IS-IS.

Guidelines and Limitations for IS-IS

IS-IS has the following configuration guidelines and limitations:

- You can configure a maximum of four IS-IS instances per VDC.
- Because the default reference bandwidth is different for Cisco NX-OS and Cisco IOS, the advertised tunnel IS-IS metric is different for these two operating systems.
- For the IS-IS Multitopology feature, one topology for IPv4 and one for IPv6 is supported.
- Unlike IOS, NXOS-ISIS works even when there is a change in bandwidth. It causes an SPF and routes updates. This result in an excessive packet drop, but port P0 continues to be active.
- If you are familiar with the Cisco IOS CLI, be aware that the Cisco NX-OS commands for this feature might differ from the Cisco IOS commands that you would use.

Default Settings for IS-IS

Table 21: Default IS-IS Parameters

Parameters	Default
Administrative distance	115

Parameters	Default
Area level	Level-1-2
DIS priority	64
Graceful restart	Enabled
Hello multiplier	3
Hello padding	Enabled
Hello time	10 seconds
IS-IS feature	Disabled
LSP interval	33
LSP MTU	1492
Maximum LSP lifetime	1200 seconds
Maximum paths	4
Metric	40
Reference bandwidth	40 Gbps

Configuring IS-IS

IS-IS Configuration Modes

Router Configuration Mode Example

This example shows how to enter router configuration mode:

```
switch#: configure terminal
switch(config)# router isis isp
switch(config-router)#
```



Note From a mode, you can enter the ? command to display the commands available in that mode.

Router Address Family Configuration Mode Example

This example shows how to enter router address family configuration mode:

```
switch(config)# router isis isp
switch(config-router)# address-family ipv4 unicast
switch(config-router-af)#
```



Note From a mode, you can enter the ? command to display the commands available in that mode.

Enabling the IS-IS Feature

You must enable the IS-IS feature before you can configure IS-IS.

Before you begin

Ensure that you are in the correct VDC (or use the **switchto vdc** command).

Procedure

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# [no] feature isis	Enables the IS-IS feature.
Step 3	(Optional) switch(config)# show feature	Displays enabled and disabled features.
Step 4	(Optional) switch(config)# copy running-config startup-config	Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

Creating an IS-IS Instance

You can create an IS-IS instance and configure the area level for that instance.

You must remove any IS-IS commands that are configured in interface mode to completely remove all configuration for the IS-IS instance

Before you begin

You must enable IS-IS.

Ensure that you are in the correct VDC (or use the **switchto vdc** command).

Procedure

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# router isis instance-tag	Creates a new IS-IS instance with the configured instance tag.
Step 3	switch(config-router)# net network-entity-title	Configures the NET for this IS-IS instance.
Step 4	(Optional) switch(config-router)# is-type {level-1 level-2 level-1-2}	Configures the area level for this IS-IS instance. The default is level-1-2.

	Command or Action	Purpose
Step 5	(Optional) switch(config)# show isis {vrf vrf-name} process	Displays a summary of IS-IS information for all IS-IS instances.
Step 6	(Optional) switch(config-router)# distance value	Sets the administrative distance for IS-IS. The range is from 1 to 255. The default is 115.
Step 7	(Optional) switch(config-router)# log-adjacency-changes	Sends a system message whenever an IS-IS neighbor changes the state.
Step 8	(Optional) switch(config-router)# lsp-mtu size	Sets the MTU for LSPs in this IS-IS instance. The range is from 128 to 4352 bytes. The default is 1492.
Step 9	(Optional) switch(config-router)# maximum-paths number	Configures the maximum number of equal-cost paths that IS-IS maintains in the route table. The range is from 1 to 16. The default is 4.
Step 10	(Optional) switch(config-router)# reference-bandwidth bandwidth value {Mbps Gbps}	Sets the default reference bandwidth used for calculating the IS-IS cost metric. The range is from 1 to 4000 Gbps. The default is 40 Gbps.
Step 11	(Optional) switch(config-if)# clear isis [instance-tag] adjacency [* system-id interface]	Clears neighbor statistics and removed adjacencies for this IS-IS instance.
Step 12	(Optional) switch(config)# copy running-config startup-config	Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

Example

The following example shows how to create an IS-IS instance in a level 2 area:

```
switch# configure terminal
switch(config)# router isis Enterprise
switch(config-router)# net 47.0004.004d.0001.0001.0c11.1111.00
switch(config-router)# is-type level 2
switch(config-router)# copy running-config startup-config
```

Restarting an IS-IS Instance

You can restart an IS-IS instance. This action clears all neighbors for the instance.

To restart an IS-IS instance and remove all associated neighbors, use the following command:

Procedure

	Command or Action	Purpose
Step 1	switch(config)# restart isis instance-tag	Restarts the IS-IS instance and removes all neighbors.

Shutting Down IS-IS

You can shut down the IS-IS instance. This action disables this IS-IS instance and retains the configuration. To shut down the IS-IS instance, use the following command in router configuration mode:

Procedure

	Command or Action	Purpose
Step 1	switch(config-router)# shutdown	Disables the IS-IS instance.

Configuring IS-IS on an Interface

You can add an interface to an IS-IS instance.

Before you begin

You must enable IS-IS.

Ensure that you are in the correct VDC (or use the **switchto vdc** command).

Procedure

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# interface <i>interface-type slot/port</i>	Enters interface configuration mode.
Step 3	(Optional) switch(config-if)# medium { broadcast p2p }	Configures the broadcast or point-to-point mode for the interface. IS-IS inherits this mode.
Step 4	switch(config-if)# { ip ipv6 } router isis <i>instance-tag</i>	Associates this IPv4 or IPv6 interface with an IS-IS instance.
Step 5	(Optional) switch(config-if)# show isis [vrf <i>vrf-name</i>] [<i>instance-tag</i>] interface [<i>interface-type slot/port</i>]	Displays IS-IS information for an interface.
Step 6	(Optional) switch(config-if)# isis circuit-type { level-1 level-2 level-1-2 }	Sets the type of adjacency that this interface participates in. Use this command only for routers that participate in both Level 1 and Level 2 areas.
Step 7	(Optional) isis metric <i>value</i> { level-1 level-2 }	Sets the IS-IS metric for this interface. The range is from 1 to 16777214. The default is 10.
Step 8	(Optional) switch(config-if)# isis passive <i>value</i> { level-1 level-2 level-1-2 }	Prevents the interface from forming adjacencies but still advertises the prefix associated with the interface.

	Command or Action	Purpose
Step 9	(Optional) switch(config)# copy running-config startup-config	Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

Example

This example shows how to add Ethernet 1/2 interface to an IS-IS instance:

```
switch# configure terminal
switch(config)# interface ethernet 1/2
switch(config-if)# ip router isis Enterprise
switch(config-if)# copy running-config startup-config
```

Configuring IS-IS Authentication in an Area

You can configure IS-IS to authenticate LSPs in an area.

Before you begin

You must enable IS-IS.

Ensure that you are in the correct VDC (or use the **switchto vdc** command).

Procedure

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# router isis <i>instance-tag</i>	Creates a new IS-IS instance with the configured instance tag.
Step 3	switch(config-router)# authentication-type {cleartext md5} {level-1 level-2}	Sets the authentication method used for a Level 1 or Level 2 area as cleartext or as an MD5 authentication digest.
Step 4	switch(config-router)# authentication key-chain <i>key</i> {level-1 level-2}	Configures the authentication key used for an IS-IS area-level authentication.
Step 5	(Optional) switch(config-router)# authentication-check {level-1 level-2}	Enables checking the authentication parameters in a received packet.
Step 6	(Optional) switch(config)# copy running-config startup-config	Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

Example

This example shows how to configure cleartext authentication on an IS-IS instance:

```
switch# configure terminal
switch(config)# router isis Enterprise
switch(config-router)# authentication-type cleartext level-2
switch(config-router)# authentication key-chain ISISKey level-2
switch(config-router)# copy running-config startup-config
```

Configuring IS-IS Authentication on an Interface

You can configure IS-IS to authenticate Hello packets on an interface.

Before you begin

You must enable IS-IS.

Ensure that you are in the correct VDC (or use the `switchto vdc` command).

Procedure

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# interface <i>interface-type</i> <i>slot/port</i>	Enters interface configuration mode.
Step 3	switch(config-if)# isis authentication-type { cleartext md5 } { level-1 level-2 }	Sets the authentication type for IS-IS on this interface as cleartext or as an MD5 authentication digest.
Step 4	switch(config-if)# isis authentication key-chain <i>key</i> { level-1 level-2 }	Configures the authentication key used for IS-IS on this interface.
Step 5	(Optional) isis authentication-check { level-1 level-2 }	Enables checking the authentication parameters in a received packet.
Step 6	(Optional) switch(config)# copy running-config startup-config	Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

Example

This example shows how to configure cleartext authentication on an IS-IS instance:

```
switch# configure terminal
switch(config)# interface ethernet 1/2
switch(config-if)# isis authentication-type cleartext level-2
switch(config-if)# isis authentication key-chain ISISKey
switch(config-if)# copy running-config startup-config
```

Configuring a Mesh Group

You can add an interface to a mesh group to limit the amount of LSP flooding for interfaces in that mesh group. You can optionally block all LSP flooding on an interface in a mesh group.

To add an interface to a mesh group, use the following command in interface configuration mode:

Procedure

	Command or Action	Purpose
Step 1	switch(config-if)# isis mesh-group { blocked <i>mesh-id</i> }	Adds this interface to a mesh group. The range is from 1 to 4294967295.

Configuring a Designated Intermediate System

You can configure a router to become the designated intermediate system (DIS) for a multiaccess network by setting the interface priority.

To configure the DIS, use the following command in interface configuration mode:

Procedure

	Command or Action	Purpose
Step 1	switch(config-if)# isis priority <i>number</i> { level-1 level-2 }	Sets the priority for DIS selection. The range is from 0 to 127. The default is 64.

Configuring Dynamic Host Exchange

You can configure IS-IS to map between the system ID and the hostname for a router using dynamic host exchange.

To configure dynamic host exchange, use the following command in router configuration mode:

Procedure

	Command or Action	Purpose
Step 1	switch(config-router)# hostname dynamic	Enables dynamic host exchange.

Setting the Overload Bit

You can configure the router to signal other routers not to use this router as an intermediate hop in their shortest path first (SPF) calculations. You can optionally configure the overload bit temporarily on startup, until BGP converges.

In addition to setting the overload bit, you might also want to suppress certain types of IP prefix advertisements from LSPs for Level 1 or Level 2 traffic.

To set the overload bit, use the following command in router configuration mode:

Procedure

	Command or Action	Purpose
Step 1	switch(config-router)# set-overload-bit { always on-startup { <i>seconds</i> wait-for bgp as-number }} [suppress [interlevel external]]	Sets the overload bit for IS-IS. The seconds range is from 5 to 86400.

Configuring the Attached Bit

You can configure the attached bit to control which Level 1/Level 2 router that the Level 1 routers use as the default route to the Level 2 area. If you disable setting the attached bit, the Level 1 routers do not use this Level 1/Level 2 router to reach the Level 2 area.

To configure the attached bit for a Level 1/Level 2 router, use the following command in router configuration mode:

Procedure

	Command or Action	Purpose
Step 1	switch(config-router)# [no] attached-bit	Configures the Level 1/Level 2 router to set the attached bit. This feature is enabled by default.

Configuring the Transient Mode for Hello Padding

You can configure the transient mode for hello padding to pad hello packets when IS-IS establishes adjacency and remove that padding after IS-IS establishes adjacency.

To configure the mode for hello padding, use the following command in router configuration mode:

Procedure

	Command or Action	Purpose
Step 1	switch(config-if)# [no] isis hello-padding	Pads the hello packet to the full MTU. The default is enabled. Use the no form of this command to configure the transient mode of hello padding.

Configuring a Summary Address

You can create aggregate addresses that are represented in the routing table by a summary address. One summary address can include multiple groups of addresses for a given level. Cisco NX-OS advertises the smallest metric of all the more-specific routes.

Before you begin

You must enable IS-IS.

Ensure that you are in the correct VDC (or use the **switchto vdc** command).

Procedure

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# router isis <i>instance-tag</i>	Creates a new IS-IS instance with the configured instance tag.
Step 3	switch(config-router)# address-family { ipv4 ipv6 } { unicast multicast }	Enters address family configuration mode.
Step 4	switch(config-router-af)# summary-address <i>ip-prefix/mask-len</i> { level-1 level-2 level-1-2 }	Configures a summary address for an IS-IS area for IPv4 or IPv6 addresses.
Step 5	(Optional) switch(config-if)# show isis [<i>vrfvrf-name</i>] { ip ipv6 } summary-address <i>ip-prefix</i> [longer-prefixes]	Displays IS-IS IPv4 or IPv6 summary address information.
Step 6	(Optional) switch(config)# copy running-config startup-config	Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

Example

This example shows how to configure an IPv4 unicast summary address for IS-IS:

```
switch# configure terminal
switch(config)# interface ethernet 1/2
switch(config-router)# address-family ipv4 unicast
switch(config-router-af)# summary-address 192.0.2.0/24 level-2
switch(config-router-af)# copy running-config startup-config
```

Configuring Redistribution

You can configure IS-IS to accept routing information from another routing protocol and redistribute that information through the IS-IS network. You can optionally assign a default route for redistributed routes.

Before you begin

You must enable IS-IS.

Ensure that you are in the correct VDC (or use the **switchto vdc** command).

Procedure

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# router isis <i>instance-tag</i>	Creates a new IS-IS instance with the configured instance tag.
Step 3	switch(config-router)# address-family { ipv4 ipv6 } unicast	Enters address family configuration mode.

	Command or Action	Purpose
Step 4	switch(config-router-af)# redistribute { bgp <i>as</i> { eigrp isis ospf ospfv3 rip } <i>instance-tag</i> static direct } route-map <i>map-name</i>	Redistributes routes from other protocols into IS-IS.
Step 5	(Optional) switch(config-router-af)# default-information originate [always] [route-map <i>map-name</i>]	Generates a default route into IS-IS.
Step 6	(Optional) switch(config-router-af)# distribute { level-1 level-2 } into { level-1 level-2 } { route-map <i>route-map</i> all }	Redistributes routes from one IS-IS level to the other IS-IS level.
Step 7	(Optional) switch(config-router-af)# show isis [vrf <i>vrf-name</i>] { ip ipv6 } route <i>ip-prefix</i> [<i>detail</i> longer-prefixes [summary detail]]	Shows the IS-IS routes.
Step 8	(Optional) switch(config)# copy running-config startup-config	Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

Example

This example shows how to redistribute EIGRP into IS-IS:

```
switch# configure terminal
switch(config)# interface ethernet 1/2
switch(config-router)# address-family ipv4 unicast
switch(config-router-af)# redistribute eigrp 201 route-map ISISmap
switch(config-router-af)# copy running-config startup-config
```

Limiting the Number of Redistributed Routes

Route redistribution can add many routes to the IS-IS route table. You can configure a maximum limit to the number of routes accepted from external protocols. IS-IS provides the following options to configure redistributed route limits:

- **Fixed limit**—Logs a message when IS-IS reaches the configured maximum. IS-IS does not accept any more redistributed routes. You can optionally configure a threshold percentage of the maximum where IS-IS logs a warning when that threshold is passed.
- **Warning only**—Logs a warning only when IS-IS reaches the maximum. IS-IS continues to accept redistributed routes.
- **Withdraw**—Starts the timeout period when IS-IS reaches the maximum. After the timeout period, IS-IS requests all redistributed routes if the current number of redistributed routes is less than the maximum limit. If the current number of redistributed routes is at the maximum limit, IS-IS withdraws all redistributed routes. You must clear this condition before IS-IS accepts more redistributed routes. You can optionally configure the timeout period.

Before you begin

You must enable IS-IS.

Ensure that you are in the correct VDC (or use the **switchto vdc** command).

Procedure

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# router eigrp <i>instance-tag</i>	Creates a new IS-IS instance with the configured instance tag.
Step 3	switch(config-router)# redistribute { bgp <i>ip</i> direct eigrp <i>id</i> isis <i>id</i> ospf <i>id</i> rip <i>id</i> static } route-map <i>map-name</i>	Redistributes the selected protocol into IS-IS through the configured route map.
Step 4	switch(config-router)# redistribute maximum-prefix <i>max</i> [<i>threshold</i>] [warning-only withdraw [<i>num-retries</i> <i>timeout</i>]]	Specifies a maximum number of prefixes that IS-IS distributes. The range is from 0 to 65536. You can optionally specify the following: <ul style="list-style-type: none"> • threshold—Percent of maximum prefixes that triggers a warning message. • warning-only—Logs a warning message when the maximum number of prefixes is exceeded. • withdraw—Withdraws all redistributed routes. You can optionally try to retrieve the redistributed routes. The num-retries range is from 1 to 12. The timeout is 60 to 600 seconds. The default is 300 seconds. Use the clear isis redistribution command if all routes are withdrawn.
Step 5	(Optional) switch(config-router)# show running-config isis	Displays the IS-IS configuration.
Step 6	(Optional) switch(config)# copy running-config startup-config	Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

Example

This example shows how to limit the number of redistributed routes into IS-IS:

```
switch# configure terminal
switch(config)# router eigrp isis Enterprise
switch(config-router)# redistribute bgp route-map FilterExternalBGP
switch(config-router)# redistribute maximum-prefix 1000 75
```

Configuring the Administrative Distance of Routes

You can set the administrative distance of routes added by IS-IS into the RIB.

Before you begin

You must enable IS-IS (see the “Enabling the IS-IS Feature” section on page 9-9).

Ensure that you are in the correct VDC (or use the **switchto vdc** command).

Procedure

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# router isis <i>instance-tag</i>	Creates a new IS-IS instance and enters router configuration mode.
Step 3	switch(config-router)# table-map <i>route-map-name</i> [filter]	Configures a table map with route map information. You can enter up to 63 alphanumeric characters for the map name. The filter keyword filters routes rejected by the route map and does not download them to the RIB.
Step 4	(Optional) switch(config)# copy running-config startup-config	Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

Disabling Strict Adjacency Mode

When both IPv4 and IPv6 address families are enabled, strict adjacency mode is enabled by default. In this mode, the device does not form an adjacency with any router that does not have both address families enabled. You can disable strict adjacency mode using the **no adjacency check** command.

Before you begin

You must enable IS-IS.

Ensure that you are in the correct VDC (or use the **switchto vdc** command).

Procedure

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# router isis <i>instance-tag</i>	Creates a new IS-IS instance with the configured instance tag.
Step 3	switch(config-router)# address-family ipv4 unicast	Enters address family configuration mode.

	Command or Action	Purpose
Step 4	switch(config-router-af)# no adjacency-check	Disables strict adjacency mode for the IPv6 address family.
Step 5	switch(config-router-af)# exit	Exits address family configuration mode.
Step 6	switch(config-router-af)# address-family ipv6 unicast	Enters address family configuration mode.
Step 7	switch(config-router-af)# no adjacency-check	Disables strict adjacency mode for the IPv6 address family.
Step 8	(Optional) switch(config-router-af)# show running-config isis	Displays the IS-IS configuration.
Step 9	(Optional) switch(config-router-af)# copy running-config startup-config	Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

Example

This example shows how to disable strict adjacency mode:

```
switch# configure terminal
switch(config)# router isis Enterprise
switch(config-router)# address-family ip4 unicast
switch(config-router-af)# no adjacency-check
switch(config-router)# exit
switch(config-router-af)# address-family ip6 unicast
switch(config-router-af)# no adjacency-check
switch(config-router-af)# show running-config isis
switch(config-router-af)# copy running-config startup-config
```

Configuring a Graceful Restart

You can configure a graceful restart for IS-IS.

Before you begin

You must enable IS-IS.

Create the VDCs and VRFs.

Ensure that you are in the correct VDC (or use the **switchto vdc** command).

Procedure

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# router isis instance-tag	Creates a new IS-IS process with the configured name.

	Command or Action	Purpose
Step 3	switch(config-router)# graceful-restart	Enables a graceful restart and the graceful restart helper functionality. Enabled by default.
Step 4	switch(config-router)# graceful-restart t3 manual time	Configures the graceful restart T3 timer. The range is from 30 to 65535 seconds. The default is 60.
Step 5	(Optional) switch(config-router)# show running-config isis	Displays the IS-IS configuration.
Step 6	(Optional) switch(config)# copy running-config startup-config	Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

Example

This example shows how to enable a graceful restart:

```
switch# configure terminal
switch(config)# router isis Enterprise
switch(config-router)# graceful restart
switch(config-router)# copy running-config startup-config
```

Configuring Virtualization

You can configure multiple IS-IS instances in each VDC. You can also create multiple VRFs within each VDC and use the same or multiple IS-IS instances in each VRF. You assign an IS-IS interface to a VRF.

You must configure a NET for the configured VRF.



Note Configure all other parameters for an interface after you configure the VRF for an interface. Configuring a VRF for an interface deletes all the configuration for that interface.

Before you begin

You must enable IS-IS.

Create the VDCs and VRFs.

Ensure that you are in the correct VDC (or use the **switchto vdc** command).

Procedure

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# vrf context vrf-name	Creates a new VRF and enters VRF configuration mode.

	Command or Action	Purpose
Step 3	switch(config)# exit	Exits VRF configuration mode.
Step 4	switch(config)# router isis <i>instance-tag</i>	Creates a new IS-IS instance with the configured instance tag.
Step 5	(Optional) switch(config-router)# vrf <i>vrf-name</i>	Enters VRF configuration mode.
Step 6	switch(config-router-vrf)# net <i>network-entity-title</i>	Configures the NET for this IS-IS instance.
Step 7	switch(config-router-vrf)# exit	Exits router VRF configuration mode.
Step 8	switch(config)# interface ethernet <i>slot/port</i>	Enters interface configuration mode.
Step 9	switch(config-if)# vrf member <i>vrf-name</i>	Adds this interface to a VRF.
Step 10	switch(config-if)# ip address <i>ip-prefix/length</i>	Configures an IP address for this interface. You must do this step after you assign this interface to a VRF.
Step 11	switch(config-if)# ip router isis <i>instance-tag</i>	Associates this IPv4 interface with an IS-IS instance.
Step 12	(Optional) switch(config-if)# show isis [vrf <i>vrf-name</i>] [<i>instance-tag</i>] interface [<i>interface-type slot/port</i>]	Displays IS-IS information for an interface. in a VRF.
Step 13	(Optional) switch(config)# copy running-config startup-config	Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

Example

The following example shows how to create a VRF and add an interface to the VRF:

```
switch# configure terminal
switch(config)# vrf context NewVRF
switch(config-vrf)# exit
switch(config)# router isis Enterprise
switch(config-router)# vrf NewVRF
switch(config-router-vrf)# net 47.0004.004d.0001.0001.0c11.1111.00
switch(config-router-vrf)# interface ethernet 1/2
switch(config-if)# vrf member NewVRF
switch(config-if)# ip address 192.0.2.1/16
switch(config-if)# ip router isis Enterprise
switch(config-if)# copy running-config startup-config
```

Tuning IS-IS

You can tune IS-IS to match your network requirements.

You can use the following optional commands in router configuration mode to tune IS-IS:

Procedure

	Command or Action	Purpose
Step 1	switch(config-router)# lsp-gen-interval [level-1 level-2] <i>lsp-max-wait</i> [<i>lsp-initial-wait</i> <i>lsp-second-wait</i>]	Configures the IS-IS throttle for LSP generation. The optional parameters are as follows: <ul style="list-style-type: none"> • <i>lsp-max-wait</i>—The maximum wait between the trigger and LSP generation. The range is from 500 to 65535 milliseconds. • <i>lsp-initial-wait</i>—The initial wait between the trigger and LSP generation. The range is from 50 to 65535 milliseconds. • <i>lsp-second-wait</i>—The second wait used for LSP throttle during backoff. The range is from 50 to 65535 milliseconds.
Step 2	switch(config-router)# max-lsp-lifetime <i>lifetime</i>	Sets the maximum LSP lifetime in seconds. The range is from 1 to 65535. The default is 1200.
Step 3	switch(config-router)# metric-style transition	Enables IS-IS to generate and accept both narrow metric-style Type Length Value (TLV) objects and wide metric-style TLV objects. The default is disabled.
Step 4	switch(config-router)# spf-interval [level-1 level-2] <i>spf-max-wait</i> [<i>spf-initial-wait</i> <i>spf-second-wait</i>]	Configures the interval between LSA arrivals. The optional parameters are as follows: <ul style="list-style-type: none"> • <i>lsp-max-wait</i>—The maximum wait between the trigger and SPF computation. The range is from 500 to 65535 milliseconds. • <i>lsp-initial-wait</i>—The initial wait between the trigger and SPF computation. The range is from 50 to 65535 milliseconds. • <i>lsp-second-wait</i>—The second wait used for SPF computation during backoff. The range is from 50 to 65535 milliseconds.
Step 5	(Optional) switch(config-router-af)# adjacency-check	Performs an adjacency check to verify that an IS-IS instance forms an adjacency only with a remote IS-IS entity that supports the same address family. This command is enabled by default.
Step 6	(Optional) switch(config-if)# isis csnp-interval <i>seconds</i> [level-1 level-2]	Sets the complete sequence number PDU (CNSP) interval in seconds for IS-IS. The range is from 1 to 65535. The default is 10.

	Command or Action	Purpose
Step 7	(Optional) switch(config-if)# isis hello-interval <i>seconds</i> [level-1 level-2]	Sets the hello interval in seconds for IS-IS. The range is from 1 to 65535. The default is 10.
Step 8	(Optional) switch(config-if)# isis hello-multiplier <i>num</i> [level-1 level-2]	Specifies the number of IS-IS hello packets that a neighbor must miss before the router tears down an adjacency. The range is from 3 to 1000. The default is 3.
Step 9	(Optional) switch(config-if)# isis lsp-interval <i>milliseconds</i>	Sets the interval in milliseconds between LSPs sent on this interface during flooding. The range is from 10 to 65535. The default is 33.

Monitoring IS-IS

To display IS-IS statistics, use the following commands:

Command	Purpose
show isis [<i>instance-tag</i>] adjacency [<i>interface</i>] [system-ID] [detail] [summary] [vrf <i>vrf-name</i>]	Displays the IS-IS adjacency statistics.
show isis [<i>instance-tag</i>] database [level-1 level-2] [detail] [summary] [<i>lsip</i>] {{ adjacency id { ip ipv6 } prefix <i>prefix</i> } [router-id <i>id</i>] [<i>zero-sequence</i>]} [vrf <i>vrf-name</i>]	Displays the IS-IS database statistics.
show isis [<i>instance-tag</i>] statistics [<i>interface</i>] [vrf <i>vrf-name</i>]	Displays the IS-IS interface statistics.
show isis ip route-map statistics redistribute { bgp <i>id</i> igrp <i>id</i> isis <i>id</i> ospf <i>id</i> rip <i>id</i> static } [vrf <i>vrf-name</i>]	Displays the IS-IS redistribution statistics.
show isis ip route-map statistics distribute { level-1 level-2 } into { level-1 level-2 } [vrf <i>vrf-name</i>]	Displays IS-IS distribution statistics for routes distributed between levels.
show isis [<i>instance-tag</i>] spf-log [detail] [vrf <i>vrf-name</i>]	Displays the IS-IS SPF calculation statistics.
show isis [<i>instance-tag</i>] traffic [<i>interface</i>] [vrf <i>vrf-name</i>]	Displays the IS-IS traffic statistics.

To clear IS-IS configuration statistics, perform one of the following tasks:

Command	Purpose
clear isis [<i>instance-tag</i>] adjacency [* [<i>interface</i>]] [system-id <i>id</i>] [vrf <i>vrf-name</i>]	Clears the IS-IS adjacency statistics.
clear { ip ipv6 } route map statistics { bgp <i>id</i> igrp <i>id</i> isis <i>id</i> ospf <i>id</i> rip <i>id</i> static } [vrf <i>vrf-name</i>]	Clears the IS-IS redistribution statistics

Command	Purpose
clear isis route-map statistics distribute {level-1 level-2} into {level-1 level-2} [vrf vrf-name]	Clears IS-IS distribution statistics for routes distributed between levels.
clear isis [instance-tag] statistics [* interface] [vrf vrf-name]	Clears the IS-IS interface statistics.
clear isis [instance-tag] traffic [* interface] [vrf vrf-name]	Clears the IS-IS traffic statistics.

Configuration Examples for IS-IS

The following example shows how to configure IS-IS:

```
router isis Enterprise
 is-type level-1
 net 49.0001.0000.0000.0003.00
 graceful-restart
 address-family ipv4 unicast
 default-information originate

interface ethernet 2/1
 ip address 192.0.2.1/24
 isis circuit-type level-1
 ip router isis Enterprise
```

Related Documents for IS-IS

Related Topic	Document Title
IS-IS CLI commands	<i>Cisco Nexus 7000 Series NX-OS Unicast Routing Command Reference</i>
VDCs and VRFs	<i>Cisco Nexus 7000 Series NX-OS Virtual Device Context Command Reference</i>

Standards for IS-IS

No new or modified standards are supported by this feature, and support for existing standards has not been modified by this feature.

Feature History for IS-IS

This table includes only the updates for those releases that have resulted in additions or changes to the feature.

Table 22: Feature History for IS-IS

Feature Name	Release	Feature Information
IS-IS	6.2(2)	Added support for configuring the administrative distance of routes.
IS-IS	6.2(2)	Added the ability to configure all IS-IS interfaces as passive by default and then activate only those interfaces where adjacencies are desired.
IS-IS	6.1(1)	Added support for IPv6.
IS-IS	6.1(1)	Added the no adjacency-check command to disable strict adjacency mode.
IS-IS	5.0(2)	Added support for BFD. See the <i>Cisco Nexus 7000 Series NX-OS Interfaces Configuration Guide</i> , for more information.
Graceful shutdown	4.2(1)	Added support to gracefully shut down an IS-IS instance or IS-IS on an interface but preserve the IS-IS configuration.
Limits on redistributed routes	4.2(1)	Added support for limiting the number of redistributed routes.
Transient mode for hello padding	4.1(2)	Added support to set or unset the hello padding mode.
Attached bit	4.1(2)	Added support to set or unset the attached bit.
IS-IS	4.0(1)	This feature was introduced.



CHAPTER 11

Configuring Basic BGP

This chapter contains the following sections:

- [Finding Feature Information, on page 253](#)
- [Information About Basic BGP, on page 253](#)
- [Prerequisites for BGP, on page 263](#)
- [Guidelines and Limitations for BGP, on page 263](#)
- [Default Settings, on page 264](#)
- [CLI Configuration Modes, on page 265](#)
- [Configuring Basic BGP, on page 266](#)
- [Verifying the Basic BGP Configuration, on page 277](#)
- [Monitoring BGP Statistics, on page 279](#)
- [Configuration Examples for Basic BGP, on page 280](#)
- [Related Documents for Basic BGP, on page 280](#)
- [MIBs, on page 280](#)
- [Feature History for BGP , on page 280](#)

Finding Feature Information

Your software release might not support all the features documented in this module. For the latest caveats and feature information, see the Bug Search Tool at <https://tools.cisco.com/bugsearch/> and the release notes for your software release. To find information about the features documented in this module, and to see a list of the releases in which each feature is supported, see the "New and Changed Information" chapter or the Feature History table in this chapter.

Information About Basic BGP

Cisco NX-OS supports BGP version 4, which includes multiprotocol extensions that allow BGP to carry routing information for IP multicast routes and multiple Layer 3 protocol address families. BGP uses TCP as a reliable transport protocol to create TCP sessions with other BGP-enabled devices.

BGP uses a path-vector routing algorithm to exchange routing information between BGP-enabled networking devices or BGP speakers. Based on this information, each BGP speaker determines a path to reach a particular destination while detecting and avoiding paths with routing loops. The routing information includes the actual route prefix for a destination, the path of autonomous systems to the destination, and additional path attributes.

BGP selects a single path, by default, as the best path to a destination host or network. Each path carries well-known mandatory, well-known discretionary, and optional transitive attributes that are used in BGP best-path analysis. You can influence BGP path selection by altering some of these attributes by configuring BGP policies.

BGP also supports load balancing or equal-cost multipath (ECMP).

For information on configuring BGP in an MPLS network, see the Cisco Nexus 7000 Series NX-OS MPLS Configuration Guide.

BGP Autonomous Systems

An autonomous system (AS) is a network controlled by a single administration entity. An autonomous system forms a routing domain with one or more interior gateway protocols (IGPs) and a consistent set of routing policies. BGP supports 16-bit and 32-bit autonomous system numbers.

Separate BGP autonomous systems dynamically exchange routing information through external BGP (eBGP) peering sessions. BGP speakers within the same autonomous system can exchange routing information through internal BGP (iBGP) peering sessions.

4-Byte AS Number Support

BGP supports 2-byte or 4-byte AS numbers. Cisco NX-OS displays 4-byte AS numbers in plain-text notation (that is, as 32-bit integers). You can configure 4-byte AS numbers as either plain-text notation (for example, 1 to 4294967295) or AS.dot notation (for example, 1.0).

Administrative Distance

An administrative distance is a rating of the trustworthiness of a routing information source. By default, BGP uses the administrative distances shown in the table.

Table 23: BGP Default Administrative Distances

Distance	Default Value	Function
External	20	Applied to routes learned from eBGP.
Internal	200	Applied to routes learned from iBGP.
Local	200	Applied to routes originated by the router.



Note The administrative distance does not influence the BGP path selection algorithm, but it does influence whether BGP-learned routes are installed in the IP routing table.

BGP Peers

A BGP speaker does not discover another BGP speaker automatically. You must configure the relationships between BGP speakers. A BGP peer is a BGP speaker that has an active TCP connection to another BGP speaker.

BGP Sessions

BGP uses TCP port 179 to create a TCP session with a peer. When a TCP connection is established between peers, each BGP peer initially exchanges all of its routes—the complete BGP routing table—with the other peer. After this initial exchange, the BGP peers send only incremental updates when a topology change occurs in the network or when a routing policy change occurs. In the periods of inactivity between these updates, peers exchange special messages called keepalives. The hold time is the maximum time limit that can elapse between receiving consecutive BGP update or keepalive messages.

Cisco NX-OS supports the following peer configuration options:

- Individual IPv4 or IPv4 address—BGP establishes a session with the BGP speaker that matches the remote address and AS number.
- IPv4 or IPv6 prefix peers for a single AS number—BGP establishes sessions with BGP speakers that match the prefix and the AS number.
- Dynamic AS number prefix peers—BGP establishes sessions with BGP speakers that match the prefix and an AS number from a list of configured AS numbers.

Dynamic AS Numbers for Prefix Peers

Cisco NX-OS accepts a range or list of AS numbers to establish BGP sessions. For example, if you configure BGP to use IPv4 prefix 192.0.2.0/8 and AS numbers 33, 66, and 99, BGP establishes a session with 192.0.2.1 with AS number 66 but rejects a session from 192.0.2.2 with AS number 50.

Cisco NX-OS does not associate prefix peers with dynamic AS numbers as either interior BGP (iBGP) or external BGP (eBGP) sessions until after the session is established.



Note The dynamic AS number prefix peer configuration overrides the individual AS number configuration that is inherited from a BGP template.

BGP Router Identifier

To establish BGP sessions between peers, BGP must have a router ID, which is sent to BGP peers in the OPEN message when a BGP session is established. The BGP router ID is a 32-bit value that is often represented by an IPv4 address. You can configure the router ID. By default, Cisco NX-OS sets the router ID to the IPv4 address of a loopback interface on the router. If no loopback interface is configured on the router, the software chooses the highest IPv4 address configured to a physical interface on the router to represent the BGP router ID. The BGP router ID must be unique to the BGP peers in a network.

If BGP does not have a router ID, it cannot establish any peering sessions with BGP peers.

BGP Path Selection

Although BGP might receive advertisements for the same route from multiple sources, BGP selects only one path as the best path. BGP puts the selected path in the IP routing table and propagates the path to its peers.



Note Beginning with Cisco NX-OS Release 6.1, BGP supports sending and receiving multiple paths per prefix and advertising such paths.

The best-path algorithm runs each time that a path is added or withdrawn for a given network. The best-path algorithm also runs if you change the BGP configuration. BGP selects the best path from the set of valid paths available for a given network.

Beginning with Cisco NX-OS Release 8.4(1), the behavior of the BGP pre-best path point of insertion (POI) is changed. In this release, the NX-OS RPM, BGP, and HMM software uses a single cost community ID (either 128 for internal routes or 129 for external routes) to identify a BGP VPNv4 route as an EIGRP originated route.

Only the routes that have the pre-best path value set to cost community ID 128 or 129 are installed in the URIB along with the cost extcommunity. Any non-eigrp originated route carrying the above described cost community ID would be installed in URIB along with pre-best path cost community. As a result, URIB would use this cost to identify the better route between the route learnt through the iBGP and backdoor-EIGRP instead of the administrative distance.

Cisco NX-OS implements the BGP best-path algorithm in the following steps:

1. Compares two paths to determine which is better.
2. Explores all paths and determines in which order to compare the paths to select the overall best path.
3. Determines whether the old and new best paths differ enough so that the new best path should be used.



Note The order of comparison determined in Part 2 is important. Consider the case where you have three paths, A, B, and C. When Cisco NX-OS compares A and B, it chooses A. When Cisco NX-OS compares B and C, it chooses B. But when Cisco NX-OS compares A and C, it might not choose A because some BGP metrics apply only among paths from the same neighboring autonomous system and not among all paths.

The path selection uses the BGP AS-path attribute. The AS-path attribute includes the list of autonomous system numbers (AS numbers) traversed in the advertised path. If you subdivide your BGP autonomous system into a collection or confederation of autonomous systems, the AS-path contains confederation segments that list these locally defined autonomous systems.

BGP Path Selection - Comparing Pairs of Paths

This first step in the BGP best-path algorithm compares two paths to determine which path is better. The following sequence describes the basic steps that Cisco NX-OS uses to compare two paths to determine the better path:

1. Cisco NX-OS chooses a valid path for comparison. (For example, a path that has an unreachable next hop is not valid.)
2. Cisco NX-OS chooses the path with the highest weight.

3. Cisco NX-OS chooses the path with the highest local preference.
4. If one of the paths is locally originated, Cisco NX-OS chooses that path.
5. Cisco NX-OS chooses the path with the shorter AS path.



Note When calculating the length of the AS-path, Cisco NX-OS ignores confederation segments and counts AS sets as 1.

6. Cisco NX-OS chooses the path with the lower origin. Interior Gateway Protocol (IGP) is considered lower than EGP.
7. Cisco NX-OS chooses the path with the lower multiexit discriminator (MED).

You can configure a number of options that affect whether or not this step is performed. In general, Cisco NX-OS compares the MED of both paths if the paths were received from peers in the same autonomous system; otherwise, Cisco NX-OS skips the MED comparison.

You can configure Cisco NX-OS to always perform the best-path algorithm MED comparison, regardless of the peer autonomous system in the paths. Otherwise, Cisco NX-OS performs a MED comparison that depends on the AS-path attributes of the two paths being compared:

- a. If a path has no AS-path or the AS-path starts with an AS_SET, the path is internal and Cisco NX-OS compares the MED to other internal paths.
- b. If the AS-path starts with an AS_SEQUENCE, the peer autonomous system is the first AS number in the sequence and Cisco NX-OS compares the MED to other paths that have the same peer autonomous system.
- c. If the AS-path contains only confederation segments or starts with confederation segments followed by an AS_SET, the path is internal and Cisco NX-OS compares the MED to other internal paths.
- d. If the AS-path starts with confederation segments that are followed by an AS_SEQUENCE, the peer autonomous system is the first AS number in the AS_SEQUENCE and Cisco NX-OS compares the MED to other paths that have the same peer autonomous system.



Note If Cisco NX-OS receives no MED attribute with the path, Cisco NX-OS considers the MED to be 0 unless you configure the best-path algorithm to set a missing MED to the highest possible value.

- e. If the non-deterministic MED comparison feature is enabled, the best-path algorithm uses the Cisco IOS style of MED comparison.
8. If one path is from an internal peer and the other path is from an external peer, Cisco NX-OS chooses the path from the external peer.
9. If the paths have different IGP metrics to their next-hop addresses, Cisco NX-OS chooses the path with the lower IGP metric.
10. Cisco NX-OS uses the path that was selected by the best-path algorithm the last time it was run.
11. If all path parameters in Step 1 through Step 9 are the same, and there is no current best path (for example, the current best path can be lost when the neighbor that offers the current best path goes down), then

the route from the BGP router with the lowest router ID is chosen. If the path includes an originator attribute, Cisco NX-OS uses that attribute as the router ID to compare to; otherwise, Cisco NX-OS uses the router ID of the peer that sent the path. If the paths have different router IDs, Cisco NX-OS chooses the path with the lower router ID.



Note When using the attribute originator as the router ID, it is possible that two paths have the same router ID. It is also possible to have two BGP sessions with the same peer router, so you could receive two paths with the same router ID.

12. Cisco NX-OS selects the path with the shorter cluster length. If a path was not received with a cluster list attribute, the cluster length is 0.
13. Cisco NX-OS chooses the path received from the peer with the lower IP address. Locally generated paths (for example, redistributed paths) have a peer IP address of 0.



Note Paths that are equal after Step 9 can be used for multipath if you configure multipath.

BGP Path Selection - Determining the Order of Comparisons

The second step of the BGP best-path algorithm implementation is to determine the order in which Cisco NX-OS compares the paths:

1. Cisco NX-OS partitions the paths into groups. Within each group, Cisco NX-OS compares the MED among all paths. Cisco NX-OS uses the same rule as in the section *Step 1—Comparing Pairs of Paths* to determine whether MED can be compared between any two paths. Typically, this comparison results in one group being chosen for each neighbor autonomous system. If you configure the **bgp bestpath med always** command, Cisco NX-OS chooses just one group that contains all the paths.
2. Cisco NX-OS determines the best path in each group by iterating through all paths in the group and keeping track of the best one so far. Cisco NX-OS compares each path with the temporary best path found so far and if the new path is better, it becomes the new temporary best path and Cisco NX-OS compares it with the next path in the group.
3. Cisco NX-OS forms a set of paths that contain the best path selected from each group in Step 2. Cisco NX-OS selects the overall best path from this set of paths by going through them as in Step 2.

BGP Path Selection - Determining the Best-Path Change Suppression

The next part of the implementation is to determine whether Cisco NX-OS uses the new best path or suppresses the new best path. The router can continue to use the existing best path if the new one is identical to the old path (if the router ID is the same). Cisco NX-OS continues to use the existing best path to avoid route changes in the network.

You can turn off the suppression feature by configuring the best-path algorithm to compare the router IDs. See the “Tuning the Best-Path Algorithm” section on page 11-10 for more information. If you configure this feature, the new best path is always preferred to the existing one.

You cannot suppress the best-path change if any of the following conditions occur:

- The existing best path is no longer valid.
- Either the existing or new best paths were received from internal (or confederation) peers or were locally generated (for example, by redistribution).
- The paths were received from the same peer (the paths have the same router ID).
- The paths have different weights, local preferences, origins, or IGP metrics to their next-hop addresses.
- The paths have different MEDs.

BGP and the Unicast RIB

BGP communicates with the unicast routing information base (unicast RIB) to store IPv4 routes in the unicast routing table. After selecting the best path, if BGP determines that the best path change needs to be reflected in the routing table, it sends a route update to the unicast RIB.

BGP receives route notifications regarding changes to its routes in the unicast RIB. It also receives route notifications about other protocol routes to support redistribution.

BGP also receives notifications from the unicast RIB regarding next-hop changes. BGP uses these notifications to keep track of the reachability and IGP metric to the next-hop addresses.

Whenever the next-hop reachability or IGP metrics in the unicast RIB change, BGP triggers a best-path recalculation for affected routes.

BGP communicates with the IPv6 unicast RIB to perform these operations for IPv6 routes.

BGP Prefix Independent Convergence

The BGP Prefix Independent Convergence (PIC) feature achieves subsecond convergence in the forwarding plane for BGP IP and Layer 3 VPN routes, when there are BGP next-hop network reachability failures.

BGP PIC has two categories:

- PIC core
- PIC edge

PIC core ensures fast convergence for BGP routes when there is a link or node failure in the core that causes a change in the IGP reachability to a remote BGP next-hop address.

PIC edge ensures fast convergence to a BGP backup path when an external (eBGP) edge link or an external neighbor node fails.

BGP PIC Feature Support Matrix

BGP PIC feature support matrix is shown in the table below:

BGP PIC	IPv4 Unicast	IPv6 Unicast	VPNv4 (per prefix)	VPNv6 (per prefix)	VPNv4 (per VRF)	VPNv6 (per VRF)
Core Unipath	Yes	Yes	No	No	Yes	No
Edge Unipath	Yes	Yes	No	No	No	No

BGP PIC	IPv4 Unicast	IPv6 Unicast	VPNv4 (per prefix)	VPNv6 (per prefix)	VPNv4 (per VRF)	VPNv6 (per VRF)
Core with Multipath equal	Yes	Yes	No	No	Yes	No
Edge Multipath equal (multiple active ECMP, only one backup)	Yes	Yes	No	No	No	No

BGP PIC Core

The BGP PIC core feature is supported by Cisco NX-OS Release 5.2 and later. The feature allows for faster convergence for traffic destined to BGP prefixes that share the same remote next hop in case of a failure in the core of the network. Both MPLS and pure IP traffic can benefit from this feature. It is enabled by default and cannot be disabled.

IPv4, VPNv4, 6PE, and VPNv6 (6VPE) support PIC core with the following constraints:

- For both IP and MPLS core, convergence for internet routes is prefix-independent on the order of BGP next hops.
- With per-VRF label allocation, VPN route convergence is also prefix-independent on the order of BGP next hops. That is, when a path to a remote PE changes, the number of VRFs on that PE determines convergence.
- With per-prefix label allocation, route convergence is not prefix-independent. Convergence moves to the order of VPN routes that are advertised by a remote PE if a failure or change occurs in the reachability to that PE.

For additional considerations when using BGP PIC core in MPLS networks, see the *Cisco Nexus 7000 Series NX-OS MPLS Configuration Guide*.

BGP PIC Edge

The BGP PIC for Edge feature improves BGP convergence after a network failure. This convergence is applicable to edge failures in an IP network. The BGP PIC Edge feature creates and stores a backup path in the routing information base (RIB) and forwarding information base (FIB) so that when a failure on an eBGP link to SP is detected (the primary path fails), the backup path can immediately take over, enabling fast fail over in the forwarding plane.



Note From Cisco NX-OS Release 7.3(0)D1(1) onwards BGP PIC Edge feature supports both IPv4 and IPv6 address families.

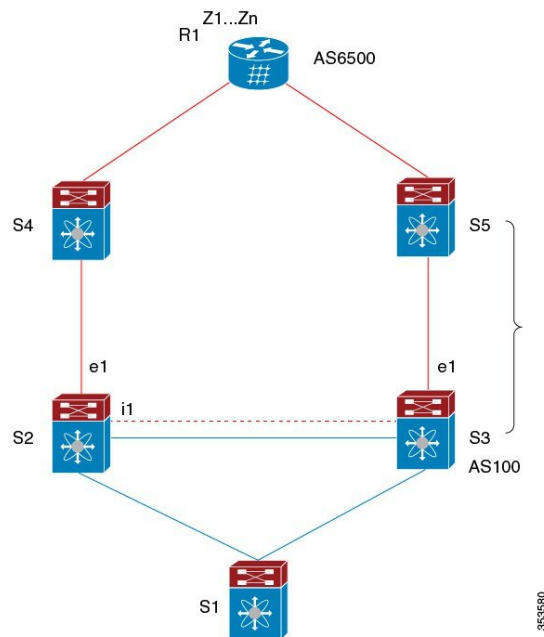
If BGP PIC edge is configured, BGP calculates an additional second best-path (the backup path) along with the primary best-path. BGP installs both best and backup paths for the prefixes with PIC support into the BGP RIB. BGP also downloads the backup path along with the RNH via APIs to the URIB, which then updates the FIB with the next hop marked as a backup. The backup path provides a fast reroute mechanism to counter a singular network failure.

This feature detects both the local interface failure and remote interface/link failure and triggers the use of the backup path.

BGP PIC Edge Unipath

A BGP PIC edge unipath topology is shown in the figure below:

Figure 32: BGP PIC Edge Unipath



In the above figure:

- eBGP sessions are between S2-S4 and S3-S5
- iBGP session is between S2-S3
- Traffic from S1 uses S2 and uses the e1 interface to reach prefixes Z1..Zn.
- S2 has two paths to reach Z1...Zn
 - Primary path via S4
 - Backup/alternate via S5

In this example, S3 advertises to S2 the prefixes Z1...Zn to reach with itself as the next hop. BGP on S2, with BGP PIC feature enabled, installs both bestpath (via S4) and backup path (via S3/S5) towards the AS6500 into the RIB and then the RIB downloads both routes to the FIB.


When the S2-S4 link goes down, the FIB on S2 detects the link failure. It automatically switches from the primary path to the backup/alternate and points to the new next hop S3. Traffic is quickly rerouted due to the local fast re-convergence in FIB. After learning the link failure event, BGP on S2 recomputes the bestpath (which is the previous backup path), removing the next hop S4 from RIB and reinstalling S3 as the primary next hop into RIB. It also computes a new backup/alternate path, if any, and notifies RIB. With the support of the BGP PIC feature, the FIB can switch to the available backup route instantly upon detection of link

failure on the primary route without waiting for BGP to select new bestpath and converge, and achieve a fast reroute.

BGP PIC Edge with Multipaths

In the presence of Equal Cost Multipath (ECMP), none of the multipaths can be selected as the backup path when BGP PIC Edge support is enabled.

Figure 33: BGP PIC Edge with Multipaths



In the above topology, there are six paths for a given prefix as follows:

- eBGP paths: e1, e2, e3
- iBGP paths: i1, i2, i3

The order of preference is $e1 > e2 > e3 > i1 > i2 > i3$.

The potential multipath situations are:

No multipaths configured

- bestpath = e1
- multipath-set = []
- backup path = e2
- PIC behavior: When e1 fails, e2 is activated.

Two-way eBGP multipaths configured:

- bestpath = e1

- multipath-set = [e1, e2]
- backup path = e3
- PIC behavior: Active multipaths are mutually backed up. When all multipaths fail, e3 is activated.

Three-way eBGP multipaths configured:

- bestpath = e1
- multipath-set = [e1, e2, e3]
- backup path = i1
- PIC behavior: Active multipaths are mutually backed up. When all multipaths fail, i1 is activated.

Four-way eiBGP multipaths configured:

- bestpath = e1
- multipath-set = [e1, e2, e3, i1]
- backup path = i2
- PIC behavior: Active multipaths are mutually backed up. When all multipaths fail, i2 is activated.

BGP Virtualization

BGP supports virtual routing and forwarding (VRF) instances. VRFs exist within virtual device contexts (VDCs). By default, Cisco NX-OS places you in the default VDC and default VRF unless you specifically configure another VDC and VRF. For more information, see the *Cisco Nexus 7000 Series NX-OS Virtual Device Context Configuration Guide*.

Prerequisites for BGP

BGP has the following prerequisites:

- You must enable BGP.
- You should have a valid router ID configured on the system.
- You must have an AS number, either assigned by a Regional Internet Registry (RIR) or locally administered.
- You must configure at least one IGP that is capable of recursive next-hop resolution.
- You must configure an address family under a neighbor for the BGP session establishment.

Guidelines and Limitations for BGP

BGP has the following configuration guidelines and limitations:

- Cisco NX-OS does not support "fast-external-falover" for the multi-hop eBGP peering. The BGP differentiates the single-hop (directly connected) and the multi-hop eBGP neighbors using the **ebgp-multihop** command. When you use the **ebgp-multihop 2** command for an eBGP peer, the BGP treats it as multi-hop session and does not trigger the "fast-external-falover". This is a known behaviour.
- The dynamic AS number prefix peer configuration overrides the individual AS number configuration inherited from a BGP template.
- If you configure a dynamic AS number for prefix peers in an AS confederation, BGP establishes sessions with only the AS numbers in the local confederation.
- BGP sessions created through a dynamic AS number prefix peer ignore any configured eBGP multihop time-to-live (TTL) value or a disabled check for directly connected peers.
- Configure a router ID for BGP to avoid automatic router ID changes and session flaps.
- Use the maximum-prefix configuration option per peer to restrict the number of routes received and system resources used.
- Configure the update source to establish a session with BGP/eBGP multihop sessions.
- Specify a BGP policy if you configure redistribution.
- Define the BGP router ID within a VRF.
- If you decrease the keepalive and hold timer values, you might experience BGP session flaps.
- You can configure a minimum route advertisement interval (MRAI) between the sending of BGP routing updates by using the **advertisement-interval** command.
- The BGP Prefix-Independent Convergence (PIC) Edge feature only supports IPv4 address family.
- Only one repair path (backup path) is supported with the BGP PIC Edge feature.

Default Settings

Table 24: Default BGP Parameters

Parameters	Default
BGP feature	Disabled
Keep alive interval	60 seconds
Hold timer	180 seconds
BGP PIC core	Enabled
BGP PIC edge	Disabled
Auto-summary	Always disabled
Synchronization	Always disabled

CLI Configuration Modes

The following sections describe how to enter each of the CLI configuration modes for BGP. From a mode, you can enter the ? command to display the commands available in that mode.

Global Configuration Mode

Use global configuration mode to create a BGP process and configure advanced features such as AS confederation and route dampening.

This example shows how to enter router configuration mode:

```
switch# configuration
switch(config)# router bgp 64496
switch(config-router)#
```

BGP supports Virtual Routing and Forwarding (VRF). You can configure BGP within the appropriate VRF if you are using VRFs in your network.

This example shows how to enter VRF configuration mode:

```
switch(config)# router bgp 64497
switch(config-router)# vrf vrf_A
switch(config-router-vrf)#
```

Address Family Configuration Mode

You can optionally configure the address families that BGP supports. Use the **address-family** command in router configuration mode to configure features for an address family. Use the **address-family** command in neighbor configuration mode to configure the specific address family for the neighbor.

You must configure the address families if you are using route redistribution, address aggregation, load balancing, and other advanced features.

The following example shows how to enter address family configuration mode from the router configuration mode:

```
switch(config)# router bgp 64496
switch(config-router)# address-family ipv6 unicast
switch(config-router-af)#
```

The following example shows how to enter VRF address family configuration mode if you are using VRFs:

```
switch(config)# router bgp 64497
switch(config-router)# vrf vrf_A
switch(config-router-vrf)# address-family ipv6 unicast
switch(config-router-vrf-af)#
```

Neighbor Configuration Mode

Cisco NX-OS provides the neighbor configuration mode to configure BGP peers. You can use neighbor configuration mode to configure all parameters for a peer.

The following example shows how to enter neighbor configuration mode:

```
switch(config)# router bgp 64496
switch(config-router)# neighbor 192.0.2.1
switch(config-router-neighbor)#
```

The following example shows how to enter VRF neighbor configuration mode:

```
switch(config)# router bgp 64497
switch(config-router)# vrf vrf_A
switch(config-router-vrf)# neighbor 192.0.2.1
switch(config-router-vrf-neighbor)#
```

Neighbor Address Family Configuration Mode

An address family configuration submode inside the neighbor configuration submode is available for entering address family-specific neighbor configuration and enabling the address family for the neighbor. Use this mode for advanced features such as limiting the number of prefixes allowed for this neighbor and removing private AS numbers for eBGP.

The following example shows how to enter neighbor address family configuration mode:

```
switch(config)# router bgp 64496
switch(config-router)# neighbor 192.0.2.1
switch(config-router-neighbor)# address-family ipv4 unicast
switch(config-router-neighbor-af)#
```

This example shows how to enter VRF neighbor address family configuration mode:

```
switch(config)# router bgp 64497
switch(config-router)# vrf vrf_A
switch(config-router-vrf)# neighbor 209.165.201.1
switch(config-router-vrf-neighbor)# address-family ipv4 unicast
switch(config-router-vrf-neighbor-af)#
```

Configuring Basic BGP

To configure a basic BGP, you must enable BGP and configure a BGP peer. Configuring a basic BGP network consists of a few required tasks and many optional tasks. You must configure a BGP routing process and BGP peers.



Note If you are familiar with the Cisco IOS CLI, be aware that the Cisco NX-OS commands for this feature might differ from the Cisco IOS commands that you would use.

Enabling BGP

You must enable BGP before you can configure BGP.

Before you begin

Ensure that you are in the correct VDC (or use the **switchto vdc** command).

Procedure

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# feature bgp	Enables BGP. Use the no feature bgp command to disable BGP and remove all associated configuration.
Step 3	(Optional) switch(config)# show feature	(Optional) Displays enabled and disabled features.
Step 4	(Optional) switch(config)# copy running-config startup-config	Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

Creating a BGP Instance

You can create a BGP instance and assign a router ID to the BGP instance. Cisco NX-OS supports 2-byte or 4-byte autonomous system (AS) numbers in plain-text notation or as.dot notation.

Before you begin

- You must enable BGP.
- BGP must be able to obtain a router ID (for example, a configured loopback address).
- Ensure that you are in the correct VDC (or use the **switchto vdc** command).

Procedure

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# router bgp <i>autonomous-system-number</i>	Enables BGP and assigns the AS number to the local BGP speaker. The AS number can be a 16-bit integer or a 32-bit integer in the form of a higher 16-bit decimal number and a lower 16-bit decimal number in xx.xx format. Use the no form of this command to disable this feature.
Step 3	switch(config-router)# router-id <i>ip-address</i>	(Optional) Configures the BGP router ID. This IP address identifies this BGP speaker.
Step 4	switch(config-router)# address-family { ipv4 ipv6 vpn4 vpn6 } { unicast multicast }	Enters global address family configuration mode for the IP or VPN address family.

	Command or Action	Purpose
Step 5	switch(config-router-af)# network ip-prefix [route-map map-name]	(Optional) Specifies a network as local to this autonomous system and adds it to the BGP routing table. For exterior protocols, the network command controls which networks are advertised. Interior protocols use the network command to determine where to send updates.
Step 6	switch(config-router-af)# show bgp all	(Optional) Displays information about all BGP address families.
Step 7	(Optional) switch(config)# copy running-config startup-config	Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

Example

This example shows how to enable BGP with the IPv4 unicast address family and manually add one network to advertise:

```
switch# configure terminal
switch(config)# router bgp 64496
switch(config-router)# address-family ipv4 unicast
switch(config-router-af)# network 192.0.2.0
switch(config-router-af)# copy running-config startup-config
```

Restarting a BGP Instance

You can restart a BGP instance and clear all peer sessions for the instance.

To restart a BGP instance and remove all associated peers, use the following command:

Procedure

	Command or Action	Purpose
Step 1	switch(config)# restart bgp instance-tag	Restarts the BGP instance and resets or reestablishes all peering sessions.

Shutting Down BGP

You can shut down the BGP protocol and gracefully disable BGP and retain the configuration.

To shut down BGP, use the following command in router configuration mode:

Procedure

	Command or Action	Purpose
Step 1	switch(config-router)# shutdown	Gracefully shuts down BGP.

Configuring BGP Peers

You can configure a BGP peer within a BGP process. Each BGP peer has an associated keepalive timer and hold timers. You can set these timers either globally or for each BGP peer. A peer configuration overrides a global configuration.



Note You must configure the address family under neighbor configuration mode for each peer.

Before you begin

- You must enable BGP.
- Ensure that you are in the correct VDC (or use the **switchto vdc** command).

Procedure

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# router bgp <i>autonomous-system-number</i>	Enables BGP and assigns the AS number to the local BGP speaker. The AS number can be a 16-bit integer or a 32-bit integer in the form of a higher 16-bit decimal number and a lower 16-bit decimal number in xx.xx format.
Step 3	switch(config-router)# neighbor { <i>ip-address</i> <i>ipv6-address</i> } remote-as <i>as-number</i>	Configures the IPv4 or IPv6 address and AS number for a remote BGP peer. The ip-address format is x.x.x.x. The ipv6-address format is A:B::C:D.
Step 4	switch(config-router-neighbor)# description <i>text</i>	(Optional) Adds a description for the neighbor. The description is an alphanumeric string up to 80 characters.
Step 5	switch(config-router-neighbor)# timers <i>keepalive-time hold-time</i>	(Optional) Adds the keepalive and hold time BGP timer values for the neighbor. The range is from 0 to 3600 seconds. The default is 60 seconds for the keepalive time and 180 seconds for the hold time.
Step 6	switch(config-router-neighbor)# shutdown	(Optional). Administratively shuts down this BGP neighbor. This command triggers an

	Command or Action	Purpose
		automatic notification and session reset for the BGP neighbor sessions.
Step 7	switch(config-router-neighbor)# address-family { ipv4 ipv6 vpn4 vpn6 } { unicast multicast }	Enters neighbor address family configuration mode for the unicast IPv4 address family.
Step 8	switch(config-router-neighbor)# weight <i>value</i>	(Optional) Sets the default weight for routes from this neighbor. The range is from 0 to 65535. All routes learned from this neighbor have the assigned weight initially. The route with the highest weight is chosen as the preferred route when multiple routes are available to a particular network. The weights assigned with the set weight route-map command override the weights assigned with this command. If you specify a BGP peer policy template, all the members of the template inherit the characteristics configured with this command.
Step 9	(Optional) switch(config-router-neighbor)# show bgp { ipv4 ipv6 vpn4 vpn6 } { unicast multicast } neighbors	(Optional) Displays information about BGP peers.
Step 10	(Optional) switch(config)# copy running-config startup-config	Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

Example

The following example shows how to configure a BGP peer:

```
switch# configure terminal
switch(config)# router bgp 64496
switch(config-router)# neighbor 192.0.2.1 remote-as 64497
switch(config-router-neighbor)# description Peer Router B
switch(config-router-neighbor)# address-family ipv4 unicast
switch(config-router-neighbor)# weight 100
switch(config-router-neighbor-af)# copy running-config startup-config
```

Configuring AS-4 Dot Notation

You can configure 4-byte autonomous system (AS) numbers in asdot notation. The default value is asplain.

Before you begin

Ensure that you are in the correct VDC (or use the **switchto vdc** command).

Procedure

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# as-format asdot	Configures the ASN notation to asdot.
Step 3	(Optional) switch(config)# copy running-config startup-config	Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

Example

The following example configures AS numbers in asdot notation.

```
switch # configure terminal
switch (config) # as-format asdot
switch (config) # copy running-config startup-config
```

Configuring Dynamic AS Numbers for Prefix Peers

You can configure multiple BGP peers within a BGP process. You can limit BGP session establishment to a single AS number or multiple AS numbers in a route map.

BGP sessions configured through dynamic AS numbers for prefix peers ignore the **ebgp-multihop** command and the **disable-connected-check** command.

You can change the list of AS numbers in the route map, but you must use the **no neighbor** command to change the route-map name. Changes to the AS numbers in the configured route map affect only new sessions.

Before you begin

- You must enable BGP.
- Ensure that you are in the correct VDC (or use the **switchto vdc** command).

Procedure

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# router bgp <i>autonomous-system-number</i>	Enables BGP and assigns the AS number to the local BGP speaker. The AS number can be a 16-bit integer or a 32-bit integer in the form of a higher 16-bit decimal number and a lower 16-bit decimal number in xx.xx format.
Step 3	switch(config-router)# neighbor <i>prefix</i> remote-as route-map <i>map-name</i>	Configures the IPv4 or IPv6 prefix and a route map for the list of accepted AS numbers for the remote BGP peers. The prefix format for IPv4

	Command or Action	Purpose
		is x.x.x.x/length. The length range is from 1 to 32. The prefix format for IPv6 is A:B::C:D/length. The length range is from 1 to 128. The <i>map-name</i> can be any case-sensitive, alphanumeric string up to 63 characters.
Step 4	switch(config-router-neighbor-af)# show bgp {ipv4 ipv6 vpnv4 vpnv6} {unicast multicast} neighbors	(Optional) Displays information about BGP peers.
Step 5	(Optional) switch(config)# copy running-config startup-config	Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

Example

This example shows how to configure dynamic AS numbers for a prefix peer:

```
switch# configure terminal
switch(config)# route-map BGPPeers
switch(config-route-map)# match as-number 64496, 64501-64510
switch(config-route-map)# match as-number as-path-list List1, List2
switch(config-route-map)# exit
switch(config)# router bgp 64496
switch(config-router)# neighbor 192.0.2.0/8 remote-as route-map BGPPeers
switch(config-router-neighbor)# description Peer Router B
switch(config-router-neighbor)# address-family ipv4 unicast
switch(config-router-neighbor-af)# copy running-config startup-config
```

Configuring BGP PIC Edge



Note The BGP PIC Edge feature only supports IPv4 address families.

Before you begin

- You must enable BGP.
- Ensure that you are in the correct VDC (or use the **switchto vdc** command).

Procedure

-
- Step 1** Enter configuration mode:
- ```
switch#configure terminal
```
- Step 2** Enable BGP and assign the autonomous system number to the local BGP speaker:

```
switch(config)# router bgp autonomous-system-number
```

**Step 3** Enter router address family configuration mode for the IPv4 unicast address family:

```
switch(config-router)# address-family ipv4 unicast
```

**Step 4** Enable BGP to install the backup path to the routing table:

```
switch(config-router-af)# additional-paths install backup
```

**Step 5** Exit router address family configuration mode:

```
switch(config-router-af)# exit
```

### Example

This example shows how to configure the device to support BGP PIC Edge in IPv4 network:

```
interface Ethernet2/2
 ip address 1.1.1.5/24
 no shutdown

interface Ethernet2/3
 ip address 2.2.2.5/24
 no shutdown

router bgp 100
 address-family ipv4 unicast
 additional-paths install backup
 neighbor 1.1.1.6 remote-as 200
 address-family ipv4 unicast
 neighbor 2.2.2.6 remote-as 100
 address-family ipv4 unicast
```

If BGP receives the same prefix (for example, 99.0.0.0/24) from the two neighbors 1.1.1.6 and 2.2.2.6, both paths will be installed in the URIB—one as the primary path and the other as the backup path.

BGP output:

```
switch(config)# show ip bgp 99.0.0.0/24
BGP routing table information for VRF default, address family IPv4 Unicast
BGP routing table entry for 99.0.0.0/24, version 4
Paths: (2 available, best #2)
Flags: (0x00001a) on xmit-list, is in urib, is best urib route

 Path type: internal, path is valid, not best reason: Internal path, backup path
 AS-Path: 200 , path sourced external to AS
 2.2.2.6 (metric 0) from 2.2.2.6 (2.2.2.6)
 Origin IGP, MED not set, localpref 100, weight 0

 Advertised path-id 1
 Path type: external, path is valid, is best path
 AS-Path: 200 , path sourced external to AS
 1.1.1.6 (metric 0) from 1.1.1.6 (99.0.0.1)
 Origin IGP, MED not set, localpref 100, weight 0

Path-id 1 advertised to peers:
 2.2.2.6
```

## URIB output:

```

URIB output:
switch(config)# show ip route 99.0.0.0/24
IP Route Table for VRF "default"
'*' denotes best ucast next-hop
'***' denotes best mcast next-hop
'[x/y]' denotes [preference/metric]
'%<string>' in via output denotes VRF <string>

99.0.0.0/24, ubest/mbest: 1/0
 *via 1.1.1.6, [20/0], 14:34:51, bgp-100, external, tag 200
 via 2.2.2.6, [200/0], 14:34:51, bgp-100, internal, tag 200 (backup)

```

## UFIB output:

```

switch# show forwarding route 123.1.1.0 detail module 8

Prefix 123.1.1.0/24, No of paths: 1, Update time: Fri Feb 7 19:00:12 2014
 Vobj id: 141 orig_as: 65002 peer_as: 65100 rnh: 10.3.0.2
 10.4.0.2 Ethernet8/4 DMAC: 0018.bad8.4dfd
 packets: 2 bytes: 3484 Repair path 10.3.0.2 Ethernet8/3
 DMAC: 0018.bad8.4dfd
 packets: 0 bytes: 1

```

## Clearing BGP Information

To clear BGP information, use the following commands:

| Command                                                                                          | Purpose                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|--------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <pre>clear bgp all {neighbor   *   as-number   peer-template name   prefix} [vrf vrf-name]</pre> | <p>Clears one or more neighbors from all address families. * clears all neighbors in all address families. The arguments are as follows:</p> <ul style="list-style-type: none"> <li>• <i>neighbor</i>—IPv4 or IPv6 address of a neighbor.</li> <li>• <i>as-number</i>—Autonomous system number. The AS number can be a 16-bit integer or a 32-bit integer in the form of higher 16-bit decimal number and a lower 16-bit decimal number in xx.xx format.</li> <li>• <i>name</i>—Peer template name. The name can be any case-sensitive, alphanumeric string up to 64 characters.</li> <li>• <i>prefix</i>—IPv4 or IPv6 prefix. All neighbors within that prefix are cleared.</li> <li>• <i>vrf-name</i>—VRF name. All neighbors in that VRF are cleared. The name can be any case-sensitive, alphanumeric string up to 64 characters.</li> </ul> |

| Command                                                                                                                                                                                                 | Purpose                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>clear bgp all dampening</b> [ <b>vrf</b> <i>vrf-name</i> ]                                                                                                                                           | Clears route flap dampening networks in all address families. The <i>vrf-name</i> can be any case-sensitive, alphanumeric string up to 64 characters.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>clear bgp all flap-statistics</b> [ <b>vrf</b> <i>vrf-name</i> ]                                                                                                                                     | Clears route flap statistics in all address families. The <i>vrf-name</i> can be any case-sensitive, alphanumeric string up to 64 characters.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>clear bgp</b> { <b>ipv4</b>   <b>ipv6</b>   <b>vpn4</b>   <b>vpn6</b> } { <b>unicast</b>   <b>multicast</b> } <b>dampening</b> [ <b>vrf</b> <i>vrf-name</i> ]                                        | Clears route flap dampening networks in the selected address family. The <i>vrf-name</i> can be any case-sensitive, alphanumeric string up to 64 characters.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>clear bgp</b> { <b>ipv4</b>   <b>ipv6</b>   <b>vpn4</b>   <b>vpn6</b> } { <b>unicast</b>   <b>multicast</b> } <b>flap-statistics</b> [ <b>vrf</b> <i>vrf-name</i> ]                                  | Clears route flap statistics in the selected address family. The <i>vrf-name</i> can be any case-sensitive, alphanumeric string up to 64 characters.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>clear bgp</b> { <b>ipv4</b>   <b>ipv6</b>   <b>vpn4</b>   <b>vpn6</b> } { <i>neighbor</i>   *   <i>as-number</i>   <b>peer-template</b> <i>name</i>   <i>prefix</i> } [ <b>vrf</b> <i>vrf-name</i> ] | Clears one or more neighbors from the selected address family. * clears all neighbors in the address family. The arguments are as follows: <ul style="list-style-type: none"> <li>• <i>neighbor</i>—IPv4 or IPv6 address of a neighbor.</li> <li>• <i>as-number</i>— Autonomous system number. The AS number can be a 16-bit integer or a 32-bit integer in the form of higher 16-bit decimal number and a lower 16-bit decimal number in xx.xx format.</li> <li>• <i>name</i>—Peer template name. The name can be any case-sensitive, alphanumeric string up to 64 characters.</li> <li>• <i>prefix</i>—IPv4 or IPv6 prefix. All neighbors within that prefix are cleared.</li> <li>• <i>vrf-name</i>—VRF name. All neighbors in that VRF are cleared. The name can be any case-sensitive, alphanumeric string up to 64 characters.</li> </ul> |

| Command                                                                                                                                                                                         | Purpose                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>clear bgp</b> { <b>ip</b> { <b>unicast</b>   <b>multicast</b> }} { <i>neighbor</i>   *<br>[ <i>as-number</i>   <b>peer-template name</b>   <i>prefix</i> ] [ <b>vrf</b><br><i>vrf-name</i> ] | Clears one or more neighbors. * clears all neighbors in the address family. The arguments are as follows: <ul style="list-style-type: none"> <li>• <i>neighbor</i>—IPv4 or IPv6 address of a neighbor.</li> <li>• <i>as-number</i>— Autonomous system number. The AS number can be a 16-bit integer or a 32-bit integer in the form of higher 16-bit decimal number and a lower 16-bit decimal number in xx.xx format.</li> <li>• <i>name</i>—Peer template name. The name can be any case-sensitive, alphanumeric string up to 64 characters.</li> <li>• <i>prefix</i>—IPv4 or IPv6 prefix. All neighbors within that prefix are cleared.</li> <li>• <i>vrf-name</i>—VRF name. All neighbors in that VRF are cleared. The name can be any case-sensitive, alphanumeric string up to 64 characters.</li> </ul> |
| <b>clear bgp dampening</b> [ <i>ip-neighbor</i>   <i>ip-prefix</i> ] [ <b>vrf</b><br><i>vrf-name</i> ]                                                                                          | Clears route flap dampening in one or more networks. The arguments are as follows: <ul style="list-style-type: none"> <li>• <i>ip-neighbor</i>—IPv4 address of a neighbor.</li> <li>• <i>ip-prefix</i>—IPv4. All neighbors within that prefix are cleared.</li> <li>• <i>vrf-name</i>—VRF name. All neighbors in that VRF are cleared. The name can be any case-sensitive, alphanumeric string up to 64 characters.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>clear bgp flap-statistics</b> [ <i>ip-neighbor</i>   <i>ip-prefix</i> ] [ <b>vrf</b><br><i>vrf-name</i> ]                                                                                    | Clears route flap statistics in one or more networks. The arguments are as follows: <ul style="list-style-type: none"> <li>• <i>ip-neighbor</i>—IPv4 address of a neighbor.</li> <li>• <i>ip-prefix</i>—IPv4. All neighbors within that prefix are cleared.</li> <li>• <i>vrf-name</i>—VRF name. All neighbors in that VRF are cleared. The name can be any case-sensitive, alphanumeric string up to 64 characters.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                |

| Command                                                                                                               | Purpose                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|-----------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>clear ip mbgp</b> {ip {unicast   multicast}} {neighbor [*  as-number   peer-template name   prefix} [vrf vrf-name] | <ul style="list-style-type: none"> <li>• <i>neighbor</i>—IPv4 or IPv6 address of a neighbor.</li> <li>• <i>as-number</i>— Autonomous system number. The AS number can be a 16-bit integer or a 32-bit integer in the form of higher 16-bit decimal number and a lower 16-bit decimal number in xx.xx format.</li> <li>• <i>name</i>—Peer template name. The name can be any case-sensitive, alphanumeric string up to 64 characters.</li> <li>• <i>prefix</i>—IPv4 or IPv6 prefix. All neighbors within that prefix are cleared.</li> <li>• <i>vrf-name</i>—VRF name. All neighbors in that VRF are cleared. The name can be any case-sensitive, alphanumeric string up to 64 characters.</li> </ul> |
| <b>clear ip mbgp dampening</b> [ip-neighbor  ip-prefix] [vrf vrf-name]                                                | <p>Clears route flap dampening in one or more networks. The arguments are as follows:</p> <ul style="list-style-type: none"> <li>• <i>ip-neighbor</i>—IPv4 address of a neighbor.</li> <li>• <i>ip-prefix</i>—IPv4. All neighbors within that prefix are cleared.</li> <li>• <i>vrf-name</i>—VRF name. All neighbors in that VRF are cleared. The name can be any case-sensitive, alphanumeric string up to 64 characters.</li> </ul>                                                                                                                                                                                                                                                                |
| <b>clear ip mbgp flap-statistics</b> [ip-neighbor  ip-prefix] [vrf vrf-name]                                          | <p>Clears route flap statistics one or more networks. The arguments are as follows:</p> <ul style="list-style-type: none"> <li>• <i>ip-neighbor</i>—IPv4 address of a neighbor.</li> <li>• <i>ip-prefix</i>—IPv4. All neighbors within that prefix are cleared.</li> <li>• <i>vrf-name</i>—VRF name. All neighbors in that VRF are cleared. The name can be any case-sensitive, alphanumeric string up to 64 characters.</li> </ul>                                                                                                                                                                                                                                                                  |

## Verifying the Basic BGP Configuration

To display the BGP configuration, perform one of the following tasks:

| Command                                      | Purpose                                                |
|----------------------------------------------|--------------------------------------------------------|
| <b>show bgp all</b> [summary] [vrf vrf-name] | Displays the BGP information for all address families. |

| Command                                                                                                                                                                                                            | Purpose                                                                                                                                     |
|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------|
| <code>show bgp convergence [vrf vrf-name]</code>                                                                                                                                                                   | Displays the BGP information for all address families.                                                                                      |
| <code>show bgp {ipv4   ipv6   vpnv4   vpnv6} {unicast   multicast} [ip-address   ipv6-prefix community [regexp expression   [community] [no-advertise] [no-export] [no-export-subconfed]]] [vrf vrf-name]</code>   | Displays the BGP routes that match a BGP community.                                                                                         |
| <code>show bgp [vrf vrf-name] {ipv4   ipv6   vpnv4   vpnv6} {unicast   multicast} [ip-address   ipv6-prefix] community-list list-name [vrf vrf-name]</code>                                                        | Displays the BGP routes that match a BGP community list.                                                                                    |
| <code>show bgp {ipv4   ipv6   vpnv4   vpnv6} {unicast   multicast} [ip-address   ipv6-prefix extcommunity [regexp expression   [generic [non-transitive   transitive] aa4:nn [exact-match]]] [vrf vrf-name]</code> | Displays the BGP routes that match a BGP extended community.                                                                                |
| <code>show bgp {ipv4   ipv6   vpnv4   vpnv6} {unicast   multicast} [ip-address   ipv6-prefix extcommunity-list list-name [exact-match]] [vrf vrf-name]</code>                                                      | Displays the BGP routes that match a BGP extended community list.                                                                           |
| <code>show bgp {ipv4   ipv6   vpnv4   vpnv6} {unicast   multicast} [ip-address   ipv6-prefix {dampening dampened-paths [regexp expression]}] [vrf vrf-name]</code>                                                 | Displays the information for BGP route dampening. Use the <b>clear bgp dampening</b> command to clear the route flap dampening information. |
| <code>show bgp {ipv4   ipv6   vpnv4   vpnv6} {unicast   multicast} [ip-address   ipv6-prefix history-paths [regexp expression] [vrf vrf-name]</code>                                                               | Displays the BGP route history paths.                                                                                                       |
| <code>show bgp {ipv4   ipv6   vpnv4   vpnv6} {unicast   multicast} [ip-address   ipv6-prefix filter-list list-name [vrf vrf-name]</code>                                                                           | Displays the information for the BGP filter list.                                                                                           |
| <code>show bgp {ipv4   ipv6   vpnv4   vpnv6} {unicast   multicast} [ip-address   ipv6-prefix] neighbors [ip-address   ipv6-prefix] [vrf vrf-name]</code>                                                           | Displays the information for BGP peers. Use the <b>clear bgp neighbors</b> command to clear these neighbors.                                |
| <code>show bgp {ipv4   ipv6   vpnv4   vpnv6} {unicast   multicast} [ip-address   ipv6-prefix] neighbors [ip-address   ipv6-prefix] {nexthop   nexthop-database} [vrf vrf-name]</code>                              | Displays the information for the BGP route next hop.                                                                                        |
| <code>show bgp paths</code>                                                                                                                                                                                        | Displays the BGP path information.                                                                                                          |
| <code>show bgp {ipv4   ipv6   vpnv4   vpnv6} {unicast   multicast} [ip-address   ipv6-prefix] policy name [vrf vrf-name]</code>                                                                                    | Displays the BGP policy information. Use the <b>clear bgp policy</b> command to clear the policy information.                               |
| <code>show bgp {ipv4   ipv6   vpnv4   vpnv6} {unicast   multicast} [ip-address   ipv6-prefix] prefix-list list-name [vrf vrf-name]</code>                                                                          | Displays the BGP routes that match the prefix list.                                                                                         |



| Command                                                                                                                                                                                                                    | Purpose                                                                                                                                                                                          |
|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>show bgp</b> { <b>ipv4</b>   <b>ipv6</b>   <b>vpn4</b>   <b>vpn6</b> } { <b>unicast</b>   <b>multicast</b> } [ <i>ip-address</i>   <i>ipv6-prefix</i> ] <b>received-paths</b> [ <b>vrf</b> <i>vrf-name</i> ]            | Displays the BGP paths stored for soft reconfiguration.                                                                                                                                          |
| <b>show bgp</b> { <b>ipv4</b>   <b>ipv6</b>   <b>vpn4</b>   <b>vpn6</b> } { <b>unicast</b>   <b>multicast</b> } [ <i>ip-address</i>   <i>ipv6-prefix</i> ] <b>regex</b> <i>expression</i> [ <b>vrf</b> <i>vrf-name</i> ]   | Displays the BGP routes that match the AS_path regular expression.                                                                                                                               |
| <b>show bgp</b> { <b>ipv4</b>   <b>ipv6</b>   <b>vpn4</b>   <b>vpn6</b> } { <b>unicast</b>   <b>multicast</b> } [ <i>ip-address</i>   <i>ipv6-prefix</i> ] <b>route-map</b> <i>map-name</i> [ <b>vrf</b> <i>vrf-name</i> ] | Displays the BGP routes that match the route map.                                                                                                                                                |
| <b>show bgp peer-policy</b> <i>name</i> [ <b>vrf</b> <i>vrf-name</i> ]                                                                                                                                                     | Displays the information about BGP peer policies.                                                                                                                                                |
| <b>show bgp peer-session</b> <i>name</i> [ <b>vrf</b> <i>vrf-name</i> ]                                                                                                                                                    | Displays the information about BGP peer sessions.                                                                                                                                                |
| <b>show bgp peer-template</b> <i>name</i> [ <b>vrf</b> <i>vrf-name</i> ]                                                                                                                                                   | Displays the information about BGP peer templates. Use the <b>clear bgp peer-template</b> command to clear all neighbors in a peer template.                                                     |
| <b>show bgp process</b>                                                                                                                                                                                                    | Displays the BGP process information.                                                                                                                                                            |
| <b>show</b> { <b>ipv</b>   <b>ipv6</b> } <b>bgp options</b>                                                                                                                                                                | Displays the BGP status and configuration information. This command has multiple options. See the <i>Cisco Nexus 7000 Series NX-OS Unicast Routing Command Reference</i> , for more information. |
| <b>show</b> { <b>ipv</b>   <b>ipv6</b> } <b>mbgp options</b>                                                                                                                                                               | Displays the BGP status and configuration information. This command has multiple options. See the <i>Cisco Nexus 7000 Series NX-OS Unicast Routing Command Reference</i> , for more information. |
| <b>show running-configuration</b> <b>bgp</b>                                                                                                                                                                               | Displays the current running BGP configuration.                                                                                                                                                  |

## Monitoring BGP Statistics

To display BGP statistics, use the following commands:

| Command                                                                                                                                                                                                          | Purpose                                                                                                             |
|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------|
| <b>show bgp</b> { <b>ipv4</b>   <b>ipv6</b>   <b>vpn4</b>   <b>vpn6</b> } { <b>unicast</b>   <b>multicast</b> } [ <i>ip-address</i>   <i>ipv6-prefix</i> ] <b>flap-statistics</b> [ <b>vrf</b> <i>vrf-name</i> ] | Displays the BGP route flap statistics. Use the <b>clear bgp flap-statistics</b> command to clear these statistics. |
| <b>show bgp sessions</b> [ <b>vrf</b> <i>vrf-name</i> ]                                                                                                                                                          | Displays the BGP sessions for all peers. Use the <b>clear bgp sessions</b> command to clear these statistics.       |
| <b>show bgp statistics</b>                                                                                                                                                                                       | Displays the BGP statistics.                                                                                        |

## Configuration Examples for Basic BGP

This example shows a basic BGP configuration:

```
switch (config) # feature bgp
switch (config) # router bgp 64496
switch (config-router) # neighbor 2001:ODB8:0:1::55 remote-as 64496
switch (config-router) # address-family ipv6 unicast
switch (config-router-af) # next-hop-self
```

This example shows a basic BGP configuration:

```
switch (config) # address-family
switch (config) # router bgp 64496
switch (config-router) # address-family ipv4 unicast
switch (config-router) # network 1.1.10 mask 255.255.255.0
switch (config-router) # neighbor 10.1.1.1 remote-as 64496
switch (config-router) # address-family ipv4 unicast
```

## Related Documents for Basic BGP

| Related Topics     | Document Title                                                                               |
|--------------------|----------------------------------------------------------------------------------------------|
| BGP CLI commands   | <i>Cisco Nexus 7000 Series NX-OS Unicast Routing Command Reference</i>                       |
| MPLS configuration | <i>Cisco Nexus 7000 Series NX-OS MPLS Configuration Guide</i>                                |
| VDCs and VRFs      | <i>Cisco Nexus 7000 Series NX-OS Virtual Device Context Configuration Guide, Release 5.x</i> |

## MIBs

| MIBs                                          | MIBs Link                                                                                                                     |
|-----------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------|
| BGP4-MIB<br>CISCO-BGP4-MIB<br>CISCO-BGP-MIBv2 | To locate and download MIBs, go to the following URL: <a href="https://cfng.cisco.com/mibs">https://cfng.cisco.com/mibs</a> . |

## Feature History for BGP

This table includes only the updates for those releases that have resulted in additions or changes to the feature.

Table 25: Feature History for BGP

| Feature Name                       | Releases | Feature Information                                                                                                                                      |
|------------------------------------|----------|----------------------------------------------------------------------------------------------------------------------------------------------------------|
| BGP PIC Edge                       | 6.2(8)   | Introduced this feature.                                                                                                                                 |
| BGP                                | 6.2(8)   | Added support for CISCO-BGP-MIBv2                                                                                                                        |
| 4-byte AS number                   | 6.2(2)   | Added the ability to configure 4-byte AS numbers in asdot notation.                                                                                      |
| BGP                                | 6.1(1)   | Added support for additional BGP paths.                                                                                                                  |
| BGP                                | 6.1(1)   | Added the ability to set the default weigh for routes from a neighbor using the <b>weight</b> command in the neighbor address family configuration mode. |
| BGP                                | 5.2(1)   | Added support for the BGP PIC core feature.                                                                                                              |
| VPN address families               | 5.2(1)   | Added support for VPN address families.                                                                                                                  |
| BFD                                | 5.0(2)   | Added support for BFD. See the <i>Cisco Nexus 7000 Series NX-OS Interfaces Configuration Guide, Release 6.x</i> for more information.                    |
| ISSU                               | 4.2(3)   | Lowered BGP minimum hold-time check to eight seconds.                                                                                                    |
| IPv6                               | 4.2(1)   | Added support for IPv6.                                                                                                                                  |
| 4-Byte AS numbers                  | 4.2(1)   | Added support for 4-byte AS numbers in plaintext notation.                                                                                               |
| Conditional advertisement          | 4.2(1)   | Added support for conditionally advertising BGP routes based on the existence of other routes in the BGP table.                                          |
| Dynamic AS number for prefix peers | 4.1(2)   | Added support for a range of AS numbers for BGP prefix peer configuration.                                                                               |
| BGP                                | 4.0(1)   | This feature was introduced.                                                                                                                             |





# CHAPTER 12

## Configuring Advanced BGP

---

This chapter contains the following sections:

- [Finding Feature Information, on page 283](#)
- [Information About Advanced BGP, on page 283](#)
- [Prerequisites for Advanced BGP, on page 295](#)
- [Guidelines and Limitations for Advanced BGP, on page 295](#)
- [Default Settings, on page 296](#)
- [Configuring Advanced BGP, on page 297](#)
- [Verifying the Advanced BGP Configuration, on page 331](#)
- [Displaying Advanced BGP Statistics, on page 333](#)
- [Related Documents, on page 333](#)
- [RFCs, on page 334](#)
- [MIBs, on page 334](#)
- [Feature History for Advanced BGP , on page 334](#)

### Finding Feature Information

Your software release might not support all the features documented in this module. For the latest caveats and feature information, see the Bug Search Tool at <https://tools.cisco.com/bugsearch/> and the release notes for your software release. To find information about the features documented in this module, and to see a list of the releases in which each feature is supported, see the "New and Changed Information" chapter or the Feature History table in this chapter.

### Information About Advanced BGP

BGP is an interdomain routing protocol that provides loop-free routing between organizations or autonomous systems. Cisco NX-OS supports BGP version 4. BGP version 4 includes multiprotocol extensions that allow BGP to carry routing information for IP multicast routes and multiple Layer 3 protocol address families. BGP uses TCP as a reliable transport protocol to create TCP sessions with other BGP-enabled devices called BGP peers. When connecting to an external organization, the router creates external BGP (eBGP) peering sessions. BGP peers within the same organization exchange routing information through internal BGP (iBGP) peering sessions.

## Peer Templates

BGP peer templates allow you to create blocks of common configuration that you can reuse across similar BGP peers. Each block allows you to define a set of attributes that a peer then inherits. You can choose to override some of the inherited attributes as well, making it a very flexible scheme for simplifying the repetitive nature of BGP configurations.

Cisco NX-OS implements three types of peer templates:

- The peer-session template defines BGP peer session attributes, such as the transport details, remote autonomous system number of the peer, and session timers. A peer-session template can also inherit attributes from another peer-session template (with locally defined attributes that override the attributes from an inherited peer-session).
- A peer-policy template defines the address-family dependent policy aspects for a peer including the inbound and outbound policy, filter-lists, and prefix-lists. A peer-policy template can inherit from a set of peer-policy templates. Cisco NX-OS evaluates these peer-policy templates in the order specified by the preference value in the inherit configuration. The lowest number is preferred over higher numbers.
- The peer template can inherit the peer-session and peer-policy templates to allow for simplified peer definitions. It is not mandatory to use a peer template but it can simplify the BGP configuration by providing reusable blocks of configuration.

## Authentication

You can configure authentication for a BGP neighbor session. This authentication method adds an MD5 authentication digest to each TCP segment sent to the neighbor to protect BGP against unauthorized messages and TCP security attacks.



---

**Note** The MD5 password must be identical between BGP peers.

---

## Route Policies and Resetting BGP Sessions

You can associate a route policy to a BGP peer. Route policies use route maps to control or modify the routes that BGP recognizes. You can configure a route policy for inbound or outbound route updates. The route policies can match on different criteria, such as a prefix or AS\_path attribute, and selectively accept or deny the routes. Route policies can also modify the path attributes.

When you change a route policy applied to a BGP peer, you must reset the BGP sessions for that peer. Cisco NX-OS supports the following three mechanisms to reset BGP peering sessions:

- Hard reset—A hard reset tears down the specified peering sessions, including the TCP connection, and deletes routes coming from the specified peer. This option interrupts packet flow through the BGP network. Hard reset is disabled by default.
- Soft reconfiguration inbound—A soft reconfiguration inbound triggers routing updates for the specified peer without resetting the session. You can use this option if you change an inbound route policy. Soft reconfiguration inbound saves a copy of all routes received from the peer before processing the routes through the inbound route policy. If you change the inbound route policy, Cisco NX-OS passes these stored routes through the modified inbound route policy to update the route table without tearing down

existing peering sessions. Soft reconfiguration inbound can use significant memory resources to store the unfiltered BGP routes. Soft reconfiguration inbound is disabled by default.

- **Route Refresh**—A route refresh updates the inbound routing tables dynamically by sending route refresh requests to supporting peers when you change an inbound route policy. The remote BGP peer responds with a new copy of its routes that the local BGP speaker processes with the modified route policy. Cisco NX-OS automatically sends an outbound route refresh of prefixes to the peer.
- BGP peers advertise the route refresh capability as part of the BGP capability negotiation when establishing the BGP peer session. Route refresh is the preferred option and enabled by default.

BGP also uses route maps for route redistribution, route aggregation, route dampening, and other features.

## eBGP

External BGP (eBGP) allows you to connect BGP peers from different autonomous systems to exchange routing updates. Connecting to external networks enables traffic from your network to be forwarded to other networks and across the Internet.

You should use loopback interfaces for establishing eBGP peering sessions because loopback interfaces are less susceptible to interface flapping. An interface flap occurs when the interface is administratively brought up or down because of a failure or maintenance issue.

### BGP Next Hop Unchanged

In an eBGP session, by default, the router changes the next-hop attribute of a BGP route to its own address when the router sends out a route. The BGP next-hop unchanged feature allows BGP to send an update to an eBGP multihop peer with the next-hop attribute unchanged.


By default, BGP puts itself as the next hop when announcing to an eBGP peer. When you enter the **set ip next-hop unchanged** command for an outbound route map that is configured for an eBGP peer, it propagates the received next hop to the eBGP peer.

## iBGP

Internal BGP (iBGP) allows you to connect BGP peers within the same autonomous system. You can use iBGP for multihomed BGP networks (networks that have more than one connection to the same external autonomous system).

The figure shows an iBGP network within a larger BGP network.

**Figure 34: iBGP Network**



iBGP networks are fully meshed. Each iBGP peer has a direct connection to all other iBGP peers to prevent network loops.

For single-hop iBGP peers with update-source configured under neighbor configuration mode, the peer supports fast external fall-over.




**Note** You should configure a separate interior gateway protocol in the iBGP network.

## AS Confederations

A fully meshed iBGP network becomes complex as the number of iBGP peers grows. You can reduce the iBGP mesh by dividing the autonomous system into multiple subautonomous systems and grouping them into a single confederation. A confederation is a group of iBGP peers that use the same autonomous system number to communicate to external networks. Each subautonomous system is fully meshed within itself and has a few connections to other subautonomous systems in the same confederation.

The figure shows the BGP network, split into two subautonomous systems and one confederation.

**Figure 35: AS Confederation**



In this example, AS10 is split into two subautonomous systems, AS1 and AS2. Each subautonomous system is fully meshed, but there is only one link between the subautonomous systems. By using AS confederations, you can reduce the number of links compared to the fully meshed autonomous system.

## Route Reflector

You can alternately reduce the iBGP mesh by using a route reflector configuration where route reflectors pass learned routes to neighbors so that all iBGP peers do not need to be fully meshed.


The figure below shows a simple iBGP configuration with four meshed iBGP speakers (routers A,B,C, and D.) Without these route reflectors, when router A receives a route from an external neighbor, it advertise the route to all three iBGP neighbors.

When you configure an iBGP peer to be a route reflector, it becomes responsible for passing iBGP learned routes to a set of iBGP neighbors.

In the figure, router B is the route reflector. When the route reflector receives routes advertised from router A, it advertises (reflects) the routes to routers C and D. Router A no longer has to advertise to both routers C and D.



Figure 36: Route Reflector



The route reflector and its client peers form a cluster. You do not have to configure all iBGP peers to act as client peers of the route reflector. You must configure any nonclient peer as fully meshed to guarantee that complete BGP updates reach all peers.

## Capabilities Negotiation

A BGP speaker can learn about BGP extensions that are supported by a peer by using the capabilities negotiation feature. Capabilities negotiation allows BGP to use only the set of features supported by both BGP peers on a link.

If a BGP peer does not support capabilities negotiation, Cisco NX-OS attempts a new session to the peer without capabilities negotiation if you have configured the address family as IPv4. Any other multiprotocol configuration (such as IPv6) requires capabilities negotiation.

## Route Dampening

Route dampening is a BGP feature that minimizes the propagation of flapping routes across an internetwork. A route flaps when it alternates between the available and unavailable states in rapid succession.

For example, consider a network with three BGP autonomous systems: AS1, AS2, and AS3. Suppose that a route in AS1 flaps (it becomes unavailable). Without route dampening, AS1 sends a withdraw message to AS2. AS2 propagates the withdrawal message to AS3. When the flapping route reappears, AS1 sends an advertisement message to AS2, which sends the advertisement to AS3. If the route repeatedly becomes unavailable, and then available, AS1 sends many withdrawal and advertisement messages that propagate through the other autonomous systems.

Route dampening can minimize flapping. Suppose that the route flaps. AS2 (in which route dampening is enabled) assigns the route a penalty of 1000. AS2 continues to advertise the status of the route to neighbors. Each time that the route flaps, AS2 adds to the penalty value. When the route flaps so often that the penalty exceeds a configurable suppression limit, AS2 stops advertising the route, regardless of how many times that it flaps. The route is now dampened.

The penalty placed on the route decays until the reuse limit is reached. At that time, AS2 advertises the route again. When the reuse limit is at 50 percent, AS2 removes the dampening information for the route.



**Note** The router does not apply a penalty to a resetting BGP peer when route dampening is enabled, even though the peer reset withdraws the route.

## Load Sharing and Multipath

BGP can install multiple equal-cost eBGP or iBGP paths into the routing table to reach the same destination prefix. Traffic to the destination prefix is then shared across all the installed paths.

The BGP best-path algorithm considers the paths as equal-cost paths if the following attributes are identical:

- Weight
- Local preference
- AS\_path
- Origin code
- Multi-exit discriminator (MED)
- IGP cost to the BGP next hop

In Cisco NX-OS releases prior to 6.1, BGP selects only one of these multiple paths as the best path and advertises the path to the BGP peers. Beginning with Cisco NX-OS Release 6.1, BGP supports sending and receiving multiple paths per prefix and advertising such paths.



---

**Note** Paths that are received from different AS confederations are considered as equal-cost paths if the external AS\_path values and the other attributes are identical.

---



---

**Note** When you configure a route reflector for iBGP multipath, and the route reflector advertises the selected best path to its peers, the next hop for the path is not modified.


---

## BGP Additional Paths

In Cisco NX-OS releases prior to 6.1, only one BGP best path is advertised, and the BGP speaker accepts only one path for a given prefix from a given peer. If a BGP speaker receives multiple paths for the same prefix within the same session, it uses the most recent advertisement.

Beginning with Cisco NX-OS Release 6.1, BGP supports the additional paths feature, which allows the BGP speaker to propagate and accept multiple paths for the same prefix without the new paths replacing any previous ones. This feature allows BGP speaker peers to negotiate whether they support advertising and receiving multiple paths per prefix and advertising such paths. A special 4-byte path ID is added to the network layer reachability information (NLRI) to differentiate multiple paths for the same prefix sent across a peer session. The following figure illustrates the BGP additional paths capability.

Figure 37: BGP Route Advertisement with the Additional Paths Capability



## Route Aggregation

You can configure aggregate addresses. Route aggregation simplifies route tables by replacing a number of more specific addresses with an address that represents all the specific addresses. For example, you can replace these three more specific addresses, 10.1.1.0/24, 10.1.2.0/24, and 10.1.3.0/24 with one aggregate address, 10.1.0.0/16.

Aggregate prefixes are present in the BGP route table so that fewer routes are advertised.



**Note** Cisco NX-OS does not support automatic route aggregation.

Route aggregation can lead to forwarding loops. To avoid this problem, when BGP generates an advertisement for an aggregate address, it automatically installs a summary discard route for that aggregate address in the local routing table. BGP sets the administrative distance of the summary discard to 220 and sets the route type to discard. BGP does not use discard routes for next-hop resolution.

Summary entry is created in the BGP table when **aggregate-address** command is configured, though it will not be eligible for advertisement until a subset of the aggregate is found in the table.

## BGP Conditional Advertisement

BGP conditional advertisement allows you to configure BGP to advertise or withdraw a route based on whether or not a prefix exists in the BGP table. This feature is useful, for example, in multihomed networks, in which you want BGP to advertise some prefixes to one of the providers only if information from the other provider is not present.

Consider an example network with three BGP autonomous systems: AS1, AS2, and AS3, where AS1 and AS3 connect to the Internet and to AS2. Without conditional advertisement, AS2 propagates all routes to both AS1 and AS3. With conditional advertisement, you can configure AS2 to advertise certain routes to AS3 only if routes from AS1 do not exist (if for example, the link to AS1 fails).

BGP conditional advertisement adds an exist or not-exist test to each route that matches the configured route map.

## BGP Next-Hop Address Tracking

BGP monitors the next-hop address of installed routes to verify next-hop reachability and to select, install, and validate the BGP best path. BGP next-hop address tracking speeds up this next-hop reachability test by triggering the verification process when routes change in the Routing Information Base (RIB) that may affect BGP next-hop reachability.

BGP receives notifications from the RIB when the next-hop information changes (event-driven notifications). BGP is notified when any of the following events occurs:

- Next hop becomes unreachable.
- Next hop becomes reachable.
- Fully recursed Interior Gateway Protocol (IGP) metric to the next hop changes.
- First hop IP address or first hop interface changes.
- Next hop becomes connected.
- Next hop becomes unconnected.
- Next hop becomes a local address.
- Next hop becomes a nonlocal address.



---

**Note** Reachability and recursed metric events trigger a best-path recalculation.

---

Event notifications from the RIB are classified as critical and noncritical. Notifications for critical and noncritical events are sent in separate batches. However, a noncritical event is sent with the critical events if the noncritical event is pending and there is a request to read the critical events.

- Critical events are related to next-hop reachability, such as the loss of next hops resulting in a switchover to a different path. A change in the IGP metric for a next hop resulting in a switchover to a different path can also be considered a critical event.
- Non-critical events are related to next hops being added without affecting the best path or changing the IGP metric to a single next hop.



---

**Note** Critical and non-critical events can be configured individually on a per address family basis. For more information on address families, see the "Configuring MPLS Layer 3 VPNs" chapter in the *Cisco Nexus 7000 Series NX-OS MPLS Configuration Guide*.

---

## Route Redistribution

You can configure BGP to redistribute static routes or routes from other protocols. You must configure a route map with the redistribution to control which routes are passed into BGP. A route map allows you to filter routes based on attributes such as the destination, origination protocol, route type, route tag, and so on.

Prior to Cisco NX-OS Release 5.2(1), when you redistribute BGP to IGP, iBGP is redistributed as well. To override this behavior, you must insert an additional deny statement into the route map. Beginning with Cisco NX-OS Release 5.2(1), redistribution varies as follows:

- In a non-MPLS VPN scenario, iBGP is not redistributed to IGP by default.
- In an MPLS VPN scenario (route distinguisher configured under a VRF), iBGP is redistributed to IGP by default.

You can use route maps to override the default behavior in both scenarios, but be careful when doing so as incorrect use of route maps can result in network loops. The following examples show how to use route maps to change the default behavior.

You can change the default behavior for scenario 1 by modifying the route map as follows:

```
route-map foo permit 10
 match route-type internal
router ospf 1
 redistribute bgp 100 route-map foo
```

Similarly, you can change the default behavior for scenario 2 by modifying the route map as follows:

```
route-map foo deny 10
 match route-type internal
router ospf 1
 vrf bar
 redistribute bgp 100 route-map foo
```

The default route should be redistributed into BGP or advertised to peers only when **default-information originate** is configured for an Address Family where the command is supported.

BGP should withdraw the default route on removal of default-information originate if it was already advertised. Also, the redistributed path should be removed for the default route.

You can delete the redistributed path for default route using the following command:

```
no default-information originate
```

## BGP Support for Importing Routes from Default VRF

You can import IP prefixes from the global routing table (the default VRF) into any other VRF by using an import policy. The VRF import policy uses a route map to specify the prefixes to be imported into a VRF. The policy can import IPv4 and IPv6 unicast prefixes.

You can configure the maximum number of prefixes that can be imported from the default VRF.



---

**Note** Routes in the BGP default VRF can be imported directly. Any other routes in the global routing table should be redistributed into BGP first.

---

## BGP Support for Exporting Routes to Default VRF

You can export IP prefixes to the default VRF (global routing table) from any other VRF using an export policy. The VRF export policy leaks a VRF route into default VRF BGP table, which will then be installed in the IPv4/IPv6 routing table. The VRF export policy uses a route map to specify the prefixes to be exported to the default VRF. The policy can export IPv4 and IPv6 unicast prefixes.

You can configure the maximum number of prefixes that can be exported to the default VRF to prevent the routing table from being overloaded.

## BFD

This feature supports bidirectional forwarding detection (BFD) for IPv4 only. BFD is a detection protocol designed to provide fast forwarding-path failure detection times. BFD provides subsecond failure detection between two adjacent devices and can be less CPU-intensive than protocol hello messages because some of the BFD load can be distributed onto the data plane on supported modules.

BFD for BGP is supported on eBGP peers and iBGP single-hop peers. Configure the update-source option in neighbor configuration mode for iBGP single-hop peers using BFD.



---

**Note** BFD is not supported on other iBGP peers or for multihop eBGP peers.

---

See the *Cisco Nexus 7000 Series NX-OS Interfaces Configuration Guide* for more information.

## Tuning BGP

You can modify the default behavior of BGP through BGP timers and by adjusting the best-path algorithm.

### BGP Timers

BGP uses different types of timers for neighbor session and global protocol events. Each established session has a minimum of two timers for sending periodic keepalive messages and for timing out sessions when peer keepalives do not arrive within the expected time. In addition, there are other timers for handling specific features. Typically, you configure these timers in seconds. The timers include a random adjustment so that the same timers on different BGP peers trigger at different times.

### Tuning the Best-Path Algorithm

You can modify the default behavior of the best-path algorithm through optional configuration parameters, including changing how the algorithm handles the multi-exit discriminator (MED) attribute and the router ID.

## Multiprotocol BGP

BGP on Cisco NX-OS supports multiple address families. Multiprotocol BGP (MP-BGP) carries different sets of routes depending on the address family. For example, BGP can carry one set of routes for IPv4 unicast routing, one set of routes for IPv4 multicast routing, and one set of routes for IPv6 multicast routing. You can use MP-BGP for reverse-path forwarding (RPF) checks in IP multicast networks.



---

**Note** Because Multicast BGP does not propagate multicast state information, you need a multicast protocol, such as Protocol Independent Multicast (PIM).

---

Use the router address-family and neighbor address-family configuration modes to support multiprotocol BGP configurations. MP-BGP maintains separate RIBs for each configured address family, such as a unicast RIB and a multicast RIB for BGP.

A multiprotocol BGP network is backward compatible but BGP peers that do not support multiprotocol extensions cannot forward routing information, such as address family identifier information, that the multiprotocol extensions carry.



---

**Note** Beginning with Cisco NX-OS Release 6.2(8), BGP supports RFC 5549 which allows an IPv4 prefix to be carried over an IPv6 next hop. Because BGP is running on every hop and all routers are capable of forwarding IPv4 and IPv6 traffic, there is no need to support IPv6 tunnels between any routers. BGP installs IPv4 over an IPv6 route to the Unicast Route Information Base (URIB).

---

## Graceful Restart and High Availability

Cisco NX-OS supports nonstop forwarding and graceful restart for BGP.

You can use nonstop forwarding (NSF) for BGP to forward data packets along known routes in the Forward Information Base (FIB) while the BGP routing protocol information is being restored following a failover. With NSF, BGP peers do not experience routing flaps. During a failover, the data traffic is forwarded through intelligent modules while the standby supervisor becomes active.

If a Cisco NX-OS router experiences a cold reboot, the network does not forward traffic to the router and removes the router from the network topology. In this scenario, BGP experiences a nongraceful restart and removes all routes. When Cisco NX-OS applies the startup configuration, BGP reestablishes peering sessions and relearns the routes.

A Cisco NX-OS router that has dual supervisors can experience a stateful supervisor switchover. During the switchover, BGP uses nonstop forwarding to forward traffic based on the information in the FIB, and the system is not removed from the network topology. A router whose neighbor is restarting is referred to as a "helper." After the switchover, a graceful restart operation begins. When it is in progress, both routers reestablish their neighbor relationship and exchange their BGP routes. The helper continues to forward prefixes pointing to the restarting peer, and the restarting router continues to forward traffic to peers even though those neighbor relationships are restarting. When the restarting router has all route updates from all BGP peers that are graceful restart capable, the graceful restart is complete, and BGP informs the neighbors that it is operational again.

When a router detects that a graceful restart operation is in progress, both routers exchange their topology tables. When the router has route updates from all BGP peers, it removes all the stale routes and runs the best-path algorithm on the updated routes.

After the switchover, Cisco NX-OS applies the running configuration, and BGP informs the neighbors that it is operational again.

For single-hop iBGP peers with update-source configured under neighbor configuration mode, the peer supports fast external fall-over.

With the additional BGP paths feature, if the number of paths advertised for a given prefix is the same before and after restart, the choice of path ID guarantees the final state and removal of stale paths. If fewer paths are advertised for a given prefix after a restart, stale paths can occur on the graceful restart helper peer.

## Low Memory Handling

BGP reacts to low memory for the following conditions:

- Minor alert—BGP does not establish any new eBGP peers. BGP continues to establish new iBGP peers and confederate peers. Established peers remain, but reset peers are not re-established.
- Severe alert—BGP shuts down select established eBGP peers every two minutes until the memory alert becomes minor. For each eBGP peer, BGP calculates the ratio of total number of paths received to the number of paths selected as best paths. The peers with the highest ratio are selected to be shut down to reduce memory usage. You must clear a shutdown eBGP peer before you can bring the eBGP peer back up to avoid oscillation.




---

**Note** You can exempt important eBGP peers from this selection process.

---

- Critical alert—BGP gracefully shuts down all the established peers. You must clear a shutdown BGP peer before you can bring the BGP peer back up.

## ISSU

Cisco NX-OS supports in-service software upgrades (ISSU). ISSU allows you to upgrade software without impacting forwarding.

The following conditions are required to support ISSU:

- Graceful restart must be enabled (default)
- Keepalive and hold timers must not be smaller than their default values

If either of these requirements is not met, Cisco NX-OS issues a warning. You can proceed with the upgrade or downgrade, but service might be disrupted.




---

**Note** Cisco NX-OS cannot guarantee ISSU for non-default timer values if the negotiated hold time between BGP peers is less than the system switchover time.

---

## Virtualization Support

Cisco NX-OS supports multiple instances of BGP that run on the same system. BGP supports virtual routing and forwarding (VRF) instances that exist within virtual device contexts (VDCs). You can configure one BGP instance in a VDC, but you can have multiple VDCs on the system.

By default, Cisco NX-OS places you in the default VDC and default VRF unless you specifically configure another VDC and VRF.



## Prerequisites for Advanced BGP

Advanced BGP has the following prerequisites:

- You must enable BGP.
- You should have a valid router ID configured on the system.
- You must have an AS number, either assigned by a Regional Internet Registry (RIR) or locally administered.
- You must have reachability (such as an interior gateway protocol [IGP], a static route, or a direct connection) to the peer that you are trying to make a neighbor relationship with.
- You must explicitly configure an address family under a neighbor for the BGP session establishment.

## Guidelines and Limitations for Advanced BGP

Advanced BGP has the following configuration guidelines and limitations:

- The dynamic AS number prefix peer configuration overrides the individual AS number configuration inherited from a BGP template.
- If you configure a dynamic AS number for prefix peers in an AS confederation, BGP establishes sessions with only the AS numbers in the local confederation.
- Command **ttl-security hops** is visible but not supported for Nexus 7K platform, it is supported only for Nexus 9K platform.
- BGP sessions created through a dynamic AS number prefix peer ignore any configured eBGP multihop time-to-live (TTL) value or a disabled check for directly connected peers.
- Configure a router ID for BGP to avoid automatic router ID changes and session flaps.
- Use the maximum-prefix configuration option per peer to restrict the number of routes received and system resources used.
- Configure the update source to establish a session with eBGP multihop sessions.
- Specify a BGP route map if you configure a redistribution.
- Configure the BGP router ID within a VRF.
- If you decrease the keepalive and hold timer values, the network might experience session flaps.
- When you redistribute BGP to IGP, iBGP is redistributed as well. To override this behavior, you must insert an additional deny statement into the route map.
- Cisco NX-OS does not support multi-hop BFD. BFD for BGP has the following limitations:
  - BFD is supported only for BGP IPv4.
  - BFD is supported only for eBGP peers and iBGP single-hop peers.
  - To enable BFD for iBGP single-hop peers, you must configure the update-source option on the physical interface.

- BFD is not supported for multi-hop iBGP peers and multi-hop eBGP peers.
- For single-hop iBGP peers with update-source configured under neighbor configuration mode, the peer supports fast external fall-over.
- The following guidelines and limitations apply to the **remove-private-as** command:
  - It applies only to eBGP peers.
  - It can be configured only in neighbor configuration mode and not in neighbor-address-family mode.
  - If the AS-path includes both private and public AS numbers, the private AS numbers are not removed.
  - If the AS-path contains the AS number of the eBGP neighbor, the private AS numbers are not removed.
  - Private AS numbers are removed only if all AS numbers in that AS-path belong to a private AS number range. Private AS numbers are not removed if a peer's AS number or a non-private AS number is found in the AS-path segment.
- BGP conditional route injection is available only for IPv4 and IPv6 unicast address families in all VRF instances.
- The **match interface** command is only supported for **redistribute** command **route-maps**.
- When sending a route advertisement to an iBGP peer, NXOS sets the interface IP address through which the announced network is reachable for the peer as the next hop instead of preserving the original next hop of the non locally originated route.

This occurs with the 'network' statement and route 'redistribution' configurations in BGP.

The knobs 'set ip next-hop redist-unchanged' or 'set ipv6 next-hop redist-unchanged' available under route-map configuration mode helps to resolve this issue. These knobs are available from Cisco NX-OS Release 6.2(12) onwards.

## Default Settings

| Parameters           | Default     |
|----------------------|-------------|
| BGP feature          | Disabled    |
| BGP additional paths | Disabled    |
| Hold timer           | 180 seconds |
| Keep alive interval  | 60 seconds  |
| Dynamic capability   | Enabled     |

# Configuring Advanced BGP

## Configuring BGP Session Templates

You can use BGP session templates to simplify the BGP configuration for multiple BGP peers with similar configuration needs. BGP templates allow you to reuse common configuration blocks. You configure BGP templates first and then apply these templates to BGP peers.

With BGP session templates, you can configure session attributes such as inheritance, passwords, timers, and security.

A peer-session template can inherit from one other peer-session template. You can configure the second template to inherit from a third template. The first template also inherits this third template. This indirect inheritance can continue for up to seven peer-session templates.

Any attributes configured for the neighbor take priority over any attributes inherited by that neighbor from a BGP template.



**Note** Use the **show bgp neighbor** command to see the template applied. See the *Cisco Nexus 7000 Series NX-OS Unicast Routing Command Reference*, for details on all commands available in the template.

### Before you begin

- You must enable BGP.
- Ensure that you are in the correct VDC (or use the **switchto vdc** command).



**Note** When editing a template, you can use the **no** form of a command at either the peer or template level to explicitly override a setting in a template. You must use the **default** form of the command to reset that attribute to the default state.

### Procedure

|               | Command or Action                                                            | Purpose                                                                                                                             |
|---------------|------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configure terminal</b>                                            | Enters configuration mode.                                                                                                          |
| <b>Step 2</b> | switch(config)# <b>router bgp</b><br><i>autonomous-system-number</i>         | Enables BGP and assigns the autonomous system number to the local BGP speaker.                                                      |
| <b>Step 3</b> | switch(config-router)# <b>template peer-session</b><br><i>template-name</i>  | Enters peer-session template configuration mode.                                                                                    |
| <b>Step 4</b> | switch(config-router-stmp)# <b>password</b> <i>number</i><br><i>password</i> | (Optional) Adds the clear text password test to the neighbor. The password is stored and displayed in type 3 encrypted form (3DES). |

|                | Command or Action                                                                 | Purpose                                                                                                                                                  |
|----------------|-----------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 5</b>  | switch(config-router-stmp)# <b>timers</b> <i>keepalive hold</i>                   | (Optional) Adds the BGP keepalive and holdtimer values to the peer-session template. The default keepalive interval is 60. The default hold time is 180. |
| <b>Step 6</b>  | switch(config-router-stmp)# <b>exit</b>                                           | Exits peer-session template configuration mode.                                                                                                          |
| <b>Step 7</b>  | switch(config-router)# <b>neighbor</b> <i>ip-address remote-as as-number</i>      | Places the router in the neighbor configuration mode for BGP routing and configures the neighbor IP address.                                             |
| <b>Step 8</b>  | switch(config-router-neighbor)# <b>inherit peer-session</b> <i>template-name</i>  | Applies a peer-session template to the peer.                                                                                                             |
| <b>Step 9</b>  | switch(config-router-neighbor)# <b>description</b> <i>text</i>                    | (Optional) Adds a description for the neighbor.                                                                                                          |
| <b>Step 10</b> | switch(config-router-neighbor)# <b>show bgp peer-session</b> <i>template-name</i> | (Optional) Displays the peer-policy template.                                                                                                            |
| <b>Step 11</b> | switch(config-router-neighbor)# <b>copy running-config startup-config</b>         | (Optional) Saves this configuration change.                                                                                                              |

### Example

This example shows how to configure a BGP peer-session template and apply it to a BGP peer:

```
switch# configure terminal
switch(config)# router bgp 65535
switch(config-router)# template peer-session BaseSession
switch(config-router-stmp)# timers 30 90
switch(config-router-stmp)# exit
switch(config-router)# neighbor 192.168.1.2 remote-as 65535
switch(config-router-neighbor)# inherit peer-session BaseSession
switch(config-router-neighbor)# description Peer Router A
switch(config-router-neighbor)# address-family ipv4 unicast
switch(config-router-neighbor)# copy running-config startup-config
```

## Configuring BGP Peer-Policy Templates

You can configure a peer-policy template to define attributes for a particular address family. You assign a preference to each peer-policy template and these templates are inherited in the order specified, for up to five peer-policy templates in a neighbor address family.

Cisco NX-OS evaluates multiple peer policies for an address family using the preference value. The lowest preference value is evaluated first. Any attributes configured for the neighbor take priority over any attributes inherited by that neighbor from a BGP template.

Peer-policy templates can configure address family-specific attributes such as AS-path filter lists, prefix lists, route reflection, and soft reconfiguration.



**Note** Use the **show bgp neighbor** command to see the template applied. See the *Cisco Nexus 7000 Series NX-OS Unicast Routing Command Reference*, for details on all commands available in the template.

#### Before you begin

- You must enable BGP.
- Ensure that you are in the correct VDC (or use the **switchto vdc** command).



**Note** When editing a template, you can use the **no** form of a command at either the peer or template level to explicitly override a setting in a template. You must use the **default** form of the command to reset that attribute to the default state.

#### Procedure

|                | Command or Action                                                                                                      | Purpose                                                                                                                        |
|----------------|------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b>  | switch# <b>configure terminal</b>                                                                                      | Enters configuration mode.                                                                                                     |
| <b>Step 2</b>  | switch(config)# <b>router bgp</b><br><i>autonomous-system-number</i>                                                   | Enables BGP and assigns the autonomous system number to the local BGP speaker.                                                 |
| <b>Step 3</b>  | switch(config-router)# <b>template peer-policy</b><br><i>template-name</i>                                             | Creates a peer-policy template.                                                                                                |
| <b>Step 4</b>  | switch(config-router-ptmp)#<br><b>advertise-active-only</b>                                                            | (Optional) Advertises only active routes to the peer.                                                                          |
| <b>Step 5</b>  | switch(config-router-ptmp)# <b>maximum-prefix</b><br><i>number</i>                                                     | (Optional) Sets the maximum number of prefixes allowed from this peer.                                                         |
| <b>Step 6</b>  | switch(config-router-ptmp)# <b>exit</b>                                                                                | Exits peer-policy template configuration mode.                                                                                 |
| <b>Step 7</b>  | switch(config-router)# <b>neighbor ip-address</b><br><b>remote-as as-number</b>                                        | Places the router in the neighbor configuration mode for BGP routing and configures the neighbor IP address.                   |
| <b>Step 8</b>  | switch(config-router-neighbor)#<br><b>address-family {ipv4   ipv6   vpnv4   vpnv6}</b><br><b>{multicast   unicast}</b> | Enters global address family configuration mode.                                                                               |
| <b>Step 9</b>  | switch(config-router-neighbor-af)# <b>inherit</b><br><b>peer-policy template-name preference</b>                       | Applies a peer-policy template to the peer address family configuration and assigns the preference value for this peer policy. |
| <b>Step 10</b> | switch(config-router-neighbor-af)# <b>show bgp</b><br><b>peer-policy template-name</b>                                 | (Optional) Displays the peer-policy template.                                                                                  |

|                | Command or Action                                                            | Purpose                                     |
|----------------|------------------------------------------------------------------------------|---------------------------------------------|
| <b>Step 11</b> | switch(config-router-neighbor-af)# <b>copy running-config startup-config</b> | (Optional) Saves this configuration change. |

### Example

This example shows how to configure a BGP peer-policy template and apply it to a BGP peer:

```
switch# configure terminal
switch(config)# router bgp 65535
switch(config-router)# template peer-session BasePolicy
switch(config-router-ptmp)# maximum-prefix 20
switch(config-router-ptmp)# exit
switch(config-router)# neighbor 192.168.1.1 remote-as 65535
switch(config-router-neighbor)# address-family ipv4 unicast
switch(config-router-neighbor-af)# inherit peer-policy BasePolicy
switch(config-router-neighbor-af)# copy running-config startup-config
```

## Configuring BGP Peer Templates

You can configure BGP peer templates to combine session and policy attributes in one reusable configuration block. Peer templates can also inherit peer-session or peer-policy templates. Any attributes configured for the neighbor take priority over any attributes inherited by that neighbor from a BGP template. You configure only one peer template for a neighbor, but that peer template can inherit peer-session and peer-policy templates.

Peer templates support session and address family attributes, such as eBGP multihop time-to-live, maximum prefix, next-hop self, and timers.



**Note** Use the **show bgp neighbor** command to see the template applied. See the *Cisco Nexus 7000 Series NX-OS Unicast Routing Command Reference*, for details on all commands available in the template.

### Before you begin

- You must enable BGP.
- Ensure that you are in the correct VDC (or use the **switchto vdc** command).



**Note** When editing a template, you can use the **no** form of a command at either the peer or template level to explicitly override a setting in a template. You must use the **default** form of the command to reset that attribute to the default state.

### Procedure

|               | Command or Action                 | Purpose                           |
|---------------|-----------------------------------|-----------------------------------|
| <b>Step 1</b> | switch# <b>configure terminal</b> | Enters global configuration mode. |

|                | Command or Action                                                                                                                                             | Purpose                                                                                                                                                  |
|----------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 2</b>  | switch(config)# <b>router bgp</b><br><i>autonomous-system-number</i>                                                                                          | Enables BGP and assigns the autonomous system number to the local BGP speaker.                                                                           |
| <b>Step 3</b>  | switch(config-router)# <b>template peer</b><br><i>template-name</i>                                                                                           | Enter peer template configuration mode.                                                                                                                  |
| <b>Step 4</b>  | switch(config-router-neighbor)# <b>inherit</b><br><b>peer-session</b> <i>template-name</i>                                                                    | (Optional) Inherits a peer-session template in the peer template.                                                                                        |
| <b>Step 5</b>  | switch(config-router-neighbor)#<br><b>address-family</b> { <b>ipv4</b>   <b>ipv6</b>   <b>vpnv4</b>   <b>vpnv6</b> }<br>{ <b>multicast</b>   <b>unicast</b> } | (Optional) Configures the global address family configuration mode.                                                                                      |
| <b>Step 6</b>  | switch(config-router-neighbor-af)# <b>inherit</b><br><b>peer</b> <i>template-name</i>                                                                         | (Optional) Applies a peer template to the neighbor address family configuration.                                                                         |
| <b>Step 7</b>  | switch(config-router-neighbor-af)# <b>exit</b>                                                                                                                | Exits BGP neighbor address family configuration mode.                                                                                                    |
| <b>Step 8</b>  | switch(config-router-neighbor)# <b>timers</b><br><i>keepalive hold</i>                                                                                        | (Optional) Adds the BGP timer values to the peer.<br><br>These values override the timer values in the peer-session template, BaseSession.               |
| <b>Step 9</b>  | switch(config-router-neighbor)# <b>exit</b>                                                                                                                   | Exits BGP peer template configuration mode.                                                                                                              |
| <b>Step 10</b> | switch(config-router)# <b>neighbor</b> <i>ip-address</i><br><b>remote-as</b> <i>as-number</i>                                                                 | Places the router in neighbor configuration mode for BGP routing and configures the neighbor IP address.                                                 |
| <b>Step 11</b> | switch(config-router-neighbor)# <b>inherit peer</b><br><i>template-name</i>                                                                                   | Inherits the peer template.                                                                                                                              |
| <b>Step 12</b> | switch(config-router-neighbor)# <b>timers</b><br><i>keepalive hold</i>                                                                                        | (Optional) Adds the BGP timer values to this neighbor.<br><br>These values override the timer values in the peer template and the peer-session template. |
| <b>Step 13</b> | switch(config-router-neighbor-af)# <b>show bgp</b><br><b>peer-template</b> <i>template-name</i>                                                               | (Optional) Displays the peer template.                                                                                                                   |
| <b>Step 14</b> | switch(config-router-neighbor-af)# <b>copy</b><br><b>running-config startup-config</b>                                                                        | (Optional) Saves this configuration change.                                                                                                              |

### Example

This example shows how to configure a BGP peer template and apply it to a BGP peer:

```
switch# configure terminal
switch(config)# router bgp 65535
switch(config-router)# template peer BasePeer
switch(config-router-neighbor)# inherit peer-session BaseSession
```

```

switch(config-router-neighbor)# address-family ipv4 unicast
switch(config-router-neighbor-af)# inherit peer-policy BasePolicy 1
switch(config-router-neighbor-af)# exit
switch(config-router-neighbor)# exit
switch(config-router)# neighbor 192.168.1.2 remote-as 65535
switch(config-router-neighbor)# inherit peer BasePeer
switch(config-router-neighbor)# copy running-config startup-config

```

## Configuring Prefix Peering

BGP supports the definition of a set of peers using a prefix for both IPv4 and IPv6. This feature allows you to not have to add each neighbor to the configuration.

When defining a prefix peering, you must specify the remote AS number with the prefix. BGP accepts any peer that connects from that prefix and autonomous system if the prefix peering does not exceed the configured maximum peers allowed.



**Note** Use the **show ip bgp neighbor** command to show the details of the configuration for that prefix peering with a list of the currently accepted instances and the counts of active, maximum concurrent, and total accepted peers.

Use the **show bgp coverage private** command to display details of the prefix peer wait timer.

### Procedure

|               | Command or Action                                                                               | Purpose                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|---------------|-------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | (Optional) switch(config-router-neighbor)#<br><b>timers prefix-peer-timeout</b> <i>interval</i> | Configures the BGP prefix peering timeout value. When a BGP peer that is part of a prefix peering disconnects, the peer structures are held for a defined prefix peer timeout value which enables the peer to reset and reconnect without danger of being blocked. The timeout range is from 0 to 1200 seconds. The default value is 30.                                                                                                                                                                                                                      |
| <b>Step 2</b> | (Optional) switch(config-router-neighbor)#<br><b>timers prefix-peer-wait</b> <i>interval</i>    | Configures the BGP prefix peering wait timer on a per-VRF basis or on the default VRF. You can use the <b>timers prefix-peer-wait</b> command to disable the peer prefix wait time so that there is no delay before BGP prefixes are inserted into the routing information base (RIB). The range of the <i>interval</i> is from 0 to 1200 seconds. The default value is 90.<br><br><b>Note</b> The timer is only applicable for BGP dynamic neighbors. It is only set when BGP is restarted or is coming up for the first time for the dynamic BGP neighbors. |



|               | Command or Action                                                               | Purpose                                                                                                                     |
|---------------|---------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------|
| <b>Step 3</b> | (Optional) switch(config-router-neighbor)#<br><b>maximum-peers</b> <i>value</i> | Configures the maximum number of peers for this prefix peering in neighbor configuration mode. The range is from 1 to 1000. |

### Example

This example shows how to configure a prefix peering that accepts up to 10 peers:

```
switch(config)# router bgp 65535
switch(config-router)# timers prefix-peer-timeout 120
switch(config-router)# neighbor 10.100.200.0/24 remote-as 65535
switch(config-router-neighbor)# maximum-peers 10
switch(config-router-neighbor)# address-family ipv4 unicast
switch(config-router-neighbor-af)#
```

## Configuring BGP Authentication

You can configure BGP to authenticate route updates from peers using MD5 digests.

To configure BGP to use MD5 authentication, use the following command in neighbor configuration mode:

### Procedure

|               | Command or Action                                                            | Purpose                                                                                                   |
|---------------|------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch(config-router-neighbor)# <b>password</b> {0<br>  3   7} <i>string</i> | Configures an MD5 password (for authentication) for BGP neighbor sessions in neighbor configuration mode. |

## Resetting a BGP Session

If you modify a route policy for BGP, you must reset the associated BGP peer sessions. If the BGP peers do not support route refresh, you can configure a soft reconfiguration for inbound policy changes. Cisco NX-OS automatically attempts a soft reset for the session.

To configure soft reconfiguration inbound, use the following command in neighbor address-family configuration mode.

### Procedure

|               | Command or Action                                                         | Purpose                                                                                                                                                                                                             |
|---------------|---------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch(config-router-neighbor-af)#<br><b>soft-reconfiguration inbound</b> | This command in neighbor address-family configuration mode, enables soft reconfiguration to store the inbound BGP route updates. This command triggers an automatic soft clear or refresh of BGP neighbor sessions. |

|               | Command or Action                                                                                            | Purpose                                                                               |
|---------------|--------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------|
| <b>Step 2</b> | switch# <b>clear bgp</b> {ipv4   ipv6   vpv4   vpv6} {unicast   multicast <i>ip-address</i> soft {in   out}} | This command in any mode resets the BGP session without tearing down the TCP session. |

## Modifying the Next-Hop Address

You can modify the next-hop address used in a route advertisement in the following ways:

- Disable next-hop calculation and use the local BGP speaker address as the next-hop address.
- Set the next-hop address as a third-party address. Use this feature in situations where the original next-hop address is on the same subnet as the peer that the route is being sent to. Using this feature saves an extra hop during forwarding.

To modify the next-hop address, use the following commands in address-family configuration mode:

### Procedure

|               | Command or Action                                                 | Purpose                                                                                                                                                         |
|---------------|-------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch(config-router-neighbor-af)#<br><b>next-hop-self</b>        | Uses the local BGP speaker address as the next-hop address in route updates. This command triggers an automatic soft clear or refresh of BGP neighbor sessions. |
| <b>Step 2</b> | switch(config-router-neighbor-af)#<br><b>next-hop-third-party</b> | Sets the next-hop address as a third-party address. Use this command for single-hop EBGP peers that do not have next-hop-self configured.                       |

## Configuring BGP Next-Hop Address Tracking

BGP next-hop address tracking is enabled by default and cannot be disabled.

You can modify the delay interval between RIB checks to increase the performance of BGP next-hop tracking.

To modify the BGP next-hop address tracking, use the following commands in address-family configuration mode:

### Procedure

|               | Command or Action                                                                                    | Purpose                                                                                                                                                                                                                                               |
|---------------|------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch(config-router-af)# <b>nexthop trigger-delay</b> {critical   non-critical} <i>milliseconds</i> | Specifies the next-hop address tracking delay timer for critical next-hop reachability routes and for noncritical routes. The range is from 1 to 4294967295 milliseconds. The critical timer default is 3000. The noncritical timer default is 10000. |

|               | Command or Action                                                 | Purpose                                                                                                                                    |
|---------------|-------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 2</b> | switch(config-router-af)# <b>nexthop route-map</b><br><i>name</i> | Specifies a route map to match the BGP next-hop addresses to. The name can be any case-sensitive, alphanumeric string up to 63 characters. |

## Configuring Next-Hop Filtering

BGP next-hop filtering allows you to specify that when a next-hop address is checked with the RIB, the underlying route for that next-hop address is passed through the route map. If the route map rejects the route, the next-hop address is treated as unreachable.

BGP marks all next hops that are rejected by the route policy as invalid and does not calculate the best path for the routes that use the invalid next-hop address.

To configure BGP next-hop filtering, use the following command in address-family configuration mode:

### Procedure

|               | Command or Action                                                 | Purpose                                                                                                                                |
|---------------|-------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch(config-router-af)# <b>nexthop route-map</b><br><i>name</i> | Specifies a route map to match the BGP next-hop route to. The name can be any case-sensitive, alphanumeric string up to 63 characters. |

## Disabling Capabilities Negotiation

You can disable capabilities negotiations to interoperate with older BGP peers that do not support capabilities negotiation.

To disable capabilities negotiation, use the following command in neighbor configuration mode:

### Procedure

|               | Command or Action                                                   | Purpose                                                                                                     |
|---------------|---------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch(config-router-neighbor)#<br><b>dont-capability-negotiate</b> | Disables capabilities negotiation. You must manually reset the BGP sessions after configuring this command. |

## Configuring BGP Additional Paths

Beginning with Cisco NX-OS Release 6.1, BGP supports sending and receiving multiple paths per prefix and advertising such paths.

### Advertising the Capability of Sending and Receiving Additional Paths

You can configure BGP to advertise the capability of sending and receiving additional paths to and from the BGP peers.

**Procedure**

|               | <b>Command or Action</b>                                                                        | <b>Purpose</b>                                                                                                                                                                                                                                  |
|---------------|-------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch(config-router-neighbor-af)#<br><b>[no]capability additional paths send [disable]</b>     | Advertises the capability to send additional paths to the BGP peer. The disable option disables the advertising capability of sending additional paths.<br><br>The no form of this command disables the capability of sending additional paths. |
| <b>Step 2</b> | switch (config-router-neighbor-af)#<br><b>[no]capability additional paths receive [disable]</b> | Advertises the capability to send additional paths to the BGP peer. The disable option disables the advertising capability of sending additional paths.<br><br>The no form of this command disables the capability of sending additional paths. |
| <b>Step 3</b> | switch(config-router-neighbor-af)# <b>show bgp neighbor</b>                                     | Displays whether the local peer has advertised the additional paths send or receive capability to the remote peer.                                                                                                                              |
| <b>Step 4</b> | switch(config-router-neighbor-af)# <b>copy running-config startup-config</b>                    | Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.                                                                                                                   |

**Example**

This example shows how to configure BGP to advertise the capability to send and receive additional paths to the BGP peer:

```
switch(config)# router bgp 100
switch(config-router)# neighbor 10.131.31.2 remote-as 100
switch(config-router-neighbor)# address-family ipv4 unicast
switch(config-router-neighbor-af)# capability additional-paths send
switch(config-router-neighbor-af)# capability additional-paths receive
switch(config-router-neighbor-af)# show bgp neighbor
switch(config-router-neighbor-af)# copy running-config startup-config
```

**Configuring the Sending and Receiving of Additional Paths**

You can configure the capability of sending and receiving additional paths to and from the BGP peers.

**Procedure**

|               | <b>Command or Action</b>                                               | <b>Purpose</b>                                                                                                                                     |
|---------------|------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch(config-router-neighbor-af)#<br><b>[no]additional-paths send</b> | Enables the send capability of additional paths for all of the neighbors under this address family for which the capability has not been disabled. |

|               | Command or Action                                                            | Purpose                                                                                                                                                                                                                                       |
|---------------|------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|               |                                                                              | The no form of this command disables the send capability.                                                                                                                                                                                     |
| <b>Step 2</b> | switch(config-router-neighbor-af)#<br>[no]additional-paths receive [disable] | Enables the receive capability of additional paths for all of the neighbors under this address family for which the capability has not been disabled.<br><br>The no form of this command disables the capability of sending additional paths. |
| <b>Step 3</b> | switch(config-router-neighbor-af)# show bgp neighbor                         | Displays whether the local peer has advertised the additional paths send or receive capability to the remote peer.                                                                                                                            |
| <b>Step 4</b> | switch(config-router-neighbor-af)# copy running-config startup-config        | Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.                                                                                                                 |

### Example

This example shows how to enable the additional paths send and receive capability for neighbors under the specified address family for which this capability has not been disabled.:

```
switch(config)# router bgp 100
switch(config-router)# neighbor 10.131.31.2 remote-as 100
switch(config-router-neighbor)# address-family ipv4 unicast
switch(config-router-neighbor-af)# additional-paths send
switch(config-router-neighbor-af)# additional-paths receive
switch(config-router-neighbor-af)# show bgp neighbor
switch(config-router-neighbor-af)# copy running-config startup-config
```

## Configuring Advertised Paths

You can specify the paths that are advertised for BGP.

### Procedure

|               | Command or Action                                                                                        | Purpose                                                                                                                                        |
|---------------|----------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch(config-route-map)# [no]set path-selection all advertise                                           | Specifies that all paths be advertised for a given prefix.<br><br>The no form of this command specifies that only the best path be advertised. |
| <b>Step 2</b> | switch(config-route-map)# show bgp neighbor {ipv4   ipv6} unicastip-address   ipv6-prefix [ vrfvrf-name] | Displays whether the local peer has advertised the additional paths send or receive capability to the remote peer.                             |

|               | Command or Action                                                   | Purpose                                                                                                                       |
|---------------|---------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 3</b> | switch(config-route-map)# <b>copy running-config startup-config</b> | Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration. |

### Example

This example shows how to specify that all paths be advertised for the specified prefix:

```
switch(config)# route-map PATH_SELECTION_RMAP
switch(config-route-map)#match ip address prefeix-list pl
switch(config-route-map)# show bgp ip4 unicast
switch(config-route-map)# copy running-config startup-config
```

## Configuring Additional Path Selection

You can configure the capability of selecting additional paths for a prefix.

### Procedure

|               | Command or Action                                                                                              | Purpose                                                                                                                                        |
|---------------|----------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch(config-router-af)# <b>[no]additional-paths selection route-map</b> <i>map-name</i>                      | Specifies that all paths be advertised for a given prefix.<br><br>The no form of this command specifies that only the best path be advertised. |
| <b>Step 2</b> | switch(config-router-af)# <b>show bgp {ipv4   ipv6} unicast</b> <i>ip-address   ipv6-prefix [ vrfvrf-name]</i> | Displays whether the local peer has advertised the additional paths send or receive capability to the remote peer.                             |
| <b>Step 3</b> | (Optional) switch(config-router-af)# <b>copy running-config startup-config</b>                                 | Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.                  |

### Example

This example shows how to specify that all paths be advertised for the specified prefix:

```
switch(config)# router bgp 100
switch(config-router)# address-family ipv4 unicast
switch(config-router-af)#additional-paths selection route-map PATH_SELECTION_RMAP
switch(config-router-af)# copy running-config startup-config
```

## Configuring eBGP

### Disabling eBGP Single-Hop Checking

You can configure eBGP to disable checking whether a single-hop eBGP peer is directly connected to the local router. Use this option for configuring a single-hop loopback eBGP session between directly connected switches.

To disable checking whether or not a single-hop eBGP peer is directly connected, use the following command in neighbor configuration mode:

#### Procedure

|               | Command or Action                                                 | Purpose                                                                                                                                           |
|---------------|-------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch(config-router-neighbor)#<br><b>disable-connected-check</b> | Disables checking whether or not a single-hop eBGP peer is directly connected. You must manually reset the BGP sessions after using this command. |

### Configuring eBGP Multihop

You can configure the eBGP time-to-live (TTL) value to support eBGP multihop. In some situations, an eBGP peer is not directly connected to another eBGP peer and requires multiple hops to reach the remote eBGP peer. You can configure the eBGP TTL value for a neighbor session to allow these multihop sessions.

To configure eBGP multihop, use the following command in neighbor configuration mode:

#### Procedure

|               | Command or Action                                                        | Purpose                                                                                                                                         |
|---------------|--------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch(config-router-neighbor)# <b>ebgp-multihop</b><br><i>ttl-value</i> | Configures the eBGP TTL value for eBGP multihop. The range is from 2 to 255. You must manually reset the BGP sessions after using this command. |

### Disabling a Fast External Fallover

By default, the Cisco Nexus 7000 Series device supports fast external fallover for neighbors in all VRFs and address-families (IPv4 or IPv6).

Typically, when a BGP router loses connectivity to a directly connected eBGP peer, BGP triggers a fast external fallover by resetting the eBGP session to the peer. You can disable this fast external fallover to limit the instability caused by link flaps.

To disable fast external fallover, use the following command in router configuration mode:

**Procedure**

|               | Command or Action                                       | Purpose                                                                               |
|---------------|---------------------------------------------------------|---------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch(config-router)# <b>no fast-external-fallover</b> | Disables a fast external fallover for eBGP peers. This command is enabled by default. |

**Limiting the AS-path Attribute**

You can configure eBGP to discard routes that have a high number of AS numbers in the AS-path attribute.

To discard routes that have a high number of AS numbers in the AS-path attribute, use the following command in router configuration mode:

**Procedure**

|               | Command or Action                                       | Purpose                                                                                                                   |
|---------------|---------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch(config-router)# <b>maxas-limit</b> <i>number</i> | Discards eBGP routes that have a number of AS-path segments that exceed the specified limit. The range is from 1 to 2000. |

**Configuring Local AS Support**

The local-AS feature allows a router to appear to be a member of a second autonomous system (AS), in addition to its real AS. Local AS allows two ISPs to merge without modifying peering arrangements. Routers in the merged ISP become members of the new autonomous system but continue to use their old AS numbers for their customers.

This feature can only be used for true eBGP peers. You cannot use this feature for two peers that are members of different confederation sub-autonomous systems.

To configure eBGP local AS support, use the following command in neighbor configuration mode:

**Procedure**

|               | Command or Action                                                                                                          | Purpose                                                                                                                                                                                                                             |
|---------------|----------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch(config-router-neighbor)# <b>local-as</b> <i>number</i> [ <b>no-prepend</b> [ <b>replace-as</b> [ <b>dual-as</b> ]]] | Configures eBGP to prepend the local AS number to the AS_PATH attribute. The AS number can be a 16-bit integer or a 32-bit integer in the form of a higher 16-bit decimal number and a lower 16-bit decimal number in xx.xx format. |

**Example**

The local-AS feature under VRF configuration mode is supported for both IBGP and EBGP neighbor relationships.

This following example shows how to configure the feature for the IBGP neighbor 10.1.2.1:

```
router bgp 65001
 vrf BGP1
```



```

local-as 65002
address-family ipv4 unicast
neighbor 10.1.2.1 remote-as 65002

```

The **local-as** command must be configured in the neighbor configuration mode for eBGP or a warning message is displayed stating that the local AS cannot be same as the remote AS. The following example shows how to configure the local-AS feature for eBGP:

```

router bgp 65001
vrf BGP1
neighbor 20.1.2.1 remote-as 65003
local-as 65001

```

## Configuring AS Confederations

To configure an AS confederation, you must specify a confederation identifier. To the outside world, the group of autonomous systems within the AS confederation look like a single autonomous system with the confederation identifier as the autonomous system number.

### Procedure

|               | Command or Action                                                                      | Purpose                                                                                                                                                                                                                                                                                                 |
|---------------|----------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch(config-router)# <b>confederation identifier</b><br><i>as-number</i>             | In router configuration mode, this command configures a BGP confederation identifier.<br><br>The command triggers an automatic notification and session reset for the BGP neighbor sessions.                                                                                                            |
| <b>Step 2</b> | switch(config-router)# <b>bgp confederation peers</b> <i>as-number [as-number2...]</i> | In router configuration mode, this command configures the autonomous systems that belong to the AS confederation.<br><br>The command specifies a list of autonomous systems that belong to the confederation and it triggers an automatic notification and session reset for the BGP neighbor sessions. |

## Configuring Route Reflector

You can configure iBGP peers as route reflector clients to the local BGP speaker, which acts as the route reflector. Together, a route reflector and its clients form a cluster. A cluster of clients usually has a single route reflector. In such instances, the cluster is identified by the router ID of the route reflector. To increase redundancy and avoid a single point of failure in the network, you can configure a cluster with more than one route reflector. You must configure all route reflectors in the cluster with the same 4-byte cluster ID so that a route reflector can recognize updates from route reflectors in the same cluster.

### Before you begin

- You must enable BGP.
- Ensure that you are in the correct VDC (or use the **switchto vdc** command).

## Procedure

|                | Command or Action                                                                                                                                                              | Purpose                                                                                                                                                                                                                 |
|----------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b>  | switch# <b>configure terminal</b>                                                                                                                                              | Enters global configuration mode.                                                                                                                                                                                       |
| <b>Step 2</b>  | switch(config)# <b>router bgp</b> <i>as-number</i>                                                                                                                             | Enters BGP mode and assigns the autonomous system number to the local BGP speaker.                                                                                                                                      |
| <b>Step 3</b>  | switch(config-router)# <b>cluster-id</b> <i>cluster-id</i>                                                                                                                     | Configures the local router as one of the route reflectors that serve the cluster. You specify a cluster ID to identify the cluster. This command triggers an automatic soft clear or refresh of BGP neighbor sessions. |
| <b>Step 4</b>  | switch(config-router)# <b>address-family</b> { <b>ipv4</b>   <b>ipv6</b>   <b>vpn4</b>   <b>vpn6</b> } { <b>unicast</b>   <b>multicast</b> }                                   | Enters router address family configuration mode for the specified address family.                                                                                                                                       |
| <b>Step 5</b>  | switch(config-router-af)# <b>client-to-client reflection</b>                                                                                                                   | (Optional) Configures client-to-client route reflection. This feature is enabled by default. This command triggers an automatic soft clear or refresh of BGP neighbor sessions.                                         |
| <b>Step 6</b>  | switch(config-router-neighbor)# <b>exit</b>                                                                                                                                    | Exits router address configuration mode.                                                                                                                                                                                |
| <b>Step 7</b>  | switch(config-router)# <b>neighbor</b> <i>ip-address</i> <b>remote-as</b> <i>as-number</i>                                                                                     | Configures the IP address and AS number for a remote BGP peer.                                                                                                                                                          |
| <b>Step 8</b>  | switch(config-router-neighbor)# <b>address-family</b> { <b>ipv4</b>   <b>ipv6</b>   <b>vpn4</b>   <b>vpn6</b> } { <b>unicast</b>   <b>multicast</b> }                          | Enters neighbor address family configuration mode for the unicast IPv4 address family.                                                                                                                                  |
| <b>Step 9</b>  | switch(config-router-neighbor-af)# <b>route-reflector-client</b>                                                                                                               | Configures the device as a BGP route reflector and configures the neighbor as its client. This command triggers an automatic notification and session reset for the BGP neighbor sessions.                              |
| <b>Step 10</b> | (Optional) switch(config-router-neighbor-af)# <b>show bgp</b> { <b>ipv4</b>   <b>ipv6</b>   <b>vpn4</b>   <b>vpn6</b> } { <b>unicast</b>   <b>multicast</b> } <b>neighbors</b> | Displays the BGP peers.                                                                                                                                                                                                 |
| <b>Step 11</b> | (Optional) switch(config-router-neighbor-af)# <b>copy running-config startup-config</b>                                                                                        | Saves this configuration change.                                                                                                                                                                                        |

## Example

This example shows how to configure the router as a route reflector and add one neighbor as a client:

```
switch(config)# router bgp 65535
switch(config-router)# neighbor 192.0.2.10 remote-as 65535
switch(config-router-neighbor)# address-family ip unicast
switch(config-router-neighbor-af)# route-reflector-client
switch(config-router-neighbor-af)# copy running-config startup-config
```

## Configuring Next-Hops on Reflected Routes Using an Outbound Route-Map

You can change the next-hop on reflected routes on a BGP route reflector using an outbound route-map. You can configure the outbound route-map to specify the peer's local address as the next-hop address.



**Note** The **next-hop-self** command does not enable this functionality for routes being reflected to clients by a route reflector. This functionality can only be enabled using an outbound route-map.

### Before you begin

You must enable BGP.

Ensure that you are in the correct VDC (or use the **switchto vdc** command).

You must enter the **set next-hop** command to configure an address family specific next-hop address. For example, for the IPv6 address family, you must enter the **set ipv6 next-hop peer-address** command.

- When setting IPv4 next-hops using route-maps—If **set ip next-hop peer-address** matches the route-map, the next-hop is set to the peer's local address. If no next-hop is set in the route-map, the next-hop is set to the one stored in the path.
- When setting IPv6 next-hops using route-maps—If **set ipv6 next-hop peer-address** matches the route-map, the next-hop is set as follows:
  - For IPv6 peers, the next-hop is set to the peer's local IPv6 address.
  - For IPv4 peers, if **update-source** is configured, the next-hop is set to the source interface's IPv6 address, if any. If no IPv6 address is configured, no next-hop is set.
  - For IPv4 peers, if **update-source** is not configured, the next-hop is set to the outgoing interface's IPv6 address, if any. If no IPv6 address is configured, no next-hop is set.

### Procedure

|               | Command or Action                                                                                                                                           | Purpose                                                                                        |
|---------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configure terminal</b>                                                                                                                           | Enters global configuration mode.                                                              |
| <b>Step 2</b> | switch(config)# <b>router bgp</b> <i>as-number</i>                                                                                                          | Enters BGP mode and assigns the autonomous system number to the local BGP speaker.             |
| <b>Step 3</b> | switch(config-router)# <b>neighbor</b> <i>ip-address</i><br><b>remote-as</b> <i>as-number</i>                                                               | Configures the IP address and AS number for a remote BGP peer.                                 |
| <b>Step 4</b> | switch(config-router-neighbor)# <b>update-source</b><br><i>interface number</i>                                                                             | (Optional) Specifies and updates the source of the BGP session.                                |
| <b>Step 5</b> | switch(config-router-neighbor)#<br><b>address-family</b> { <b>ipv4</b>   <b>ipv6</b>   <b>vpn4</b>   <b>vpn6</b> }<br>{ <b>unicast</b>   <b>multicast</b> } | Enters router address family configuration mode for the specified address family.              |
| <b>Step 6</b> | switch(config-router-neighbor-af)#<br><b>route-reflector-client</b>                                                                                         | Configures the device as a BGP route reflector and configures the neighbor as its client. This |

|               | Command or Action                                                                                                                                                              | Purpose                                                                                     |
|---------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------|
|               |                                                                                                                                                                                | command triggers an automatic notification and session reset for the BGP neighbor sessions. |
| <b>Step 7</b> | switch(config-router-neighbor-af)# <b>route-map map-name out</b>                                                                                                               | Applies the configured BGP policy to outgoing routes.                                       |
| <b>Step 8</b> | (Optional) switch(config-router-neighbor-af)# <b>show bgp {ipv4   ipv6   vpnv4   vpnv6} {unicast   multicast} [ip-address   ipv6-prefix] route-map map-name [vrf vrf-name]</b> | Displays the BGP routes that match the route map.                                           |
| <b>Step 9</b> | (Optional) switch(config-router-neighbor-af)# <b>copy running-config startup-config</b>                                                                                        | Saves this configuration change.                                                            |

### Example

This example shows how to configure the next-hop on reflected routes on a BGP route reflector using an outbound route-map:

```
switch(config)# interface loopback 300
switch(config-if)# ip address 192.0.2.11/32
switch(config-if)# ipv6 address 2001::a0c:1a65/64
switch(config-if)# ip router ospf 1 area 0.0.0.0
switch(config-if)# exit
switch(config)# route-map setrrnh permit 10
switch(config-route-map)# set ip next-hop peer-address
switch(config-route-map)# exit
switch(config)# route-map setrrnhv6 permit 10
switch(config-route-map)# set ipv6 next-hop peer-address
switch(config-route-map)# exit
switch(config)# router bgp 200
switch(config-router)# neighbor 192.0.2.12 remote-as 200
switch(config-router-neighbor)# update-source loopback 300
switch(config-router-neighbor)# address-family ipv4 unicast
switch(config-router-neighbor-af)# route-reflector-client
switch(config-router-neighbor-af)# route-map setrrnh out
switch(config-router-neighbor-af)# exit
switch(config-router-neighbor)# address-family ipv6 unicast
switch(config-router-neighbor-af)# route-reflector-client
switch(config-router-neighbor-af)# route-map setrrnhv6 out
```

## Configuring Route Dampening

You can configure route dampening to minimize route flaps propagating through your iBGP network.

To configure route dampening, use the following command in address-family or VRF address family configuration mode.

**Procedure**

|               | Command or Action                                                                                                                                | Purpose                                                                                                                                                                                                                                                                                                                                                          |
|---------------|--------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch(config-router-af)# <b>dampening</b> [ <i>half-life reuse-limit suppress-limit max-suppress-time</i>   <b>route-map</b> <i>map-name</i> }] | Disables capabilities negotiation. The parameter values are as follows: <ul style="list-style-type: none"> <li>• <i>half-life</i>—The range is from 1 to 45</li> <li>• <i>reuse-limit</i>—The range is from 1 to 20000.</li> <li>• <i>suppress-limit</i>—The range is from 1 to 20000.</li> <li>• <i>max-suppress-time</i>—The range is from 1 to 255</li> </ul> |

## Configuring Load Sharing and ECMP

You can configure the maximum number of paths that BGP adds to the route table for equal-cost multipath load balancing.

To configure the maximum number of paths, use the following command in router address-family configuration mode:

| Command                                                                        | Purpose                                                                                                                                                                                                                                                                                                              |
|--------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| switch(config-router-af)# <b>maximum-paths</b> [ <b>ibgp</b> ] <i>maxpaths</i> | Configures the maximum number of equal-cost paths for load sharing. The range is from 1 to 16. The default is 1. Starting from Cisco NX-OS Release 8.4(1), the range is from 1 to 64 on M3- and F3-Series I/O modules. Starting from Cisco NX-OS Release 8.4(2), the range is from 1 to 64 on F4-Series I/O modules. |

## Configuring Maximum Prefixes

You can configure the maximum number of prefixes that BGP can receive from a BGP peer. If the number of prefixes exceeds this value, you can optionally configure BGP to generate a warning message or tear down the BGP session to the peer.

To configure the maximum allowed prefixes for a BGP peer, use the following command in neighbor address-family configuration mode:

**Procedure**

|               | Command or Action                                                                                                                         | Purpose                                                                                                                                                                                      |
|---------------|-------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch(config-router-neighbor-af)# <b>maximum-prefix</b> <i>maximum</i> [ <i>threshold</i> ] [ <i>restarttime</i>   <b>warning-only</b> ] | Configure the maximum number of prefixes from a peer. The parameter ranges are as follows: <ul style="list-style-type: none"> <li>• <i>maximum</i>—The range is from 1 to 300000.</li> </ul> |

|  | Command or Action | Purpose                                                                                                                                                                                                                                                                                                                                  |
|--|-------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|  |                   | <ul style="list-style-type: none"> <li>• <i>threshold</i>—The range is from 1 to 100 percent. The default is 75 percent.</li> <li>• <i>time</i>—The range is from 1 to 65535 minutes.</li> </ul> <p>This command triggers an automatic notification and session reset for the BGP neighbor sessions if the prefix limit is exceeded.</p> |

## Configuring Dynamic Capability

You can configure dynamic capability for a BGP peer.

To configure dynamic capability, use the following command in neighbor configuration mode:

| Command                                                      | Purpose                                                                                                                      |
|--------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------|
| switch(config-router-neighbor)#<br><b>dynamic-capability</b> | Enables dynamic capability. This command triggers an automatic notification and session reset for the BGP neighbor sessions. |

## Configuring Aggregate Addresses

You can configure aggregate address entries in the BGP route table.

To configure an aggregate address, use the following command in router address-family configuration mode:

| Command                                                                                                                                                                                                                             | Purpose                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>aggregate-address</b><br><i>ip-prefix/length</i> [ <b>as-set</b> ]<br>[ <b>summary-only</b> ] [ <b>advertise-map</b><br><i>map-name</i> ] [ <b>attribute-map</b><br><i>map-name</i> ] [ <b>suppress-map</b><br><i>map-name</i> ] | Creates an aggregate address. The path advertised for this route is an autonomous system set that consists of all elements contained in all paths that are being summarized: <ul style="list-style-type: none"> <li>• The <b>as-set</b> keyword generates autonomous system set path information and community information from contributing paths.</li> <li>• The <b>summary-only</b> keyword filters all more specific routes from updates.</li> <li>• The <b>advertise-map</b> <i>map-name</i> keyword and argument specify the route map used to select attribute information from selected routes.</li> <li>• The <b>attribute-map</b> <i>map-name</i> keyword and argument specify the route map used to select attribute information from the aggregate.</li> <li>• The <b>suppress-map</b> <i>map-name</i> keyword and argument conditionally filter more specific routes.</li> </ul> |

## Unsuppressing the Advertisement of Aggregated Routes

You can configure BGP to advertise routes that are suppressed by the **aggregate-address** command.

To unsuppress the advertising of aggregated routes, use the following command in router neighbor address-family configuration mode:

### Procedure

|               | Command or Action                                                          | Purpose                                                                                  |
|---------------|----------------------------------------------------------------------------|------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch(config-route-neighbor-af)#<br><b>unsuppress-map</b> <i>map-name</i> | Advertises selective routes that are suppressed by the <b>aggregate-address</b> command. |

## Configuring BGP Conditional Route Injection

You can configure BGP conditional route injection to inject specific routes based on the administrative policy or traffic engineering information and control the packets being forwarded to these specific routes, which are injected into the BGP routing table only if the configured conditions are met. This feature allows you to improve the accuracy of common route aggregation by conditionally injecting or replacing less specific prefixes with more specific prefixes. Only prefixes that are equal to or more specific than the original prefix can be injected.



**Note** The injected prefixes inherit the attributes of the aggregated route.

### Before you begin

- You must enable BGP
- Ensure that you are in the correct VDC (or use the **switchto vdc** command).

### Procedure

|               | Command or Action                                                                                                                    | Purpose                                                                                                                                                                                                                                           |
|---------------|--------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configure terminal</b>                                                                                                    | Enters global configuration mode.                                                                                                                                                                                                                 |
| <b>Step 2</b> | switch(config)# <b>router bgp</b> <i>as-number</i>                                                                                   | Enters BGP configuration mode and assigns the autonomous system number to the local BGP speaker.                                                                                                                                                  |
| <b>Step 3</b> | switch(config-router)# <b>address-family</b> { <b>ipv4</b>   <b>ipv6</b> } <b>unicast</b>                                            | Enters address family configuration mode.                                                                                                                                                                                                         |
| <b>Step 4</b> | switch(config-router-af)# <b>inject-map</b> <i>inject-map-name</i> <b>exist-map</b> <i>exist-map-name</i> [ <b>copy-attributes</b> ] | Specifies the inject-map and exist-map routes for conditional route injection. These maps install one or more prefixes into a BGP routing table. The <i>exist-map</i> route map specifies the prefixes that BGP tracks, and the <i>inject-map</i> |

|                | Command or Action                                                                                | Purpose                                                                                                                                                                                                                  |
|----------------|--------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|                |                                                                                                  | route map defines the prefixes that are created and installed into the local BGP table.<br><br>Use the <b>copy-attributes</b> keyword to specify that the injected route inherits the attributes of the aggregate route. |
| <b>Step 5</b>  | switch(config-router-af)# <b>exit</b>                                                            | Exits address family configuration mode.                                                                                                                                                                                 |
| <b>Step 6</b>  | switch(config-router)# <b>exit</b>                                                               | Exits BGP configuration mode.                                                                                                                                                                                            |
| <b>Step 7</b>  | switch(config)# <b>ip prefix-list</b> <i>list-name seq sequence-number permit network-length</i> | Configures a prefix list. Repeat this step for every prefix list to be created.                                                                                                                                          |
| <b>Step 8</b>  | switch(config)# <b>route-map</b> <i>map-name permit sequence-number</i>                          | Configures a route-map and enters route-map configuration mode.                                                                                                                                                          |
| <b>Step 9</b>  | switch(config-route-map)# <b>match ip address prefix-list</b> <i>prefix-list-name</i>            | Specifies the aggregate route to which a more specific route will be injected.                                                                                                                                           |
| <b>Step 10</b> | switch(config-route-map)# <b>match ip route-source prefix-list</b> <i>prefix-list-name</i>       | Specifies the match conditions for the source fo the route.                                                                                                                                                              |
| <b>Step 11</b> | switch(config-route-map)# <b>exit</b>                                                            | Exits route-map configuration mode.                                                                                                                                                                                      |
| <b>Step 12</b> | switch(config)# <b>ip prefix-list</b> <i>list-name seq sequence-number permit network-length</i> | Configures a prefix list. Repeat this step for every prefix list to be created.                                                                                                                                          |
| <b>Step 13</b> | switch(config)# <b>route-map</b> <i>map-name permit sequence-number</i>                          | Configures a route map and enters route-map configuration mode.                                                                                                                                                          |
| <b>Step 14</b> | switch(config-route-map)# <b>set ip address prefix-list</b>                                      | Specifies the routes to be injected.                                                                                                                                                                                     |
| <b>Step 15</b> | (Optional) switch(config-route-map)# <b>show bgp {ipv4   ipv6} unicast injected-routes</b>       | Displays injected routes in the routing table.                                                                                                                                                                           |
| <b>Step 16</b> | (Optional) switch(config)# <b>copy running-config startup-config</b>                             | Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.                                                                                            |

## Configuring BGP Conditional Advertisement

You can configure BGP conditional advertisement to limit the routes that BGP propagates. You define the following two route maps:

- **Advertise map**—Specifies the conditions that the route must match before BGP considers the conditional advertisement. This route map can contain any appropriate match statements.
- **Exist map or nonexist map**—Defines the prefix that must exist in the BGP table before BGP propagates a route that matches the advertise map. The nonexist map defines the prefix that must not exist in the BGP table before BGP propagates a route that matches the advertise map. BGP processes only the permit statements in the prefix list match statements in these route maps.



If the route does not pass the condition, BGP withdraws the route if it exists in the BGP table.

### Before you begin

- You must enable BGP.
- Ensure that you are in the correct VDC (or use the **switchto vdc** command).

### Procedure

|               | Command or Action                                                                                                                                      | Purpose                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|---------------|--------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configure terminal</b>                                                                                                                      | Enters configuration mode.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Step 2</b> | switch(config)# <b>router bgp</b> <i>as-number</i>                                                                                                     | Enters BGP mode and assigns the autonomous system number to the local BGP speaker.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Step 3</b> | switch(config-router)# <b>neighbor</b> <i>ip-address</i><br><b>remote-as</b> <i>as-number</i>                                                          | Places the router in neighbor configuration mode for BGP routing and configures the neighbor IP address.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Step 4</b> | switch(config-router-neighbor)#<br><b>address-family</b> { <i>ipv4 ipv6 vpnv4 vpnv6</i> }<br>{ <b>unicast multicast</b> }                              | Enters address family configuration mode.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Step 5</b> | switch(config-router-neighbor-af)#<br><b>advertise-map</b> <i>adv-map</i> { <b>exist-map</b><br><i>exist-rmap non-exist-map</i> <i>nonexist-rmap</i> } | Configures BGP to conditionally advertise routes based on the two configured route maps: <ul style="list-style-type: none"> <li>• <b>adv-map</b>—Specifies a route map with match statements that the route must pass before BGP passes the route to the next route map. The adv-map is a case-sensitive, alphanumeric string up to 63 characters.</li> <li>• <b>exist-rmap</b>—Specifies a route map with match statements for a prefix list. A prefix in the BGP table must match a prefix in the prefix list before BGP advertises the route. The exist-rmap is a case-sensitive, alphanumeric string up to 63 characters.</li> <li>• <b>nonexist-rmap</b>—Specifies a route map with match statements for a prefix list. A prefix in the BGP table must not match a prefix in the prefix list before BGP advertises the route. The nonexist-rmap is a case-sensitive, alphanumeric string up to 63 characters.</li> </ul> |
| <b>Step 6</b> | (Optional) switch(config-router-neighbor-af)#<br><b>show ip bgp neighbor</b>                                                                           | Displays information about BGP and the configured conditional advertisement route maps.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |

|               | Command or Action                                                                          | Purpose                          |
|---------------|--------------------------------------------------------------------------------------------|----------------------------------|
| <b>Step 7</b> | (Optional) switch(config-router-neighbor-af)#<br><b>copy running-config startup-config</b> | Saves this configuration change. |

### Example

This example shows how to configure BGP conditional advertisement:

```
switch# configure terminal
switch(config)# router bgp 65535
switch(config-router)# neighbor 192.0.2.2 remote-as 65537
switch(config-router-neighbor)# address-family ipv4 unicast
switch(config-router-neighbor-af)# advertise-map advertise exist-map exist
switch(config-router-neighbor-af)# exit
switch(config-router-neighbor)# exit
switch(config-router)# exit
switch(config)# route-map advertise
switch(config-route-map)# match as-path pathList
switch(config-route-map)# exit
switch(config)# route-map exit
switch(config-route-map)# match ip address prefix-list plist
switch(config-route-map)# exit
switch(config)# ip prefix-list plist permit 209.165.201.0/27
```

## Configuring Route Redistribution

You can configure BGP to accept routing information from another routing protocol and redistribute that information through the BGP network. Optionally, you can assign a default metric for redistributed routes.

### Before you begin

- You must enable BGP.
- Ensure that you are in the correct VDC (or use the **switchto vdc** command).

### Procedure

|               | Command or Action                                                                                                                                                                                                | Purpose                                                                            |
|---------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configure terminal</b>                                                                                                                                                                                | Enters global configuration mode.                                                  |
| <b>Step 2</b> | switch(config)# <b>router bgp</b> <i>as-number</i>                                                                                                                                                               | Enters BGP mode and assigns the autonomous system number to the local BGP speaker. |
| <b>Step 3</b> | switch(config-router)# <b>address-family</b> { <b>ipv4</b>   <b>ipv6</b>   <b>vpn4</b>   <b>vpn6</b> } { <b>unicast</b>   <b>multicast</b> }                                                                     | Enters address family configuration mode.                                          |
| <b>Step 4</b> | switch(config-router-af)# <b>redistribute</b> { <b>direct</b>   { <b>eigrp</b>   <b>isis</b>   <b>ospf</b>   <b>ospfv3</b>   <b>rip</b> } <i>instance-tag</i>   <b>static</b> } <b>route-map</b> <i>map-name</i> | Redistributes routes from other protocols into BGP.                                |
| <b>Step 5</b> | (Optional) switch(config-router-af)# <b>default-metric</b> <i>value</i>                                                                                                                                          | Generates a default metric into BGP.                                               |

|               | Command or Action                                                                             | Purpose                          |
|---------------|-----------------------------------------------------------------------------------------------|----------------------------------|
| <b>Step 6</b> | (Optional) <code>switch(config-router-neighbor-af)# copy running-config startup-config</code> | Saves this configuration change. |

### Example

This example shows how to redistribute EIGRP into BGP:

```
switch# configure terminal
switch(config)# router bgp 65535
switch(config-router)# address-family ipv4 unicast
switch(config-router-af)# redistribute eigrp 201 route-map Eigrpmap
switch(config-router-af)# copy running-config startup-config
```

## Advertising the Default Route

You can configure BGP to advertise the default route (network 0.0.0.0).

### Before you begin

You must enable BGP.

Ensure that you are in the correct VDC (or use the `switchto vdc` command).

### Procedure

|               | Command or Action                                                                        | Purpose                                                                                                                       |
|---------------|------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | <code>switch# configure terminal</code>                                                  | Enters global configuration mode.                                                                                             |
| <b>Step 2</b> | <code>switch(config)# route-map allow permit</code>                                      | Enters router map configuration mode and defines the conditions for redistributing routes                                     |
| <b>Step 3</b> | <code>switch(config-route-map)# exit</code>                                              | Exits router map configuration mode.                                                                                          |
| <b>Step 4</b> | <code>switch(config)# ip route ip-address network-mask null null-interface-number</code> | Configures the IP address.                                                                                                    |
| <b>Step 5</b> | <code>switch(config)# router bgp as-number</code>                                        | Enters BGP mode and assigns the AS number to the local BGP speaker.                                                           |
| <b>Step 6</b> | <code>switch(config-router)# address-family {ipv4   ipv6   vpnv4   vpnv6} unicast</code> | Enters address family configuration mode.                                                                                     |
| <b>Step 7</b> | <code>switch(config-router-af)# default-information originate</code>                     | Advertises the default route.                                                                                                 |
| <b>Step 8</b> | <code>switch(config-router-af)# redistribute static route-map allow</code>               | Redistributes the default route.                                                                                              |
| <b>Step 9</b> | (Optional) <code>switch(config)# copy running-config startup-config</code>               | Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration. |

## Configuring Route Import from Default VRF to any other VRF

Perform the following steps to import routes from default VRF to any other non-default VRF.

### Before you begin

- Enable BGP.
- Ensure that you are in the correct VDC.

### Procedure

---

**Step 1** Enter the global configuration mode:

```
switch#configure terminal
```

**Step 2** Enable BGP:

```
switch(config)#feature bgp
```

**Step 3** Create a new VRF and enter VRF configuration mode:

```
switch(config)#vrf context vrf-name
```

**Step 4** Enter the IPv4 / IPv6 unicast address family configuration mode:

```
switch(config-vrf)# address-family {ipv4 | ipv6} unicast
```

**Step 5** Configure an import policy for a VRF to import prefixes from the default VRF:

```
switch(config-vrf-af)# import vrf default [prefix-limit] map route-map
```

*prefix-limit* limits the number of routes that can be imported. Default value is 1000.

*route-map* specifies the route-map to be imported and can be case-sensitive, alphanumeric string up to 63 characters.

---

## Configuring Route Export from BGP VRF to Default VRF

Perform the following steps to export routes from non-default VRF to Default VRF.

### Before you begin

- Enable BGP.
- Ensure that you are in the correct VDC.

### Procedure

---

**Step 1** Enter the global configuration mode:

```
switch#configure terminal
```

**Step 2** Enable BGP:

```
switch(config)#feature bgp
```

**Step 3** Create a new VRF and enter VRF configuration mode:

```
switch(config)#vrf context vrf-name
```

**Step 4** Enter the IPv4 / IPv6 unicast address family configuration mode:

```
switch(config-vrf)# address-family {ipv4 | ipv6} unicast
```

**Step 5** Export IPv4 or IPv6 prefixes from non-default VRF to default VRF, filtered by *route-map*:

```
switch(config-vrf-af)# export vrf default [prefix-limit] map route-map
```

*prefix-limit* limits the number of routes that can be exported, in order to avoid the global table being overloaded. Default value is 1000.

*route-map* can be case-sensitive, alphanumeric string up to 63 characters. It specifies the route-map.

If the route map does not exist, the command will be accepted but processed at a later time when the route map is created.

## Example

The following example shows how to export the route map, BgpMap, to default VRF, and verify the configuration.

```
switch# configure terminal
switch(config)# feature bgp
switch(config)#vrf context vpn1
switch(config-vrf)# address-family ipv4 unicast
switch(config-vrf-af)# export vrf default 3 map BgpMap
switch(config-vrf-af)# exit
switch(config)#show bgp process vrf vpn1
```

Information regarding configured VRFs:

BGP Information for VRF vpn1

```
VRF Id : 3
VRF state : UP
Router-ID : 20.0.0.1
Configured Router-ID : 0.0.0.0
Confed-ID : 0
Cluster-ID : 0.0.0.0
No. of configured peers : 2
No. of pending config peers : 0
No. of established peers : 2
VRF RD : 100:1
```

Information for address family IPv4 Unicast in VRF vpn1

```
Table Id : 3
Table state : UP
Peers Active-peers Routes Paths Networks Aggregates
1 1 6 6 0 0
```

Redistribution

```

static, route-map allow

Export RT list:
 100:1
 1000:1
Import RT list:
 100:1
Label mode: per-prefix
Aggregate label: 492287
Import default limit : 1000
Import default prefix count : 2
Import default map : allow
Export default limit : 1000
Export default prefix count : 3
Export default map : allow

```

## Configuring Multiprotocol BGP

You can configure MP-BGP to support multiple address families, including IPv4 and IPv6 unicast and multicast routes.

### Before you begin

- You must enable BGP.
- Ensure that you are in the correct VDC (or use the **switchto vdc** command).

### Procedure

|               | Command or Action                                                                                                                                           | Purpose                                                                                                  |
|---------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configure terminal</b>                                                                                                                           | Enters configuration mode.                                                                               |
| <b>Step 2</b> | switch(config)# <b>router bgp</b> <i>as-number</i>                                                                                                          | Enters BGP mode and assigns the autonomous system number to the local BGP speaker.                       |
| <b>Step 3</b> | switch(config-router)# <b>neighbor</b> <i>ip-address</i><br><b>remote-as</b> <i>as-number</i>                                                               | Places the router in neighbor configuration mode for BGP routing and configures the neighbor IP address. |
| <b>Step 4</b> | switch(config-router-neighbor)#<br><b>address-family</b> { <b>ipv4</b>   <b>ipv6</b>   <b>vpn4</b>   <b>vpn6</b> }<br>{ <b>unicast</b>   <b>multicast</b> } | Enters address family configuration mode.                                                                |
| <b>Step 5</b> | (Optional) switch(config-router-neighbor-af)#<br><b>copy running-config startup-config</b>                                                                  | Saves this configuration change.                                                                         |

### Example

This example shows how to enable advertising and receiving IPv4 and IPv6 routes for multicast RPF for a neighbor:

```

switch# configure terminal
switch(config)# interface ethernet 2/1
switch(config-if)# ipv6 address 2001:0DB8::1

```

```

switch(config-if)# router bgp 65535
switch(config-router)# neighbor 192.168.1.2 remote-as 35537
switch(config-router-neighbor)# address-family ipv4 multicast
switch(config-router-neighbor-af)# exit
switch(config-router-neighbor)# address-family ipv6 multicast
switch(config-router-neighbor-af)# copy running-config startup-config

```

## Configuring Policy-Based Administrative Distance

You can configure a distance for external BGP (eBGP) and internal BGP (iBGP) routes that match a policy described in the configured route map. The distance configured in the route map is downloaded to the unicast RIB along with the matching routes. BGP uses the best path to determine the administrative distance when downloading next hops in the unicast RIB table. If there is no match or a deny clause in the policy, BGP uses the distance configured in the distance command or the default distance for routes.

The policy-based administrative distance feature is useful when there are two or more different routes to the same destination from two different routing protocols.

### Before you begin

You must enable BGP.

Ensure that you are in the correct VDC (or use the **switchto vdc** command).

### Procedure

|               | Command or Action                                                                                 | Purpose                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|---------------|---------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configure terminal</b>                                                                 | Enters global configuration mode.                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Step 2</b> | switch(config)# <b>ip prefix-list</b> <i>name seq number permit prefix-length</i>                 | Creates a prefix list to match IP packets or routes with the permit keyword.                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Step 3</b> | switch(config)# <b>route-map</b> <i>map-tag permit sequence-number</i>                            | Creates a route map and enters route-map configuration mode with the permit keyword. If the match criteria for the route is met in the policy, the packet is policy routed.                                                                                                                                                                                                                                                                                    |
| <b>Step 4</b> | switch(config-route-map)# <b>match ip address prefix-list</b> <i>prefix-list-name</i>             | Matches IPv4 network routes based on a prefix list. The prefix-list name can be any alphanumeric string up to 63 characters.                                                                                                                                                                                                                                                                                                                                   |
| <b>Step 5</b> | switch(config-route-map)# <b>set distance</b> <i>&lt;value1&gt; &lt;value2&gt; &lt;value3&gt;</i> | Specifies the administrative distance for interior BGP (iBGP) or exterior BGP (eBGP) routes and BGP routes originated in the local autonomous system. The range is from 1 to 255.<br><br>After you enter the value for the external administrative distance, you must enter the value for the administrative distance for the internal routes or/and the value for the administrative distance for the local routes depending on your requirement; so that the |

|                | Command or Action                                                                                                       | Purpose                                                                                                                                                                                                                                                                                                                 |
|----------------|-------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|                |                                                                                                                         | internal/local routes are also considered in the route administration.                                                                                                                                                                                                                                                  |
| <b>Step 6</b>  | switch(config-route-map)# <b>exit</b>                                                                                   | Exits route-map configuration mode.                                                                                                                                                                                                                                                                                     |
| <b>Step 7</b>  | switch(config)# <b>router bgp</b> <i>as-number</i>                                                                      | Enters BGP mode and assigns the AS number to the local BGP speaker.                                                                                                                                                                                                                                                     |
| <b>Step 8</b>  | switch(config-router)# <b>address-family</b> { <b>ipv4</b>   <b>ipv6</b>   <b>vpnv4</b>   <b>vpnv6</b> } <b>unicast</b> | Enters address family configuration mode.                                                                                                                                                                                                                                                                               |
| <b>Step 9</b>  | switch(config-router-af)# <b>table-map</b> <i>map-name</i>                                                              | Configures the selective administrative distance for a route map for BGP routes before forwarding them to the RIB table. The table-map name can be any alphanumeric string up to 63 characters.<br><br><b>Note</b> You can also configure the <b>table-map</b> command under the VRF address-family configuration mode. |
| <b>Step 10</b> | (Optional) switch(config-router-af)# <b>show forwarding distribution</b>                                                | Displays forwarding information distribution.                                                                                                                                                                                                                                                                           |
| <b>Step 11</b> | (Optional) switch(config)# <b>copy running-config startup-config</b>                                                    | Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.                                                                                                                                                                                           |

## Tuning BGP

You can tune BGP characteristics through a series of optional parameters.

### Procedure

|               | Command or Action                                                                                                                                                                                                                                           | Purpose                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|---------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | Required: switch(config-router)# <b>bestpath</b> [ <b>always-compare-med</b>   <b>as-path multipath-relax</b>   <b>compare-routerid</b>   <b>cost-community ignore</b>   <b>med</b> { <b>confed</b>   <b>missing-as-worst</b>   <b>non-deterministic</b> }] | Modifies the best-path algorithm. The optional parameters are as follows: <ul style="list-style-type: none"> <li>• <b>always-compare-med</b>—Compares MED on paths from different autonomous systems.</li> <li>• <b>as-path multipath-relax</b>—Allows load sharing across the providers with different (but equal-length) AS paths. Without this option, the AS paths must be identical for load sharing.</li> <li>• <b>compare-routerid</b>—Compares the router IDs for identical eBGP paths.</li> </ul> |



|               | Command or Action                                                                                                  | Purpose                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|---------------|--------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|               |                                                                                                                    | <ul style="list-style-type: none"> <li>• <b>cost-community ignore</b>—Ignores the cost community for BGP best-path calculations. For more information on the BGP cost community, see the “Configuring MPLS Layer 3 VPN Load Balancing” chapter of the Cisco Nexus 7000 Series NX-OS MPLS Configuration Guide.</li> <li>• <b>med confed</b>—Forces bestpath to do a MED comparison only between paths originated within a confederation.</li> <li>• <b>med missing-as-worst</b>—Treats a missing MED as the highest MED.</li> <li>• <b>med non-deterministic</b>—Does not always pick the best MED path from among the paths from the same autonomous system.</li> </ul> |
| <b>Step 2</b> | switch(config-router)# <b>enforce-first-as</b>                                                                     | Enforces the neighbor autonomous system to be the first AS number listed in the AS_path attribute for eBGP.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Step 3</b> | switch(config-router)# <b>log-neighbor-changes</b>                                                                 | Generates a system message when a neighbor changes state.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Step 4</b> | switch(config-router)# <b>router-id id</b>                                                                         | Manually configures the router ID for this BGP speaker.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Step 5</b> | switch(config-router)# <b>timers [bestpath-delay delay   bgp keepalive holdtime   prefix-peer-timeout timeout]</b> | <p>Sets the BGP timer values. The optional parameters are as follows:</p> <ul style="list-style-type: none"> <li>• <b>delay</b>—Initial best-path timeout value after a restart. The range is from 0 to 3600 seconds. The default value is 300.</li> <li>• <b>keepalive</b>—BGP session keepalive time. The range is from 0 to 3600 seconds. The default value is 60.</li> <li>• <b>holdtime</b>—BGP session hold time. The range is from 0 to 3600 seconds. The default value is 180.</li> <li>• <b>timeout</b>—Prefix peer timeout value. The range is from 0 to 1200 seconds. The default value is 30.</li> </ul>                                                    |

|                | Command or Action                                                                              | Purpose                                                                                                                                                                                                                                                                                                                                                                        |
|----------------|------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 6</b>  | switch(config-router-af)# <b>distance</b><br><i>ebgp-distance ibgp-distance local-distance</i> | Sets the administrative distance for BGP. The range is from 1 to 255. The defaults are as follows: <ul style="list-style-type: none"> <li>• <b>ebgp-distance</b>—20.</li> <li>• <b>ibgp-distance</b>—200.</li> <li>• <b>local-distance</b>—220. Local-distance is the administrative distance used for aggregate discard routes when they are installed in the RIB.</li> </ul> |
| <b>Step 7</b>  | switch(config-router-neighbor)# <b>description</b><br><i>string</i>                            | Sets a descriptive string for this BGP peer. The string can be up to 80 alphanumeric characters.                                                                                                                                                                                                                                                                               |
| <b>Step 8</b>  | switch(config-router-neighbor)# <b>low-memory exempt</b>                                       | Exempts this BGP neighbor from a possible shutdown due to a low memory condition.                                                                                                                                                                                                                                                                                              |
| <b>Step 9</b>  | switch(config-router-neighbor)# <b>transport connection-mode passive</b>                       | Allows a passive connection setup only. This BGP speaker does not initiate a TCP connection to a BGP peer. You must manually reset the BGP sessions after configuring this command.                                                                                                                                                                                            |
| <b>Step 10</b> | <b>remove-private-as</b>                                                                       | Removes private AS numbers from outbound route updates to an eBGP peer. This command triggers an automatic soft clear or refresh of BGP neighbor sessions. <p><b>Note</b> See the “Guidelines and Limitations for Advanced BGP” section for more information on this command.</p>                                                                                              |
| <b>Step 11</b> | switch(config-router-neighbor)# <b>update-source</b> <i>interface-type number</i>              | Configures the BGP speaker to use the source IP address of the configured interface for BGP sessions to the peer. This command triggers an automatic notification and session reset for the BGP neighbor sessions. Single-hop iBGP peers support fast external failover when <b>update-source</b> is configured.                                                               |
| <b>Step 12</b> | switch(config-router-neighbor)# <b>suppress-inactive</b>                                       | Advertises the best (active) routes only to the BGP peer. This command triggers an automatic soft clear or refresh of BGP neighbor sessions.                                                                                                                                                                                                                                   |
| <b>Step 13</b> | switch(config-router-neighbor)# <b>default-originate</b> [ <b>route-map</b> <i>map-name</i> ]  | Generates a default route to the BGP peer.                                                                                                                                                                                                                                                                                                                                     |
| <b>Step 14</b> | switch(config-router-neighbor)# <b>filter-list</b><br><i>list-name {in out}</i>                | Applies an AS path filter list to this BGP peer for inbound or outbound route updates. This                                                                                                                                                                                                                                                                                    |

|                | Command or Action                                                               | Purpose                                                                                                                                                          |
|----------------|---------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|                |                                                                                 | command triggers an automatic soft clear or refresh of BGP neighbor sessions.                                                                                    |
| <b>Step 15</b> | switch(config-router-neighbor)# <b>prefix-list</b><br><i>list-name {in out}</i> | Applies a prefix list to this BGP peer for inbound or outbound route updates. This command triggers an automatic soft clear or refresh of BGP neighbor sessions. |
| <b>Step 16</b> | switch(config-router-neighbor)#<br><b>send-community</b>                        | Sends the community attribute to this BGP peer. This command triggers an automatic soft clear or refresh of BGP neighbor sessions.                               |
| <b>Step 17</b> | switch(config-router-neighbor)#<br><b>send-community extended</b>               | Sends the extended community attribute to this BGP peer. This command triggers an automatic soft clear or refresh of BGP neighbor sessions.                      |

## Configuring a Graceful Restart

You can configure a graceful restart and enable the graceful restart helper feature for BGP.

### Before you begin

You must enable BGP.

Create the VDCs and VRFs.

Ensure that you are in the correct VDC (or use the **switchto vdc** command).

### Procedure

|               | Command or Action                                                                                | Purpose                                                                                                                                                                                                                                                                  |
|---------------|--------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configure terminal</b>                                                                | Enters configuration mode.                                                                                                                                                                                                                                               |
| <b>Step 2</b> | switch(config)# <b>router bgp</b> <i>as-number</i>                                               | Creates a new BGP process with the configured autonomous system number.                                                                                                                                                                                                  |
| <b>Step 3</b> | switch(config-router)# <b>graceful-restart</b>                                                   | Enables a graceful restart and the graceful restart helper functionality. This command is enabled by default.<br><br>This command triggers an automatic notification and session reset for the BGP neighbor sessions.                                                    |
| <b>Step 4</b> | switch(config-router)# <b>graceful-restart</b><br><b>{restart-time time stalepath-time time}</b> | Configures the graceful restart timers.<br><br>The optional parameters are as follows: <ul style="list-style-type: none"> <li>• <b>restart-time</b>—Maximum time for a restart sent to the BGP peer. The range is from 1 to 3600 seconds. The default is 120.</li> </ul> |

|               | Command or Action                                                           | Purpose                                                                                                                                                                                                                                                                                                               |
|---------------|-----------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|               |                                                                             | <ul style="list-style-type: none"> <li>• <b>stalepath-time</b>—Maximum time that BGP keeps the stale routes from the restarting BGP peer. The range is from 1 to 3600 seconds. The default is 300.</li> </ul> <p>This command triggers an automatic notification and session reset for the BGP neighbor sessions.</p> |
| <b>Step 5</b> | switch(config-router)# <b>graceful-restart-helper</b>                       | Enables the graceful restart helper functionality. Use this command if you have disabled graceful restart but you still want to enable graceful restart helper functionality. This command triggers an automatic notification and session reset for the BGP neighbor sessions.                                        |
| <b>Step 6</b> | (Optional) switch(config-router)# <b>show running-config bgp</b>            | Displays the BGP configuration.                                                                                                                                                                                                                                                                                       |
| <b>Step 7</b> | (Optional) switch(config-router)# <b>copy running-config startup-config</b> | Saves this configuration change.                                                                                                                                                                                                                                                                                      |

### Example

This example shows how to enable a graceful restart:

```
switch# configure terminal
switch(config)# router bgp 65535
switch(config-router)# graceful-restart
switch(config-router)# copy running-config startup-config
```

## Configuring Virtualization

You can configure one BGP process in each VDC. You can create multiple VRFs within each VDC and use the same BGP process in each VRF.

### Before you begin

- You must enable BGP.
- Ensure that you are in the correct VDC (or use the **switchto vdc** command).

### Procedure

|               | Command or Action                                  | Purpose                                              |
|---------------|----------------------------------------------------|------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configure terminal</b>                  | Enters configuration mode.                           |
| <b>Step 2</b> | switch(config)# <b>vrf context</b> <i>vrf-name</i> | Creates a new VRF and enters VRF configuration mode. |
| <b>Step 3</b> | switch(config-vrf)# <b>exit</b>                    | Exits VRF configuration mode.                        |

|               | Command or Action                                                                                 | Purpose                                                                               |
|---------------|---------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------|
| <b>Step 4</b> | switch(config)# <b>router bgp</b> <i>as-number</i>                                                | Creates a new BGP process with the configured autonomous system number.               |
| <b>Step 5</b> | switch(config-router)# <b>vrf</b> <i>vrf-name</i>                                                 | Enters the router VRF configuration mode and associates this BGP instance with a VRF. |
| <b>Step 6</b> | switch(config-router-vrf)# <b>neighbor</b> <i>ip-address</i><br><b>remote-as</b> <i>as-number</i> | Configures the IP address and AS number for a remote BGP peer.                        |
| <b>Step 7</b> | (Optional) switch(config-router-neighbor-af)#<br><b>copy running-config startup-config</b>        | Saves this configuration change.                                                      |

### Example

This example shows how to create a VRF and configure the router ID in the VRF:

```
switch# configure terminal
switch(config)# vrf context NewVRF
switch(config-vrf)# exit
switch(config)# router bgp 65535
switch(config-router)# vrf NewVRF
switch(config-router-vrf)# neighbor 209.165.201.1 remote-as 65535
switch(config-router-vrf-neighbor)# copy running-config startup-config
```

## Verifying the Advanced BGP Configuration

To display the BGP configuration, perform one of the following tasks:

| Command                                                                                                                                                                                                                                                                                                                                                    | Purpose                                                      |
|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------|
| <b>show bgp all</b> [ <b>summary</b> ] [ <b>vrf</b> <i>vrf-name</i> ]                                                                                                                                                                                                                                                                                      | Displays the BGP information for all address families.       |
| <b>show bgp convergence vrf</b> <i>vrf-name</i>                                                                                                                                                                                                                                                                                                            | Displays the BGP information for all address families.       |
| <b>show bgp</b> { <b>ipv4</b>   <b>ipv6</b>   <b>vpn4</b>   <b>vpn6</b> } { <b>unicast</b>   <b>multicast</b> }<br>[ <i>ip-address</i>   <i>ipv6-prefix</i> ] <b>community</b> { <b>regexp</b> <i>expression</i>   [ <b>community</b> ]<br>[ <b>no-advertise</b> ] [ <b>no-export</b> ] [ <b>no-export-subconfed</b> ]} [ <b>vrf</b> <i>vrf-name</i> ]     | Displays the BGP routes that match a BGP community.          |
| <b>show bgp</b> [ <b>vrf</b> <i>vrf-name</i> ] { <b>ipv4</b>   <b>ipv6</b>   <b>vpn4</b>   <b>vpn6</b> } { <b>unicast</b>   <b>multicast</b> }<br>[ <i>ip-address</i>   <i>ipv6-prefix</i> ] <b>community-list</b> <i>list-name</i> [ <b>vrf</b> <i>vrf-name</i> ]                                                                                         | Displays the BGP routes that match a BGP community list.     |
| <b>show bgp</b> { <b>ipv4</b>   <b>ipv6</b>   <b>vpn4</b>   <b>vpn6</b> } { <b>unicast</b>   <b>multicast</b> }<br>[ <i>ip-address</i>   <i>ipv6-prefix</i> ] <b>extcommunity</b> { <b>regexp</b> <i>expression</i>   <b>generic</b><br>[ <b>non-transitive</b>   <b>transitive</b> ] <i>aa4:nn</i> [ <b>exact-match</b> ]} [ <b>vrf</b> <i>vrf-name</i> ] | Displays the BGP routes that match a BGP extended community. |

|                                                                                                                                                                                                                                                                    |                                                                                                                                                                                                                                           |
|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>show bgp</b> { <b>ipv4</b>   <b>ipv6</b>   <b>vpn4</b>   <b>vpn6</b> } { <b>unicast</b>   <b>multicast</b> }<br>[ <i>ip-address</i>   <i>ipv6-prefix</i> ] <b>extcommunity-list</b> <i>list-name</i> [ <b>exact-match</b> ]}<br>[ <b>vrf</b> <i>vrf-name</i> ]  | Displays the BGP routes that match a BGP extended community list.                                                                                                                                                                         |
| <b>show bgp</b> { <b>ipv4</b>   <b>ipv6</b>   <b>vpn4</b>   <b>vpn6</b> } { <b>unicast</b>   <b>multicast</b> }<br>[ <i>ip-address</i>   <i>ipv6-prefix</i> ] { <b>dampening dampened-paths</b> [ <b>regex</b> <i>expression</i> ]} [ <b>vrf</b> <i>vrf-name</i> ] | Displays the information for BGP route dampening. Use the <b>clear bgp dampening</b> command to clear the route flap dampening information.                                                                                               |
| <b>show bgp</b> { <b>ipv4</b>   <b>ipv6</b>   <b>vpn4</b>   <b>vpn6</b> } { <b>unicast</b>   <b>multicast</b> }<br>[ <i>ip-address</i>   <i>ipv6-prefix</i> ] <b>history-paths</b> [ <b>regex</b> <i>expression</i> ] [ <b>vrf</b> <i>vrf-name</i> ]               | Displays the BGP route history paths.                                                                                                                                                                                                     |
| <b>show bgp</b> { <b>ipv4</b>   <b>ipv6</b>   <b>vpn4</b>   <b>vpn6</b> } { <b>unicast</b>   <b>multicast</b> }<br>[ <i>ip-address</i>   <i>ipv6-prefix</i> ] <b>filter-list</b> <i>list-name</i> [ <b>vrf</b> <i>vrf-name</i> ]                                   | Displays the information for the BGP filter list.                                                                                                                                                                                         |
| <b>show bgp</b> { <b>ipv4</b>   <b>ipv6</b>   <b>vpn4</b>   <b>vpn6</b> } { <b>unicast</b>   <b>multicast</b> }<br>[ <i>ip-address</i>   <i>ipv6-prefix</i> ] <b>neighbors</b> [ <i>ip-address</i>   <i>ipv6-prefix</i> ] [ <b>vrf</b> <i>vrf-name</i> ]           | Displays the information for BGP peers. Use the <b>clear bgp neighbors</b> command to clear these neighbors.                                                                                                                              |
| <b>show bgp</b> { <b>ipv4</b>   <b>ipv6</b>   <b>vpn4</b>   <b>vpn6</b> } { <b>unicast</b>   <b>multicast</b> }<br>[ <i>ip-address</i>   <i>ipv6-prefix</i> ] { <b>nexthop</b>   <b>nexthop-database</b> } [ <b>vrf</b> <i>vrf-name</i> ]                          | <b>show bgp</b> { <b>ipv4</b>   <b>ipv6</b>   <b>vpn4</b>   <b>vpn6</b> } { <b>unicast</b>   <b>multicast</b> }<br>[ <i>ip-address</i>   <i>ipv6-prefix</i> ] { <b>nexthop</b>   <b>nexthop-database</b> } [ <b>vrf</b> <i>vrf-name</i> ] |
| <b>show bgp paths</b>                                                                                                                                                                                                                                              | Displays the BGP path information.                                                                                                                                                                                                        |
| <b>show bgp</b> { <b>ipv4</b>   <b>ipv6</b>   <b>vpn4</b>   <b>vpn6</b> } { <b>unicast</b>   <b>multicast</b> }<br>[ <i>ip-address</i>   <i>ipv6-prefix</i> ] <b>policy</b> <i>name</i> [ <b>vrf</b> <i>vrf-name</i> ]                                             | Displays the BGP policy information. Use the <b>clear bgp policy</b> command to clear the policy information.                                                                                                                             |
| <b>show bgp</b> { <b>ipv4</b>   <b>ipv6</b>   <b>vpn4</b>   <b>vpn6</b> } { <b>unicast</b>   <b>multicast</b> }<br>[ <i>ip-address</i>   <i>ipv6-prefix</i> ] <b>prefix-list</b> <i>list-name</i> [ <b>vrf</b> <i>vrf-name</i> ]                                   | Displays the BGP routes that match the prefix list.                                                                                                                                                                                       |
| <b>show bgp</b> { <b>ipv4</b>   <b>ipv6</b>   <b>vpn4</b>   <b>vpn6</b> } { <b>unicast</b>   <b>multicast</b> }<br>[ <i>ip-address</i>   <i>ipv6-prefix</i> ] <b>received-paths</b> [ <b>vrf</b> <i>vrf-name</i> ]                                                 | Displays the BGP paths stored for soft reconfiguration.                                                                                                                                                                                   |
| <b>show bgp</b> { <b>ipv4</b>   <b>ipv6</b>   <b>vpn4</b>   <b>vpn6</b> } { <b>unicast</b>   <b>multicast</b> }<br>[ <i>ip-address</i>   <i>ipv6-prefix</i> ] <b>regex</b> <i>expression</i> [ <b>vrf</b> <i>vrf-name</i> ]                                        | Displays the BGP routes that match the AS_path regular expression.                                                                                                                                                                        |
| <b>show bgp</b> { <b>ipv4</b>   <b>ipv6</b>   <b>vpn4</b>   <b>vpn6</b> } { <b>unicast</b>   <b>multicast</b> }<br>[ <i>ip-address</i>   <i>ipv6-prefix</i> ] <b>route-map</b> <i>map-name</i> [ <b>vrf</b> <i>vrf-name</i> ]                                      | Displays the BGP routes that match the route map.                                                                                                                                                                                         |
| <b>show bgp peer-policy</b> <i>name</i> [ <b>vrf</b> <i>vrf-name</i> ]                                                                                                                                                                                             | Displays the information about BGP peer policies.                                                                                                                                                                                         |
| <b>show bgp peer-session</b> <i>name</i> [ <b>vrf</b> <i>vrf-name</i> ]                                                                                                                                                                                            | Displays the information about BGP peer sessions.                                                                                                                                                                                         |

|                                                                                           |                                                                                                                                                                                                  |
|-------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>show bgp peer-template</b> <i>name</i> [ <b>vrf</b> <i>vrf-name</i> ]                  | Displays the information about BGP peer templates. Use the <b>clear bgp peer-template</b> command to clear all neighbors in a peer template.                                                     |
| <b>show bgp process</b>                                                                   | Displays the BGP process information.                                                                                                                                                            |
| <b>show</b> { <b>ipv4</b>   <b>ipv6</b>   <b>vpn4</b>   <b>vpn6</b> } <b>bgp options</b>  | Displays the BGP status and configuration information. This command has multiple options. See the <i>Cisco Nexus 7000 Series NX-OS Unicast Routing Command Reference</i> , for more information. |
| <b>show</b> { <b>ipv4</b>   <b>ipv6</b>   <b>vpn4</b>   <b>vpn6</b> } <b>mbgp options</b> | Displays the BGP status and configuration information. This command has multiple options. See the <i>Cisco Nexus 7000 Series NX-OS Unicast Routing Command Reference</i> , for more information. |
| <b>show running-configuration bgp</b>                                                     | Displays the current running BGP configuration.                                                                                                                                                  |

## Displaying Advanced BGP Statistics

To display advanced BGP statistics, use the following commands:

| Command                                                                                                                                                                                                          | Purpose                                                                                                             |
|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------|
| <b>show bgp</b> { <b>ipv4</b>   <b>ipv6</b>   <b>vpn4</b>   <b>vpn6</b> } { <b>unicast</b>   <b>multicast</b> } [ <i>ip-address</i>   <i>ipv6-prefix</i> ] <b>flap-statistics</b> [ <b>vrf</b> <i>vrf-name</i> ] | Displays the BGP route flap statistics. Use the <b>clear bgp flap-statistics</b> command to clear these statistics. |
| <b>show bgp sessions</b> [ <b>vrf</b> <i>vrf-name</i> ]                                                                                                                                                          | Displays the BGP sessions for all peers. Use the <b>clear bgp sessions</b> command to clear these statistics.       |
| <b>show bgp statistics</b>                                                                                                                                                                                       | Displays the BGP statistics.                                                                                        |

## Related Documents

| Related Topic    | Document Title                                                                  |
|------------------|---------------------------------------------------------------------------------|
| BGP CLI commands | <b>Cisco Nexus 7000 Series NX-OS Unicast Routing Command Reference</b>          |
| VDCs and VRFs    | <b>Cisco Nexus 7000 Series NX-OS Virtual Device Context Configuration Guide</b> |

## RFCs

| RFC      | Title                                                                                                                        |
|----------|------------------------------------------------------------------------------------------------------------------------------|
| RFC 2918 | Route Refresh Capability for BGP-4 <a href="http://www.faqs.org/rfcs/rfc2918.html">http://www.faqs.org/rfcs/rfc2918.html</a> |

## MIBs

| MIBs                                          | MIBs Link                                                                                                                       |
|-----------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------|
| BGP4-MIB<br>CISCO-BGP4-MIB<br>CISCO-BGP-MIBv2 | To locate and download MIBs, go to the following URL: <a href="https://cfnng.cisco.com/mibs">https://cfnng.cisco.com/mibs</a> . |

## Feature History for Advanced BGP

This table includes only the updates for those releases that have resulted in additions or changes to the feature.

**Table 26: Feature History for Advanced BGP**

| Feature Name                         | Release     | Feature Information                                                                                                                |
|--------------------------------------|-------------|------------------------------------------------------------------------------------------------------------------------------------|
| BGP                                  | 7.3(0)D1(1) | Added support for exporting routes to Default VRF                                                                                  |
| BGP                                  | 6.2(8)      | Added support for CISCO-BGP-MIBv2                                                                                                  |
| BGP                                  | 6.2(8)      | Added support for RFC 5549                                                                                                         |
| BGP Next Hop Unchanged               | 6.2(8)      | Introduced this feature.                                                                                                           |
| BGP                                  | 6.2(2)      | Added BFD support for the IPv6 address family.                                                                                     |
| BGP                                  | 6.2(2)      | Added the ability to configure BGP to advertise the default route and introduced the <b>default-information originate</b> command. |
| BGP                                  | 6.2(2)      | Added the ability to advertise routes that are suppressed by the <b>aggregate-address</b> command.                                 |
| Policy-based administrative distance | 6.2(2)      | Introduced this feature.                                                                                                           |



| Feature Name                       | Release | Feature Information                                                                                                      |
|------------------------------------|---------|--------------------------------------------------------------------------------------------------------------------------|
| BGP conditional route injection    | 6.2(2)  | Introduced this feature.                                                                                                 |
| BGP AS-path multipath relax        | 6.0(1)  | Added the <b>as-path multipath-relax</b> option to the <b>bestpath</b> command.                                          |
| BGP outbound route-maps            | 6.0(1)  | Added support for setting next-hops on reflected routes using an outbound route-map.                                     |
| BGP cost community ignore          | 5.2(1)  | Added the <b>cost-community ignore</b> option to the <b>bestpath</b> command.                                            |
| VPN address families               | 5.2(1)  | Added support for VPN address families.                                                                                  |
| BGP                                | 5.1(1)  | No change from Release 5.0.                                                                                              |
| BFD                                | 5.0(2)  | Added support for BFD. See the <i>Cisco Nexus 7000 Series NX-OS Interfaces Configuration Guide</i> for more information. |
| ISSU                               | 4.2(3)  | Lowered the BGP minimum hold-time check to eight seconds.                                                                |
| Next-hop addressing                | 4.2(1)  | Added support for the BGP next-hop address tracking and filtering.                                                       |
| 4-Byte AS numbers                  | 4.2(1)  | Added support for 4-byte AS numbers in plaintext notation.                                                               |
| Conditional advertisement          | 4.2(1)  | Added support for conditionally advertising BGP routes based on the existence of other routes in the BGP table.          |
| Dynamic AS number for prefix peers | 4.1(2)  | Added support for a range of AS numbers for the BGP prefix peer configuration.                                           |
| BGP                                | 4.0(1)  | This feature was introduced.                                                                                             |





# CHAPTER 13

## Configuring RIP

---

This chapter contains the following sections:

- [Finding Feature Information, on page 337](#)
- [Information About RIP, on page 337](#)
- [Prerequisites for RIP, on page 340](#)
- [Guidelines and Limitations for RIP, on page 340](#)
- [Default Settings for RIP Parameters, on page 340](#)
- [Configuring RIP, on page 340](#)
- [Verifying the RIP Configuration, on page 351](#)
- [Displaying RIP Statistics, on page 352](#)
- [Configuration Examples for RIP, on page 352](#)
- [Related Documents for RIP, on page 353](#)
- [Standards for RIP, on page 353](#)
- [Feature History for RIP, on page 353](#)

## Finding Feature Information

Your software release might not support all the features documented in this module. For the latest caveats and feature information, see the Bug Search Tool at <https://tools.cisco.com/bugsearch/> and the release notes for your software release. To find information about the features documented in this module, and to see a list of the releases in which each feature is supported, see the "New and Changed Information" chapter or the Feature History table in this chapter.

## Information About RIP

RIP uses User Datagram Protocol (UDP) data packets to exchange routing information in small internetworks. RIPv2 supports IPv4. RIPv2 uses an optional authentication feature supported by the RIPv2 protocol.



---

**Note** Cisco NX-OS does not support IPv6 for RIP.

---

RIP uses the following two message types:

- Request—Sent to the multicast address 224.0.0.9 to request route updates from other RIP-enabled routers.

- **Response**—Sent every 30 seconds by default. The router also sends response messages after it receives a request message. The response message contains the entire RIP route table. RIP sends multiple response packets for a request if the RIP routing table cannot fit in one response packet.

RIP uses a hop count for the routing metric. The hop count is the number of routers that a packet can traverse before reaching its destination. A directly connected network has a metric of 1; an unreachable network has a metric of 16. This small range of metrics makes RIP an unsuitable routing protocol for large networks.

## RIPv2 Authentication

You can configure authentication on RIP messages to prevent unauthorized or invalid routing updates in your network. Cisco NX-OS supports a simple password or an MD5 authentication digest.

You can configure the RIP authentication per interface by using key-chain management for the authentication keys. Key-chain management allows you to control changes to the authentication keys used by an MD5 authentication digest or simple text password authentication. See the *Cisco Nexus 7000 Series NX-OS Security Configuration Guide*, for more details about creating key-chains.

To use an MD5 authentication digest, you configure a password that is shared at the local router and all remote RIP neighbors. Cisco NX-OS creates an MD5 one-way message digest based on the message itself and the encrypted password and sends this digest with the RIP message (Request or Response). The receiving RIP neighbor validates the digest by using the same encrypted password. If the message has not changed, the calculation is identical and the RIP message is considered valid.

An MD5 authentication digest also includes a sequence number with each RIP message to ensure that no message is replayed in the network.


## Split Horizon

You can use split horizon to ensure that RIP never advertises a route out of the interface where it was learned.

Split horizon is a method that controls the sending of RIP update and query packets. When you enable split horizon on an interface, Cisco NX-OS does not send update packets for destinations that were learned from this interface. Controlling update packets in this manner reduces the possibility of routing loops.

You can use split horizon with poison reverse to configure an interface to advertise routes learned by RIP as unreachable over the interface that learned the routes.

**Figure 38: Sample RIP Network with Split Horizon Poison Reverse Enabled**



Router C learns about route X and advertises that route to Router B. Router B in turn advertises route X to Router A, but sends a route X unreachable update back to Router C.

By default, split horizon is enabled on all interfaces.

## Route Filtering

You can configure a route policy on a RIP-enabled interface to filter the RIP updates. Cisco NX-OS updates the route table with only those routes that the route policy allows.

## Route Summarization

You can configure multiple summary aggregate addresses for a specified interface. Route summarization simplifies route tables by replacing a number of more-specific addresses with an address that represents all the specific addresses. For example, you can replace 10.1.1.0/24, 10.1.2.0/24, and 10.1.3.0/24 with one summary address, 10.1.0.0/16.

If more specific routes are in the routing table, RIP advertises the summary address from the interface with a metric equal to the maximum metric of the more specific routes.



---

**Note** Cisco NX-OS does not support automatic route summarization.

---

## Route Redistribution

You can use RIP to redistribute static routes or routes from other protocols. You must configure a route map with the redistribution to control which routes are passed into RIP. A route policy allows you to filter routes based on attributes such as the destination, origination protocol, route type, route tag, and so on. For more information, see [Configuring Route Policy Manager, on page 397](#).

Whenever you redistribute routes into a RIP routing domain, Cisco NX-OS does not, by default, redistribute the default route into the RIP routing domain. You can generate a default route into RIP, which can be controlled by a route policy.

You also configure the default metric that is used for all imported routes into RIP.

## Load Balancing

You can use load balancing to allow a router to distribute traffic over all the router network ports that are the same distance from the destination address. Load balancing increases the usage of network segments and increases effective network bandwidth.

Cisco NX-OS supports the Equal Cost Multiple Paths (ECMP) feature with up to 16 equal-cost paths in the RIP route table and the unicast RIB. You can configure RIP to load balance traffic across some or all of those paths.

## High Availability for RIP

Cisco NX-OS supports stateless restarts for RIP. After a reboot or supervisor switchover, Cisco NX-OS applies the running configuration and RIP immediately sends request packets to repopulate its routing table.

## Virtualization Support

Cisco NX-OS supports multiple instances of the RIP protocol that run on the same system. RIP supports virtual routing and forwarding (VRF) instances. VRFs exist within virtual device contexts (VDCs).

You can configure up to four RIP instances on a VDC. By default, Cisco NX-OS places you in the default VDC and default VRF unless you specifically configure another VDC and VRF.

See the *Cisco Nexus 7000 Series NX-OS Virtual Device Context Configuration Guide*.

## Prerequisites for RIP

- You must enable RIP.

## Guidelines and Limitations for RIP

- Cisco NX-OS does not support RIPv1. If Cisco NX-OS receives a RIPv1 packet, it logs a message and drops the packet.
- Cisco NX-OS does not establish adjacencies with RIPv1 routers.
- If you are familiar with the Cisco IOS CLI, be aware that the Cisco NX-OS commands for this feature might differ from the Cisco IOS commands that you would use.

## Default Settings for RIP Parameters

Default RIP Parameters

| Parameters                       | Default  |
|----------------------------------|----------|
| Maximum paths for load balancing | 8        |
| RIP feature                      | Disabled |
| Split horizon                    | Enabled  |

## Configuring RIP

### Enabling RIP

#### Before you begin

Confirm that you are in the correct VDC. To change the VDC, use the **switchto vdc** command.

**Procedure**

|               | <b>Command or Action</b>                                             | <b>Purpose</b>                                                                                                                |
|---------------|----------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configure terminal</b>                                    | Enters global configuration mode.                                                                                             |
| <b>Step 2</b> | switch(config)# <b>[no] feature rip</b>                              | Enables the RIP feature.<br><br>Use the <b>no</b> form of this command to disable this feature.                               |
| <b>Step 3</b> | (Optional) switch(config)# <b>copy running-config startup-config</b> | Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration. |
| <b>Step 4</b> | (Optional) switch(config)# <b>show feature</b>                       | Displays enables and disabled features.                                                                                       |

**Example**

The following example enables RIP:

```
switch # configure terminal
switch(config)# feature rip
switch(config)# copy running-config startup-config
```

## Creating a RIP Instance

You can create a RIP instance and configure the address family for that instance.

**Before you begin**

You must enable RIP.

Ensure that you are in the correct VDC (or use the **switchto vdc** command).

**Procedure**

|               | <b>Command or Action</b>                                  | <b>Purpose</b>                                                                                                                                                                                                             |
|---------------|-----------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configure terminal</b>                         | Enters global configuration mode.                                                                                                                                                                                          |
| <b>Step 2</b> | switch(config)# <b>[no] router rip instance-tag</b>       | Creates a new RIP instance with the configured instance-tag.<br><br>Use the <b>no</b> form of this command to disable this feature.<br><br><b>Note</b> You must also remove any RIP commands configured in interface mode. |
| <b>Step 3</b> | switch(config-router)# <b>address-family ipv4 unicast</b> | Configures the address family for this RIP instance and enters address-family configuration mode.                                                                                                                          |

|               | Command or Action                                                                                              | Purpose                                                                                                                                                   |
|---------------|----------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 4</b> | (Optional) switch(config-router-af)# <b>show ip rip</b> [instance <i>instance-tag</i> ] [vrf <i>vrf-name</i> ] | Displays a summary of RIP information for all RIP instances.                                                                                              |
| <b>Step 5</b> | (Optional) switch(config-router-af)# <b>distance</b> <i>value</i>                                              | Sets the administrative distance for RIP, in address-family configuration mode. The range is from 1 to 255.                                               |
| <b>Step 6</b> | (Optional) switch(config-router-af)# <b>maximum-paths</b> <i>number</i>                                        | Configures the maximum number of equal-cost paths that RIP maintains in the route table, in address-family configuration mode. The range is from 1 to 16. |
| <b>Step 7</b> | (Optional) switch(config-router-af)# <b>copy running-config startup-config</b>                                 | Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.                             |

### Example

The following example creates a RIP instance for IPv4 and sets the number of equal-cost paths for load balancing:

```
switch# configure terminal
switch(config)# router rip Enterprise
switch(config-router)# address-family ipv4 unicast
switch(config-router-af)# max-paths 10
switch(config-router-af)# copy running-config startup-config
```

## Restarting a RIP Instance

You can restart a RIP instance and remove all associated neighbors for the instance.

### Procedure

|               | Command or Action                                                    | Purpose                                                                                                                       |
|---------------|----------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configure terminal</b>                                    | Enters global configuration mode.                                                                                             |
| <b>Step 2</b> | switch(config)# <b>restart rip</b> <i>instance-tag</i>               | Restarts the RIP instance and removes all neighbors.                                                                          |
| <b>Step 3</b> | (Optional) switch(config)# <b>copy running-config startup-config</b> | Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration. |

### Example

The following example restarts a RIP instance:



```
switch # configure terminal
switch(config)# restart rip Enterprise
switch(config)# copy running-config startup-config
```

## Configuring RIP on an Interface

### Before you begin

- You must enable RIP.
- Confirm that you are in the correct VDC. To change the VDC, use the **switchto vdc** command.

### Procedure

|               | Command or Action                                                                                                                                                                | Purpose                                                                                                                       |
|---------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configure terminal</b>                                                                                                                                                | Enters global configuration mode.                                                                                             |
| <b>Step 2</b> | switch(config)# <b>interface</b> <i>interface-type slot/port</i>                                                                                                                 | Enters interface configuration mode.                                                                                          |
| <b>Step 3</b> | switch(config-if)# <b>ip router rip</b> <i>instance-tag</i>                                                                                                                      | Associates this interface with a RIP instance.                                                                                |
| <b>Step 4</b> | (Optional) switch(config-if)# <b>copy running-config startup-config</b>                                                                                                          | Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration. |
| <b>Step 5</b> | (Optional) switch(config-if)# <b>show ip rip</b> [ <i>instance instance-tag</i> ] <b>interface</b> [ <i>interface-type slot/port</i> ] [ <i>vrf vrf-name</i> ] [ <b>detail</b> ] | Displays RIP information for an interface.                                                                                    |

### Example

The following example configures RIP on an Ethernet interface:

```
switch # configure terminal
switch(config)# interface ethernet 1/2
switch(config-if)# no switchport
switch(config-if)# ip router rip Enterprise
switch(config-if)# show ip rip Enterprise ethernet 1/2
switch(config-if)# copy running-config startup-config
```

## Configuring RIP Authentication

You can configure authentication for RIP packets on an interface.

### Before you begin

- You must enable RIP.
- Ensure that you are in the correct VDC (or use the **switchto vdc** command).

- Configure a keychain if necessary before enabling authentication. For details about implementing key chains, see the *Cisco Nexus 7000 Series NX-OS Security Configuration Guide*.

### Procedure

|               | Command or Action                                                       | Purpose                                                                                                                       |
|---------------|-------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configure terminal</b>                                       | Enters global configuration mode.                                                                                             |
| <b>Step 2</b> | switch(config)# <b>interface</b> <i>interface-type slot/port</i>        | Enters interface configuration mode.                                                                                          |
| <b>Step 3</b> | switch(config-if)# <b>ip rip authentication mode</b> {text   md5}       | Sets the authentication type for RIP on this interface as cleartext or MD5 authentication digest.                             |
| <b>Step 4</b> | switch(config-if)# <b>ip rip authentication keychain</b> <i>key</i>     | Configures the authentication key used for RIP on this interface.                                                             |
| <b>Step 5</b> | (Optional) switch(config-if)# <b>copy running-config startup-config</b> | Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration. |

### Example

The following example creates a key chain and configures MD5 authentication on a RIP interface:

```
switch# configure terminal
switch(config)# key chain RIPKey
switch(config)# key-string myrip
switch(config)# accept-lifetime 00:00:00 Jan 01 2000 infinite
switch(config)# send-lifetime 00:00:00 Jan 01 2000 infinite
switch(config)# interface ethernet 1/2
switch(config-if)# ip rip authentication mode md5
switch(config-if)# ip rip authentication keychain RIPKey
switch(config-if)# copy running-config startup-config
```

## Configuring a Passive Interface

You can configure a RIP interface to receive routes but not send route updates by setting the interfaces to passive mode. You can configure a RIP interface in passive mode in the interface configuration mode.

### Procedure

|               | Command or Action                                                | Purpose                               |
|---------------|------------------------------------------------------------------|---------------------------------------|
| <b>Step 1</b> | switch# <b>configure terminal</b>                                | Enters global configuration mode.     |
| <b>Step 2</b> | switch(config)# <b>interface</b> <i>interface-type slot/port</i> | Enters interface configuration mode.  |
| <b>Step 3</b> | switch(config-if)# <b>ip rip passive-interface</b>               | Sets the interface into passive mode. |

|               | Command or Action                                                       | Purpose                                                                                                                       |
|---------------|-------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 4</b> | (Optional) switch(config-if)# <b>copy running-config startup-config</b> | Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration. |

### Example

The following example configures a RIP interface in passive mode:

```
switch # configure terminal
switch(config)# interface ethernet 1/2
switch(config-if)# ip rip passive-interface
switch(config-if)# copy running-config startup-config
```

## Configuring Split Horizon with Poison Reverse

You can configure an interface to advertise routes learned by RIP as unreachable over the interface that learned the routes by enabling poison reverse. You can configure split horizon with poison reverse on an interface using the interface configuration mode.

### Procedure

|               | Command or Action                                                       | Purpose                                                                                                                       |
|---------------|-------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configure terminal</b>                                       | Enters global configuration mode.                                                                                             |
| <b>Step 2</b> | switch(config)# <b>interface</b> <i>interface-type slot/port</i>        | Enters interface configuration mode.                                                                                          |
| <b>Step 3</b> | switch(config-if)# <b>ip rip poison-reverse</b>                         | Enables split horizon with poison reverse. Split horizon with poison reverse is disabled by default.                          |
| <b>Step 4</b> | (Optional) switch(config-if)# <b>copy running-config startup-config</b> | Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration. |

### Example

The following example restarts a RIP instance:

```
switch # configure terminal
switch(config)# interface ethernet 1/2
switch(config-if)# ip rip poison-reverse
switch(config-if)# copy running-config startup-config
```

## Configuring Route Summarization

You can create aggregate addresses that are represented in the routing table by a summary address. Cisco NX-OS advertises the summary address metric that is the smallest metric of all the more-specific routes. To configure a summary address on an interface, use the interface configuration mode.

### Procedure

|               | Command or Action                                                          | Purpose                                                                                                                       |
|---------------|----------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configure terminal</b>                                          | Enters global configuration mode.                                                                                             |
| <b>Step 2</b> | switch(config)# <b>interface</b> <i>interface-type slot/port</i>           | Enters interface configuration mode.                                                                                          |
| <b>Step 3</b> | switch(config-if)# <b>ip rip summary-address</b> <i>ip-prefix/mask-len</i> | Configured a summary address for RIP for IPv4 addresses.                                                                      |
| <b>Step 4</b> | (Optional) switch(config-if)# <b>copy running-config startup-config</b>    | Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration. |

### Example

The following example restarts a RIP instance:

```
switch# configure terminal
switch(config)# interface ethernet 1/2
switch(config-if)# ip router rip summary-address 192.0.2.0/24
switch(config-if)# copy running-config startup-config
```

## Configuring Route Redistribution

You can configure RIP to accept routing information from another routing protocol and redistribute that information through the RIP network. Redistributed routes can optionally be assigned a default route.

### Before you begin

- You must enable RIP.
- Ensure that you are in the correct VDC (or use the **switchto vdc** command).
- Configure a route map before configuring redistribution.

### Procedure

|               | Command or Action                                     | Purpose                                                      |
|---------------|-------------------------------------------------------|--------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configure terminal</b>                     | Enters global configuration mode.                            |
| <b>Step 2</b> | switch(config)# <b>router rip</b> <i>instance-tag</i> | Creates a new RIP instance with the configured instance-tag. |

|               | Command or Action                                                                                                                                                                                                              | Purpose                                                                                                                       |
|---------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 3</b> | switch(config-router)# <b>address-family ipv4 unicast</b>                                                                                                                                                                      | Enters address family configuration mode.                                                                                     |
| <b>Step 4</b> | switch(config-router-af)# <b>redistribute</b> { <b>bgp as</b>   <b>direct</b>   <b>eigrp</b>   <b>isis</b>   <b>ospf</b>   <b>ospfv3</b>   <b>rip</b> } <i>instance-tag</i>   <b>static</b> } <b>route-map</b> <i>map-name</i> | Redistributes routes from other protocols into RIP.                                                                           |
| <b>Step 5</b> | (Optional) switch(config-router-af)# <b>default-information originate</b> [ <b>always</b> ] [ <b>route-map</b> <i>map-name</i> ]                                                                                               | Generates a default route into RIP, optionally controlled by a route map.                                                     |
| <b>Step 6</b> | (Optional) switch(config-router-af)# <b>default-metric</b> <i>value</i>                                                                                                                                                        | Sets the default metric for all redistributed routes. The range is from 1 to 15. The default is 1.                            |
| <b>Step 7</b> | (Optional) switch(config-router-af)# <b>copy running-config startup-config</b>                                                                                                                                                 | Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration. |
| <b>Step 8</b> | (Optional) switch(config-router-af)# <b>show ip rip route</b> [ <i>ip-prefix</i> [ <b>longer-prefixes</b>   <b>shorter-prefixes</b> ]] [ <b>vrf</b> <i>vrf-name</i> ] [ <b>summary</b> ]                                       | Shows the routes in RIP.                                                                                                      |

### Example

The following example shows how to redistribute EIGRP into RIP:

```
switch# configure terminal
switch(config)# router rip Enterprise
switch(config-router)# address-family ipv4 unicast
switch(config-router-af)# redistribute eigrp 201 route-map RIPmap
switch(config-router-af)# copy running-config startup-config
```

## Configuring Cisco NX-OS RIP for Compatibility with Cisco IOS RIP

Beginning with Cisco NX-OS Release 6.1, you can configure Cisco NX-OS RIP to behave like Cisco IOS RIP in the way that routes are advertised and processed.

Directly connected routes are treated with cost 1 in Cisco NX-OS RIP and with cost 0 in Cisco IOS RIP. When routes are advertised in Cisco NX-OS RIP, the receiving device adds a minimum cost of +1 to all received routes and installs the routes in its routing table. In Cisco IOS RIP, this cost increment is done on the sending router, and the receiving router installs the routes without any modification. This difference in behavior can cause issues when both Cisco NX-OS and Cisco IOS devices are working together. You can prevent these compatibility issues by configuring Cisco NX-OS RIP to advertise and process routes like Cisco IOS RIP

### Before you begin

- You must enable RIP.
- Ensure that you are in the correct VDC (or use the **switchto vdc** command).

**Procedure**

|               | <b>Command or Action</b>                                                    | <b>Purpose</b>                                                                                                                                                                                                                                                                                                                                                    |
|---------------|-----------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configure terminal</b>                                           | Enters global configuration mode.                                                                                                                                                                                                                                                                                                                                 |
| <b>Step 2</b> | switch(config)# <b>router rip</b> <i>instance-tag</i>                       | Creates a new RIP instance with the configured instance-tag. You can enter 100, 201, or up to 20 alphanumeric characters for the instance tag.                                                                                                                                                                                                                    |
| <b>Step 3</b> | switch(config-router)# [ <b>no</b> ] <b>metric direct 0</b>                 | Configures all directly connected routes with cost 0 instead of the default of cost 1 in order to make Cisco NX-OS RIP compatible with Cisco IOS RIP in the way that routes are advertised and processed.<br><br><b>Note</b> This command must be configured on all Cisco NX-OS devices that are present in any RIP network that also contains Cisco IOS devices. |
| <b>Step 4</b> | (Optional) switch(config-router)# <b>show running-config rip</b>            | Displays the current running RIP configuration.                                                                                                                                                                                                                                                                                                                   |
| <b>Step 5</b> | (Optional) switch(config-router)# <b>copy running-config startup-config</b> | Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration                                                                                                                                                                                                                                      |

**Example**

The following example shows how to disable NX-OS RIP compatibility with Cisco IOS RIP by returning all direct routes from cost 0 to cost 1:

```
switch# configure terminal
switch(config)# router rip 100
switch(config-router)# no metric direct 0
switch(config-router)# show running-config rip
switch(config-router)# copy running-config startup-config
```

## Configuring Virtualization

You can configure multiple RIP instances in each VDC. You can also create multiple VRFs within each VDC and use the same or multiple RIP instances in each VRF. You assign a RIP interface to a VRF.



**Note** Configure all other parameters for an interface after you configure the VRF for an interface. Configuring a VRF for an interface deletes all the configurations for an interface.

**Before you begin**

- You must enable RIP.
- Create the VDCs.

**Procedure**

|                | <b>Command or Action</b>                                                                                                                                 | <b>Purpose</b>                                                                                                                 |
|----------------|----------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b>  | switch# <b>configure terminal</b>                                                                                                                        | Enters global configuration mode.                                                                                              |
| <b>Step 2</b>  | switch(config)# <b>vrf vrf-name</b>                                                                                                                      | Creates a new VRF.                                                                                                             |
| <b>Step 3</b>  | switch(config-vrf)# <b>exit</b>                                                                                                                          | Exits VRF configuration mode.                                                                                                  |
| <b>Step 4</b>  | switch(config)# <b>router rip instance-tag</b>                                                                                                           | Creates a new RIP instance with the configured instance tag.                                                                   |
| <b>Step 5</b>  | switch(config-router)# <b>vrf vrf-name</b>                                                                                                               | Creates a new VRF and enters VRF configuration mode.                                                                           |
| <b>Step 6</b>  | (Optional) switch(config-router-vrf)# <b>address-family ipv4 unicast</b>                                                                                 | (Optional) Configures the VRF address family for this RIP instance.                                                            |
| <b>Step 7</b>  | (Optional) switch(-router-vrf-af)# <b>redistribute {bgp as   direct   {eigrp   isis   ospf   ospfv3   rip} instance-tag   static} route-map map-name</b> | Redistributes routes from other protocols into RIP.<br><br>See <a href="#">Configuring Route Policy Manager, on page 397</a> . |
| <b>Step 8</b>  | switch(config-router-vrf-af)# <b>interface ethernet slot/port</b>                                                                                        | Enters interface configuration mode.                                                                                           |
| <b>Step 9</b>  | switch(config-if)# <b>no switchport</b>                                                                                                                  | Configures the interface as a Layer 3 routed interface.                                                                        |
| <b>Step 10</b> | switch(config-if)# <b>vrf member vrf-name</b>                                                                                                            | Adds this interface to a VRF.                                                                                                  |
| <b>Step 11</b> | switch(config-if)# <b>ip-address ip-prefix/length</b>                                                                                                    | Configures an IP address for this interface. You must perform this step after you assign this interface to a VRF.              |
| <b>Step 12</b> | switch(config-if)# <b>ip router rip instance-tag</b>                                                                                                     | Associates this interface with a RIP instance.                                                                                 |
| <b>Step 13</b> | (Optional) switch(config-if)# <b>copy running-config startup-config</b>                                                                                  | Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.  |
| <b>Step 14</b> | (Optional) switch(config-if)# <b>show ip rip [instance instance-tag] interface [interface-type slot/port] [vrf vrf-name]</b>                             | Displays RIP information for an interface in a VRF.                                                                            |

## Example

The following example shows how to create a VRF and add an interface to the VRF:

```
switch# configure terminal
switch(config)# vrf context RemoteOfficeVRF
switch(config-vrf)# exit
switch(config)# router rip Enterprise
switch(config-router)# vrf RemoteOfficeVRF
switch(config-router-vrf)# address-family ipv4 unicast
switch(config-router-vrf-af)# redistribute eigrp 201 route-map RIPmap
switch(config-router-vrf-af)# interface ethernet 1/2
switch(config-if)# vrf member RemoteOfficeVRF
switch(config-if)# ip address 192.0.2.1/16
switch(config-if)# ip router rip Enterprise
switch(config-if)# copy running-config startup-config
```

## Tuning RIP

You can tune RIP to match your network requirements. RIP uses several timers that determine the frequency of routing updates, the length of time before a route becomes invalid, and other parameters. You can adjust these timers to tune routing protocol performance to better suit your internet work needs.



**Note** You must configure the same values for the RIP timers on all RIP-enabled routers in your network.

### Procedure

|               | Command or Action                                                                                                              | Purpose                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|---------------|--------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | (Optional) switch(config-router-af)# <b>timers</b><br><b>basic</b> <i>update timeout holddown</i><br><i>garbage-collection</i> | <p><b>Note</b> This is set in the address-family configuration mode.</p> <p>Sets the RIP timers in seconds. The parameters are as follows:</p> <ul style="list-style-type: none"> <li>• <i>update</i>—The range is from 5 to any positive integer. The default is 30.</li> <li>• <i>timeout</i>—The time that Cisco NX-OS waits before declaring a route as invalid. If Cisco NX-OS does not receive route update information for this route before the timeout interval ends, Cisco NX-OS declares the route as invalid. The range is from 1 to any positive integer. The default is 180.</li> <li>• <i>holddown</i>—The time during which Cisco NX-OS ignores better route information for an invalid route. The range is from 0 to any positive integer. The default is 180.</li> </ul> |



|               | Command or Action                                                                                                                                              | Purpose                                                                                                                                                                                                                                                          |
|---------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|               |                                                                                                                                                                | <ul style="list-style-type: none"> <li><i>garbage-collection</i>—The time from when Cisco NX-OS marks a route as invalid until Cisco NX-OS removes the route from the routing table. The range is from 1 to any positive integer. The default is 120.</li> </ul> |
| <b>Step 2</b> | switch(config-router-af)# <b>exit</b>                                                                                                                          | Exits address-family configuration mode.                                                                                                                                                                                                                         |
| <b>Step 3</b> | switch(config-router)# <b>exit</b>                                                                                                                             | Exits router configuration mode.                                                                                                                                                                                                                                 |
| <b>Step 4</b> | switch(config)# <b>interface</b> <i>type number</i>                                                                                                            | Enters interface configuration mode.                                                                                                                                                                                                                             |
| <b>Step 5</b> | (Optional) switch(config-if)# <b>ip rip metric-offset</b> <i>value</i>                                                                                         | <p><b>Note</b> This is set in the interface configuration mode.</p> <p>Adds a value to the metric for every router received on this interface. The range is from 1 to 15. The default is 1.</p>                                                                  |
| <b>Step 6</b> | (Optional) switch(config-if)# <b>ip rip route-filter</b> { <b>prefix-list</b> <i>list-name</i>   <b>route-map</b> <i>map-name</i>   [ <b>in</b>   <b>out</b> ] | <p><b>Note</b> This is set in the interface configuration mode.</p> <p>Specifies a route map to filter incoming or outgoing RIP updates.</p>                                                                                                                     |

### Example

The following optional examples show how to tune RIP:

```
switch(config-router-af)# timers basic 40 120 120 100
switch(config-router-af)# exit
switch(config-router)# exit
switch(config)# exit
switch(config)# interface ethernet 1/2
switch(config-if)# ip rip metric-offset 10
switch(config-if)# ip rip route-filter route-map InputMap in
```

## Verifying the RIP Configuration

Use one of the following commands to verify the configuration:

| Command                                                                                                                                    | Purpose                                     |
|--------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------|
| <b>show ip rip instance</b> [ <i>instance-tag</i> ] [ <b>vrf</b> <i>vrf-name</i> ]                                                         | Displays the status for an instance of RIP. |
| <b>show ip rip</b> [ <b>instance</b> <i>instance-tag</i> ] <b>interface</b> <i>slot/port</i> <b>detail</b> [ <b>vrf</b> <i>vrf-name</i> ]  | Displays the RIP status for an interface    |
| <b>show ip rip</b> [ <b>instance</b> <i>instance-tag</i> ] <b>neighbor</b> [ <i>interface-type number</i> ] [ <b>vrf</b> <i>vrf-name</i> ] | Displays the RIP neighbor table             |

| Command                                                                                                                                                                                           | Purpose                                         |
|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------|
| <b>show ip rip</b> [ <i>instance instance-tag</i> ] <b>route</b><br>[ <i>ip-prefix/length</i> [ <b>longer-prefixes</b>   <b>shorter-prefixes</b> ]]<br>[ <b>summary</b> ] [ <b>vrf vrf-name</b> ] | Displays the RIP route table                    |
| <b>show running-configuration rip</b>                                                                                                                                                             | Displays the current running RIP configuration. |

## Displaying RIP Statistics

Use one of the following commands to display RIP statistics:

| Command                                                                                                                                                                                                                                                                            | Purpose                                                                                                   |
|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------|
| <b>show ip rip</b> [ <i>instance instance-tag</i> ] <b>policy statistics</b><br><b>redistribute</b> { <b>bgp as</b>   <b>direct</b>   { <b>eigrp</b>   <b>isis</b>   <b>ospf</b>  <br><b>ospfv3</b>   <b>rip</b> } [ <i>instance-tag</i>   <b>static</b> ] [ <b>vrf vrf-name</b> ] | Displays the RIP policy status.<br>Use the <b>clear ip rip policy</b> command to clear policy statistics. |
| <b>show ip rip</b> [ <i>instance instance-tag</i> ] <b>statistics</b><br><i>interface-type number</i> ] [ <b>vrf vrf-name</b> ]                                                                                                                                                    | Displays the RIP statistics.<br>Use the <b>clear ip rip statistics</b> command to clear RIP statistics.   |

Use the **clear ip rip policy** command to clear policy statistics.

Use the **clear ip rip statistics** command to clear RIP statistics.

## Configuration Examples for RIP

This example creates the Enterprise RIP instance in a VRF and adds Ethernet interface 1/2 to this RIP instance. The example also configures authentication for Ethernet interface 1/2 and redistributes EIGRP into this RIP domain.

```
vrf context NewVRF
!
 feature rip
 router rip Enterprise
 vrf NewVRF
 address-family ip unicast
 redistribute eigrp 201 route-map RIPmap
 max-paths 10
!
interface ethernet 1/2
 vrf NewVRF
 ip address 192.0.2.1/16
 ip router rip Enterprise
 ip rip authentication mode md5
 ip rip authentication keychain RIPKey
 .
```

## Related Documents for RIP

| Related Topic | Document Title                                                                  |
|---------------|---------------------------------------------------------------------------------|
| RIP CLI       | <i>Cisco Nexus 7000 Series NX-OS Unicast Routing Command Reference</i>          |
| VDCs and VRFs | <i>Cisco Nexus 7000 Series NX-OS Virtual Device Context Configuration Guide</i> |

## Standards for RIP

No new or modified standards are supported by this feature, and support for existing standards has not been modified by this feature.

## Feature History for RIP

This table includes only the updates for those releases that have resulted in additions or changes to the feature.

| Feature Name | Releases | Feature Information                                                                                                                                  |
|--------------|----------|------------------------------------------------------------------------------------------------------------------------------------------------------|
| RIP          | 6.1(1)   | Added the ability to configure Cisco NX-OS RIP to be behaviorally compatible with Cisco IOS RIP in the way that routes are advertised and processed. |
| RIP          | 4.0(1)   | This feature was introduced.                                                                                                                         |





# CHAPTER 14

## Configuring Static Routing

---

This chapter contains the following sections:

- [Finding Feature Information, on page 355](#)
- [Information About Static Routing, on page 355](#)
- [Prerequisites for Static Routing, on page 358](#)
- [Guidelines and Limitations for Static Routing, on page 358](#)
- [Default Settings for Static Routing Parameters, on page 358](#)
- [Configuring Static Routing, on page 358](#)
- [Verifying the Static Routing Configuration, on page 364](#)
- [Related Documents for Static Routing, on page 364](#)
- [Feature History for Static Routing, on page 364](#)

### Finding Feature Information

Your software release might not support all the features documented in this module. For the latest caveats and feature information, see the Bug Search Tool at <https://tools.cisco.com/bugsearch/> and the release notes for your software release. To find information about the features documented in this module, and to see a list of the releases in which each feature is supported, see the "New and Changed Information" chapter or the Feature History table in this chapter.

### Information About Static Routing

Routers forward packets using either route information from route table entries that you manually configure or the route information that is calculated using dynamic routing algorithms.

Static routes, which define explicit paths between two routers, cannot be automatically updated; you must manually reconfigure static routes when network changes occur. Static routes use less bandwidth than dynamic routes. No CPU cycles are used to calculate and analyze routing updates.

You can supplement dynamic routes with static routes where appropriate. You can redistribute static routes into dynamic routing algorithms but you cannot redistribute routing information calculated by dynamic routing algorithms into the static routing table.

You should use static routes in environments where network traffic is predictable and where the network design is simple. You should not use static routes in large, constantly changing networks because static routes cannot react to network changes. Most networks use dynamic routes to communicate between routers but

might have one or two static routes configured for special cases. Static routes are also useful for specifying a gateway of last resort (a default router to which all unroutable packets are sent).

## Administrative Distance

An administrative distance is the metric used by routers to choose the best path when there are two or more routes to the same destination from two different routing protocols. An administrative distance guides the selection of one routing protocol (or static route) over another, when more than one protocol adds the same route to the unicast routing table. Each routing protocol is prioritized in order of most to least reliable using an administrative distance value.

Static routes have a default administrative distance of 1. A router prefers a static route to a dynamic route because the router considers a route with a low number to be the shortest. If you want a dynamic route to override a static route, you can specify an administrative distance for the static route. For example, if you have two dynamic routes with an administrative distance of 120, you would specify an administrative distance that is greater than 120 for the static route if you want the dynamic route to override the static route.

## Directly Connected Static Routes

You must specify only the output interface (the interface on which all packets are sent to the destination network) in a directly connected static route. The router assumes the destination is directly attached to the output interface and the packet destination is used as the next-hop address. The next hop can be an interface, only for point-to-point interfaces. For broadcast interfaces, the next hop must be an IPv4 or IPv6 address.

## Fully Specified Static Routes

You must specify either the output interface (the interface on which all packets are sent to the destination network) or the next-hop address in a fully specified static route. You can use a fully specified static route when the output interface is a multi-access interface and you need to identify the next-hop address. The next-hop address must be directly attached to the specified output interface.

## Floating Static Routes

A floating static route is a static route that the router uses to back up a dynamic route. You must configure a floating static route with a higher administrative distance than the dynamic route that it backs up. In this instance, the router prefers a dynamic route to a floating static route. You can use a floating static route as a replacement if the dynamic route is lost.



---

**Note** By default, a router prefers a static route to a dynamic route because a static route has a smaller administrative distance than a dynamic route.

---

## Remote Next-Hops for Static Routes

You can specify the next-hop address of a neighboring router which is not directly connected to the router for static routes with remote (non-directly attached) next-hops. If a static route has remote next-hops during

data-forwarding, the next-hops are recursively used in the unicast routing table to identify the corresponding directly attached next-hop(s) that have reachability to the remote next-hops.

## Reliable Static Routing Backup Using Object Tracking Deployment

You can configure Cisco NX-OS to initiate a backup connection from an alternative port if the circuit to the primary gateway is interrupted. You can ensure reliable deployment backups in the case of certain catastrophic events, such as an Internet circuit failure or peer device failure.

Reliable static routing backup using object tracking can determine the state of the primary connection without having to enable a dynamic routing protocol. It also provides a reliable backup solution that can be used for critical circuits that must not go down without automatically engaging a backup circuit.

In a typical scenario, the primary interface of the remote router forwards traffic from the remote LAN to the main office. If the router loses the connection to the main office, the status of the tracked object changes from up to down. When this change occurs, the router removes the routing table entry for the primary interface and installs the preconfigured floating static route on the secondary interface. The router's secondary interface then forwards traffic to the preconfigured destination. The backup circuit can be configured to use the Internet. When the state of the tracked object changes from down to up, the router reinstalls the routing table entry for the primary interface and removes the floating static route for the secondary interface.

### IP Service Level Agreements

This feature uses IP service level agreements (IP SLAs), a network monitoring feature set, to generate ICMP pings to monitor the state of the connection to the primary gateway. An IP SLA is configured to ping a target, such as a publicly routable IP address or a target inside the corporate network. The pings are routed from the primary interface only. A track object is created to monitor the status of the IP SLA configuration. The track object informs the client, the static route, if a state change occurs. The preconfigured floating static route on the secondary interface is installed when the state changes from up to down.



---

**Note** User Datagram Protocol (UDP) echo, or any other protocol supported by IP SLAs, can be used instead of ICMP pings.

---

For more information on IP SLAs, see the *Cisco Nexus 7000 Series NX-OS IP SLAs Configuration Guide*.

### BFD

This feature supports bidirectional forwarding detection (BFD). BFD is a detection protocol designed to provide fast forwarding-path failure detection times. BFD provides subsecond failure detection between two adjacent devices and can be less CPU-intensive than protocol hello messages because some of the BFD load can be distributed onto the data plane on supported modules. See the , for more information.

### Virtualization Support

Static routes support virtual routing and forwarding (VRF) instances. VRFs exist within virtual device contexts (VDCs). By default, Cisco NX-OS places you in the default VDC and default VRF unless you specifically configure another VDC and VRF. For more information, see the *Cisco Nexus 7000 Series NX-OS Virtual Device Context Configuration Guide*.

## Prerequisites for Static Routing

If the next-hop address for a static route is unreachable, the static route will not be added to the unicast routing table.

## Guidelines and Limitations for Static Routing

- You can specify an interface as the next-hop address for a static route only for point-to-point interfaces such as generic routing encapsulation (GRE) tunnels.
- Starting from Cisco NX-OS Release 8.2(4), static IPv6 route with next-hop as the VxLAN route is supported.
- The forward referencing of static routes is not supported for track objects.
- If you are familiar with the Cisco IOS CLI, be aware that the Cisco NX-OS commands for this feature might differ from the Cisco IOS commands that you would use.

## Default Settings for Static Routing Parameters

### Default Static Routing Parameters

| Parameters              | Default  |
|-------------------------|----------|
| Administrative distance | 1        |
| RIP feature             | Disabled |

## Configuring Static Routing

### Configuring a Static Route for IPv4

#### Before you begin

Confirm that you are in the correct VDC. To change the VDC, use the **switchto vdc** command.

#### Procedure

|               | Command or Action                                                                                                     | Purpose                                                                                                                                    |
|---------------|-----------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configure terminal</b>                                                                                     | Enters global configuration mode.                                                                                                          |
| <b>Step 2</b> | switch(config)# <b>ip route</b> { <i>ip-prefix</i>   <i>ip-addr/ip-mask</i> } [{ <i>next-hop</i>   <i>nh-prefix</i> ] | Configures a static route and the interface for this static route. Use ? to display a list of supported interfaces. You can specify a null |



|               | Command or Action                                                          | Purpose                                                                                                                                                                            |
|---------------|----------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|               | <code>[interface next-hop   nh-prefix] [tag tag-value] [pref]</code>       | interface by using null 0. You can optionally configure the next-hop address. The preference value sets the administrative distance. The range is from 1 to 255. The default is 1. |
| <b>Step 3</b> | (Optional) <code>switch(config)# copy running-config startup-config</code> | Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.                                                      |
| <b>Step 4</b> | (Optional) <code>switch(config)# show ip static-route</code>               | Displays information about static routes.                                                                                                                                          |

### Configuration Example

Configuring a Static Route for a null interface.

```
switch# configure terminal
switch(config)# ip static-route 1.1.1.1/32 null 0
switch(config)# copy running-config startup-config
```

Use the **no ip static-route** command to remove the static route.

## Configuring a Static Route for IPv6

### Before you begin

Confirm that you are in the correct VDC. To change the VDC, use the **switchto vdc** command.

### Procedure

|               | Command or Action                                                                                                                                                                             | Purpose                                                                                                                                                                                                                                                                                                                       |
|---------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | <code>switch# configure terminal</code>                                                                                                                                                       | Enters global configuration mode.                                                                                                                                                                                                                                                                                             |
| <b>Step 2</b> | <code>switch(config)# ipv6 route ip6-prefix {nh-prefix   link-local-nh-prefix}   {nh-prefix [interface]   link-local-nh-prefix [interface]} [name nexthop-name] [tag tag-value] [pref]</code> | Configures a static route and the interface for this static route. Use ? to display a list of supported interfaces. You can specify a null interface by using null 0. You can optionally configure the next-hop address. The preference value sets the administrative distance. The range is from 1 to 255. The default is 1. |
| <b>Step 3</b> | (Optional) <code>switch(config)# copy running-config startup-config</code>                                                                                                                    | Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.                                                                                                                                                                                                 |
| <b>Step 4</b> | (Optional) <code>switch(config)# show ipv6 static-route</code>                                                                                                                                | Displays information about static routes.                                                                                                                                                                                                                                                                                     |

**Example**

The following example configures a static route for IPv6:

```
switch# configure terminal
switch(config)# ipv6 route 2001:0DB8::/48 6::6 null 0
```

**Configuring a Static Route over a VLAN**

You can configure a static route without next hop support over a VLAN, also known as a switch virtual switch (SVI).

**Before you begin**

Ensure that the access port is part of the VLAN.

**Procedure**

|               | <b>Command or Action</b>                                                                                      | <b>Purpose</b>                                                                                                                                                                                               |
|---------------|---------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configure terminal</b>                                                                             | Enters global configuration mode.                                                                                                                                                                            |
| <b>Step 2</b> | switch(config)# <b>feature interface-vlan</b>                                                                 | Enables VLAN interface mode.                                                                                                                                                                                 |
| <b>Step 3</b> | switch(config)# <b>interface vlan</b> <i>vlan-id</i>                                                          | Creates a switch virtual interface (SVI) and enters interface configuration mode.<br><br>The range for the <i>vlan-id</i> argument is from 1 to 4094, except for the VLANs reserved for the internal switch. |
| <b>Step 4</b> | switch(config-if)# <b>ip address</b> <i>ip-addr/length</i>                                                    | Configures an IP address for the VLAN.                                                                                                                                                                       |
| <b>Step 5</b> | switch(config-if)# <b>ip route</b> <i>ip-addr/length</i> <i>vlan-id</i>                                       | Adds an interface static route without a next hop on the SVI.<br><br>The IP address is the address that is configured on the interface that is connected to the switch.                                      |
| <b>Step 6</b> | (Optional) switch(config-if)# <b>ip route</b> <i>ip-addr/length</i> <i>vlan-id</i> <i>next-hop-ip-address</i> | Configures explicit next hop address when you set up a /32 static route over an interface VLAN.<br><br>The IP address is the address that is configured on the interface that is connected to the switch.    |
| <b>Step 7</b> | (Optional) switch(config-if)# <b>show ip route</b>                                                            | Displays routes from the Unicast Route Information Base (URIB).                                                                                                                                              |
| <b>Step 8</b> | (Optional) switch(config)# <b>copy running-config startup-config</b>                                          | Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.                                                                                |

### Example

This example shows how to configure a static route without a next hop over an SVI:

```
switch# configure terminal
switch(config)# feature interface-vlan
switch(config)# interface vlan 10
switch(config-if)# ip address 192.0.2.1/8
switch(config-if)# ip route 209.165.200.224/27 vlan 10 <===209,165.200.224 is the IP
address of the interface that is configured on the interface that is directly connected to
the switch.
switch(config-if)# copy running-config startup-config
```



---

**Note** When you set up a /32 static route over an interface VLAN, you have to configure an explicit next hop by using the **ip route ip-addr/length vlan-id next-hop-ip-address** command.

---

This example shows how to configure an explicit next hop when you set up a /32 static route over an interface VLAN:

```
switch# configure terminal
switch(config)# feature interface-vlan
switch(config)# interface vlan 10
switch(config-if)# ip address 209.165.202.128/27
switch(config-if)# ip route 209.165.202.130/32 vlan 10 209.165.202.130
switch(config-if)# copy running-config startup-config
```

### What to do next

Use the **no ip static-route** command to remove the static route.

## Configuring Reliable Static Routing Backup Using Object Tracking

You can configure Cisco NX-OS to use Internet Control Message Protocol (ICMP) pings to identify when a connection goes down and initiate a backup connection from any alternative port.

### Before you begin

- Configure both a primary interface and a backup interface to used for reliable static routing backup.
- Configure an IP SLA with policy-based routing object tracking to be used for reliable static routing backup.
- Configure a routing policy for static routing to be used for reliable static routing backup.
- Create a track object to be associated with the static route using the **track object-id interface** command
- Ensure that you are in the correct VDC (or use the **switchto vdc** command).



---

**Note** If you attempt to configure a static route associated with a track object before you create the track object, the static route command is not accepted by the switch.

---

**Procedure**

|               | <b>Command or Action</b>                                                                                             | <b>Purpose</b>                                                                                                                                                                                                                                                                                                                                                 |
|---------------|----------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configure terminal</b>                                                                                    | Enters global configuration mode.                                                                                                                                                                                                                                                                                                                              |
| <b>Step 2</b> | switch(config)# <b>{ip   ipv6} route ip-prefix ip-mask ip-addr track object-number</b>                               | Configures a static route associated with the track object. The object-number argument specifies that the static route is installed only if the configured track object is up.                                                                                                                                                                                 |
| <b>Step 3</b> | switch(config)# <b>show {ip   ipv6} static-route track-table</b>                                                     | Displays information about the IPv4 or IPv6 static-route track table.                                                                                                                                                                                                                                                                                          |
| <b>Step 4</b> | switch(config)# <b>show track track-number</b>                                                                       | Displays information about a specific tracked object.                                                                                                                                                                                                                                                                                                          |
| <b>Step 5</b> | switch(config)# <b>{ip   ipv6} route network-number network-mask {ip-address   interface} [distance] [name name]</b> | Configures a floating IPv4 or IPv6 static route on the secondary interface.<br><br>The network prefix and mask length must be the same as the static route previously configured for the primary interface associated with a track object. The floating static route should have a higher value of preference than the route associated with the track object. |

## Configuring Virtualization for IPv4

**Before you begin**

Confirm that you are in the correct VDC. To change the VDC, use the **switchto vdc** command.

**Procedure**

|               | <b>Command or Action</b>                                                                                                                                                 | <b>Purpose</b>                                                                                                                                                                                                                                                                                                                |
|---------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configure terminal</b>                                                                                                                                        | Enters global configuration mode.                                                                                                                                                                                                                                                                                             |
| <b>Step 2</b> | switch(config)# <b>vrf context vrf-name</b>                                                                                                                              | Creates a VRF and enters VRF configuration mode.                                                                                                                                                                                                                                                                              |
| <b>Step 3</b> | switch(config-vrf)# <b>ip route {ip-prefix   ip-addr ip-mask} {next-hop   nh-prefix   interface [sub-intf-separator sub-intf-num ] next-hop } [tag tag-value] [pref]</b> | Configures a static route and the interface for this static route. Use ? to display a list of supported interfaces. You can specify a null interface by using null 0. You can optionally configure the next-hop address. The preference value sets the administrative distance. The range is from 1 to 255. The default is 1. |
| <b>Step 4</b> | (Optional) switch(config-vrf)# <b>copy running-config startup-config</b>                                                                                                 | Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.                                                                                                                                                                                                 |

|               | Command or Action                                                       | Purpose                                |
|---------------|-------------------------------------------------------------------------|----------------------------------------|
| <b>Step 5</b> | (Optional) switch(config-vrf)# <b>show ip static-route vrf vrf-name</b> | Displays information on static routes. |

### Example

The following example configures VRF for IPv4.

```
switch # configure terminal
switch(config)# vrf context StaticVrf
switch(config-vrf)# ip route 192.0.2.0/8 ethernet 1/2 10.0.0.2
switch(config-vrf)# show running-config startup-config
```

## Configuring Virtualization for IPv6

### Before you begin

Confirm that you are in the correct VDC. To change the VDC, use the **switchto vdc** command.

### Procedure

|               | Command or Action                                                                                                                                                                                                         | Purpose                                                                                                                                                                                                                                                                                                                       |
|---------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configure terminal</b>                                                                                                                                                                                         | Enters global configuration mode.                                                                                                                                                                                                                                                                                             |
| <b>Step 2</b> | switch(config)# <b>vrf context vrf-name</b>                                                                                                                                                                               | Creates a VRF and enters VRF configuration mode.                                                                                                                                                                                                                                                                              |
| <b>Step 3</b> | switch(config-vrf)# <b>ipv6 route ip6-prefix {nh-prefix   link-local-nh-prefix }   {next-hop   link-local-net-hop   interface [sub-intf-separator sub-intf-num] next-hop } [name nexthop-name] [tag tag-value] [pref]</b> | Configures a static route and the interface for this static route. Use ? to display a list of supported interfaces. You can specify a null interface by using null 0. You can optionally configure the next-hop address. The preference value sets the administrative distance. The range is from 1 to 255. The default is 1. |
| <b>Step 4</b> | (Optional) switch(config-vrf)# <b>copy running-config startup-config</b>                                                                                                                                                  | Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.                                                                                                                                                                                                 |
| <b>Step 5</b> | (Optional) switch(config-vrf)# <b>show ipv6 static-route vrf vrf-name</b>                                                                                                                                                 | Displays information on static routes.                                                                                                                                                                                                                                                                                        |

### Example

The following example configures virtualization for IPv6:

```
switch # configure terminal
switch(config)# vrf context StaticVrf
switch(config-vrf)# ipv6 route 2001:0DB8::/48 6::6 ethernet 2/1 2b11::2f01:4c
switch(config-vrf)# copy running-config startup-config
```

## Verifying the Static Routing Configuration

Use one of the following commands to verify the configuration:

| Command                                                | Purpose                                                               |
|--------------------------------------------------------|-----------------------------------------------------------------------|
| <code>show {ip   ipv6} static-route</code>             | Displays the configured static routes.                                |
| <code>show ipv6 static-route vrf vrf-name</code>       | Displays static route information for each VRF.                       |
| <code>show {ip   ipv6} static-route track-table</code> | Displays information about the IPv4 or IPv6 static-route track table. |
| <code>show track track-number</code>                   | Displays information about a specific tracked object.                 |

## Related Documents for Static Routing

| Related Topic      | Document Title                                                                  |
|--------------------|---------------------------------------------------------------------------------|
| Static Routing CLI | <i>Cisco Nexus 7000 Series NX-OS Unicast Routing Command Reference</i>          |
| VDCs               | <i>Cisco Nexus 7000 Series NX-OS Virtual Device Context Configuration Guide</i> |

## Feature History for Static Routing

This table includes only the updates for those releases that have resulted in additions or changes to the feature.

| Feature Name                                         | Releases | Feature Information                                                   |
|------------------------------------------------------|----------|-----------------------------------------------------------------------|
| Static IPv6 Route                                    | 8.2(4)   | Added support for static IPv6 route with next-hop as the VxLAN route. |
| Static Route over VLAN                               | 6.2(2a)  | This feature was introduced.                                          |
| Reliable static routing backup using object tracking | 6.2(2)   | This feature was introduced.                                          |
| Static routing                                       | 6.0(1)   | Updated for F2 Series modules.                                        |
| Layer 3 routing using a mixed chassis                | 5.1(1)   | This feature was introduced.                                          |
| Static routing                                       | 5.1(1)   | Added the <b>name</b> option to the <b>ip route</b> command.          |

| Feature Name   | Releases | Feature Information                                                                                                        |
|----------------|----------|----------------------------------------------------------------------------------------------------------------------------|
| BFD            | 5.0(2)   | Added support for BFD. See the <i>Cisco Nexus 7000 Series NX-OS Interfaces Configuration Guide</i> , for more information. |
| Static routing | 4.0(1)   | This feature was introduced.                                                                                               |







## CHAPTER 15

# Configuring the Interoperability of Modules for Unicast Routing

---

This chapter contains the following sections:

- [Finding Feature Information, on page 367](#)
- [Configuring the Interoperability of Modules for Unicast Routing, on page 367](#)
- [Information About the Interoperability of Modules for Unicast Routing, on page 368](#)
- [Guidelines and Limitations for the Interoperability of Modules for Unicast Routing, on page 368](#)
- [Configuring the Interoperability of Modules for Unicast Routing, on page 368](#)
- [Verifying the Configuration for the Interoperability of Modules for Unicast Routing, on page 369](#)
- [Configuration Examples for the Interoperability of Modules for Unicast Routing, on page 369](#)
- [Related Documents for the Interoperability of Modules for Unicast Routing , on page 370](#)
- [Feature History for the Interoperability of Modules for Unicast Routing, on page 370](#)

## Finding Feature Information

Your software release might not support all the features documented in this module. For the latest caveats and feature information, see the Bug Search Tool at <https://tools.cisco.com/bugsearch/> and the release notes for your software release. To find information about the features documented in this module, and to see a list of the releases in which each feature is supported, see the "New and Changed Information" chapter or the Feature History table in this chapter.

## Configuring the Interoperability of Modules for Unicast Routing

This chapter describes how to configure the interoperability of F1 Series modules with M Series modules for unicast routing on the Cisco NX-OS device.

# Information About the Interoperability of Modules for Unicast Routing

A mixed chassis is a Cisco Nexus 7000 Series chassis that contains at least one F1 Series module and at least one M Series module. Because the F1 Series module processes only Layer 2 traffic, you must configure it to pass Layer 3 traffic through the chassis.

## Guidelines and Limitations for the Interoperability of Modules for Unicast Routing

The interoperability of modules for unicast routing has the following configuration guidelines and limitations:

- You cannot use F2, F2e, or F3 Series modules in the Cisco Nexus 7000 Series chassis to perform proxy Layer 3 routing for F1 series modules.
- To support the coexistence of an F2e Series module with an M Series module in the same VDC, the F2e Series module operates in a proxy mode so that all Layer 3 traffic is sent to an M Series module in the same VDC. For F2e proxy mode, having routing adjacencies connected through F2e interfaces with an M1 Series module is not supported. However, routing adjacencies connected through F2e interfaces with an M2 Series module is supported.

## Configuring the Interoperability of Modules for Unicast Routing

To configure a Layer 3 gateway in a mixed chassis, you use the proxy routing functionality. You enable routing on a specific VLAN by configuring a VLAN interface, and the system automatically provides load-balanced routing functionality. See the *Cisco Nexus 7000 Series NX-OS Interfaces Configuration Guide* for more information about Layer 3 routing and VLAN interfaces.

For interoperability between F1 Series and M Series modules, use the following procedure to specify which physical interfaces on the M Series modules you want to use for Layer 3 routing.

### Before you begin

You must configure a VLAN interface for each VLAN on the F1 Series module that you want to use with the proxy-routing functionality in a mixed chassis.

You must have interfaces from both the M Series modules and the F1 Series modules in the same VDC.

### Procedure

|               | Command or Action                                                                                                                | Purpose                                                                                                                          |
|---------------|----------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configure terminal</b>                                                                                                | Enters global configuration mode.                                                                                                |
| <b>Step 2</b> | switch(config)# <b>hardware proxy layer-3 routing {use   exclude} {module mod-number   interface slot/port} [module-type f1]</b> | Configures specific modules and physical interfaces on the M Series module to provide the proxy routing on the F1 Series module. |

|               | Command or Action                                                    | Purpose                                                                                                                       |
|---------------|----------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 3</b> | (Optional) switch(config)# <b>show hardware proxy layer-3 detail</b> | Displays information about the proxy Layer 3 functionality.                                                                   |
| <b>Step 4</b> | (Optional) switch(config)# <b>copy running-config startup-config</b> | Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration. |

## Verifying the Configuration for the Interoperability of Modules for Unicast Routing

To display the interoperability of modules for unicast routing configuration, perform one of the following tasks:

| Command                                                      | Purpose                                                                                                                                                                                                         |
|--------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>show hardware proxy layer-3 counters {brief   detail}</b> | Displays the number of packets sent by F1 Series modules to each M Series module for proxy forwarding.<br><br><b>Note</b> Enter the <b>clear hardware proxy layer-3 counters</b> command to clear the counters. |
| <b>show hardware proxy layer-3 detail</b>                    | Displays information about proxy routing from an F1 Series module to an M Series module in a chassis that contains both types of modules.                                                                       |

## Configuration Examples for the Interoperability of Modules for Unicast Routing

This example shows how to specify physical interfaces on M Series modules to perform proxy routing on F1 Series modules in a mixed chassis:

```
switch# configure terminal
switch(config)# hardware proxy layer-3 routing use module 1, 7
switch(config)# show hardware proxy layer-3 detail
```

## Related Documents for the Interoperability of Modules for Unicast Routing

| Related Topic                                       | Document Title                                                                  |
|-----------------------------------------------------|---------------------------------------------------------------------------------|
| Interoperability of modules for unicast routing CLI | <i>Cisco Nexus 7000 Series NX-OS Unicast Routing Command Reference</i>          |
| VDCs                                                | <i>Cisco Nexus 7000 Series NX-OS Virtual Device Context Configuration Guide</i> |

## Feature History for the Interoperability of Modules for Unicast Routing

This table includes only the updates for those releases that have resulted in additions or changes to the feature.

**Table 27: Feature History for the Interoperability of Modules for Unicast Routing**

| Feature Name                                    | Release | Feature Information                  |
|-------------------------------------------------|---------|--------------------------------------|
| Interoperability of modules for unicast routing | 6.1(1)  | Added support for M2 Series modules. |
| Interoperability of modules for unicast routing | 5.1(1)  | This feature was introduced.         |



## CHAPTER 16

# Configuring Layer 3 Virtualization

This chapter contains the following sections:

- [Finding Feature Information, on page 371](#)
- [Information About Layer 3 Virtualization, on page 371](#)
- [Guidelines and Limitations for VRF, on page 375](#)
- [Default Settings for VRF, on page 376](#)
- [Configuring VRFs, on page 376](#)
- [Verifying the VRF Configuration, on page 380](#)
- [Configuration Examples for VRF, on page 380](#)
- [Related Documents for VRF, on page 381](#)
- [Standards for VRF, on page 382](#)
- [Feature History for VRF, on page 382](#)

## Finding Feature Information

Your software release might not support all the features documented in this module. For the latest caveats and feature information, see the Bug Search Tool at <https://tools.cisco.com/bugsearch/> and the release notes for your software release. To find information about the features documented in this module, and to see a list of the releases in which each feature is supported, see the "New and Changed Information" chapter or the Feature History table in this chapter.


## Information About Layer 3 Virtualization

Cisco NX-OS supports a hierarchy of virtualization that can divide the physical system resources into multiple virtual device contexts (VDCs). Each VDC acts as a standalone device with both Layer 2 and Layer 3 services available. You can configure up to 4 VDCs, including the default VDC. See the *Cisco Nexus 7000 Series NX-OS Virtual Device Context Configuration Guide, Release 5.x*, for more information on VDCs.

Cisco NX-OS further virtualizes each VDC to support virtual routing and forwarding instances (VRFs). You can configure multiple VRFs in a VDC. Each VRF contains a separate address space with unicast and multicast route tables for IPv4 and IPv6 and makes routing decisions independent of any other VRF.

The figure shows multiple independent VRFs in two different VDCs.

Figure 39: Multiple VRFs in VDCs



A VRF name is local to a VDC, so you can configure two VRFs with the same name if the VRFs exist in different VDCs. In Figure 14-1, VRF A in VDC 2 is independent of VRF B and VRF A in VDC n.

Each router has a default VRF and a management VRF. All Layer 3 interfaces and routing protocols exist in the default VRF until you assign them to another VRF. The mgmt0 interface exists in the management VRF and is shared among multiple VDCs. Each VDC has a unique IP address for the mgmt0 interface (see the *Cisco Nexus 7000 Series NX-OS Fundamentals Configuration Guide, Release 6.x*).

#### Management VRF

- The management VRF is for management purposes only.
- Only the mgmt 0 interface can be in the management VRF.
- The mgmt 0 interface cannot be assigned to another VRF.
- The mgmt 0 interface is shared among multiple VDCs.
- No routing protocols can run in the management VRF (static only).

#### Default VRF


- All Layer 3 interfaces exist in the default VRF until they are assigned to another VRF.
- Routing protocols run in the default VRF context unless another VRF context is specified.
- The default VRF uses the default routing context for all show commands.
- The default VRF is similar to the global routing table concept in Cisco IOS.

## VRF and Routing

All unicast and multicast routing protocols support VRFs. When you configure a routing protocol in a VRF, you set routing parameters for the VRF that are independent of routing parameters in another VRF for the same routing protocol instance.

You can assign interfaces and route protocols to a VRF to create virtual Layer 3 networks. An interface exists in only one VRF. Figure 9-1 shows one physical network split into two virtual networks with two VRFs. Routers Z, A, and B exist in VRF Red and form one address domain. These routers share route updates that do not include router C because router C is configured in a different VRF.

**Figure 40: VRFs in a Network**



By default, Cisco NX-OS uses the VRF of the incoming interface to select which routing table to use for a route lookup. You can configure a route policy to modify this behavior and set the VRF that Cisco NX-OS uses for incoming packets.




---

**Note** Do not use the **export map** command in the VRF mode for prefix filtering. When a route-target export is configured, all routes are exported and then imported to VRFs with a matching route-target import. In this case, the export map does not filter routes, but it can be used to set attributes for the selected routes. If you need to export only the selected routes, remove the route-target export and use the export map to filter routes; and set the route-target-ext-community so that the VRFs with the matching route-target import imports these routes.

---

## VRF-Aware Services

A fundamental feature of the Cisco NX-OS architecture is that every IP-based feature is VRF aware.

The following VRF-aware services can select a particular VRF to reach a remote server or to filter information based on the selected VRF:

- AAA
- Call Home
- DNS
- GLBP
- HSRP
- HTTP
- NetFlow
- NTP

- RADIUS
- Ping and Traceroute
- SSH
- SNMP
- Syslog
- TACACS+
- TFTP
- VRRP
- XML


See the appropriate configuration guide for each service for more information on configuring VRF support in that service.

## Reachability

Reachability indicates which VRF contains the routing information necessary to get to the server providing the service. For example, you can configure an SNMP server that is reachable on the management VRF. When you configure that server address on the router, you also configure which VRF that Cisco NX-OS must use to reach the server.

The figure shows an SNMP server that is reachable over the management VRF. You configure router A to use the management VRF for SNMP server host 192.0.2.1.

**Figure 41: Service VRF Reachability**




## Filtering

Filtering allows you to limit the type of information that goes to a VRF-aware service based on the VRF. For example, you can configure a syslog server to support a particular VRF. The figure shows two syslog servers with each server supporting one VRF. syslog server A is configured in VRF Red, so Cisco NX-OS sends only system messages generated in VRF Red to syslog server A.



Figure 42: Service VRF Filtering




## Combining Reachability and Filtering

You can combine reachability and filtering for VRF-aware services. You configure the VRF that Cisco NX-OS uses to connect to that service as well as the VRF that the service supports. If you configure a service in the default VRF, you can optionally configure the service to support all VRFs.

The figure shows an SNMP server that is reachable on the management VRF. You can configure the SNMP server to support only the SNMP notifications from VRF Red, for example.

Figure 43: Service VRF Reachability Filtering



## Guidelines and Limitations for VRF

- To completely disable selective VRF download in F3 modules in all VDCs, use the **no hardware forwarding selective-vrf** command in global configuration mode. You must reload the device after applying this command.
- When you make an interface a member of an existing VRF, Cisco NX-OS removes all Layer 3 configurations. You should configure all Layer 3 parameters after adding an interface to a VRF.
- You should add the mgmt0 interface to the management VRF and configure the mgmt0 IP address and other parameters after you add it to the management VRF.
- If you configure an interface for a VRF before the VRF exists, the interface is operationally down until you create the VRF.
- Cisco NX-OS creates the default and management VRFs by default. You should make the mgmt0 interface a member of the management VRF.
- The **write erase boot** command does not remove the management VRF configurations. You must use the **write erase** command and then the **write erase boot** command.
- If you are familiar with the Cisco IOS CLI, be aware that the Cisco NX-OS commands for this feature might differ from the Cisco IOS commands that you would use.

## Default Settings for VRF

| Parameters      | Default             |
|-----------------|---------------------|
| Configured VRFs | Default, management |
| routing context | Default VRF         |

## Configuring VRFs

### Creating a VRF

Commands available in global configuration mode are also available in VRF configuration mode.

#### Before you begin

Ensure that you are in the correct VDC (or use the `switchto vdc` command).

#### Procedure

|               | Command or Action                                                                                                                                                    | Purpose                                                                                                                                                                                                                    |
|---------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | <code>switch# configure terminal</code>                                                                                                                              | Enters global configuration mode.                                                                                                                                                                                          |
| <b>Step 2</b> | <code>switch(config)# vrf context name</code>                                                                                                                        | Creates a new VRF and enters VRF configuration mode. The name can be any case-sensitive, alphanumeric string up to 32 characters.                                                                                          |
| <b>Step 3</b> | (Optional) <code>switch(config-vrf)# ip route {ip-prefix   ip-addr ip-mask} {[next-hop   nh-prefix]   [interface next-hop   nh-prefix]} [tag tag-value [pref]</code> | Configures a static route and the interface for this static route. You can optionally configure the next-hop address. The preference value sets the administrative distance. The range is from 1 to 255. The default is 1. |
| <b>Step 4</b> | (Optional) <code>switch(config-vrf)# show vrf [vrf-name]</code>                                                                                                      | Displays VRF information.                                                                                                                                                                                                  |
| <b>Step 5</b> | <code>switch(config-vrf)# exit</code>                                                                                                                                | Exits the current configuration mode.                                                                                                                                                                                      |
| <b>Step 6</b> | (Optional) <code>switch(config)# copy running-config startup-config</code>                                                                                           | Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.                                                                                              |

#### Example

This example shows how to create a VRF and add a static route to the VRF:

```

switch# configure terminal
switch(config)# vrf context Enterprise
switch(config-vrf)# ip route 192.0.2.0/8 ethernet 1/2
switch(config-vrf)# exit
switch(config)# copy running-config startup-config

```

## Assigning VRF Membership to an Interface

### Before you begin

- Ensure that you are in the correct VDC or use **switchto vdc** command).
- Assign the IP address for an interface after you have configured the interface for a VRF

### Procedure

|               | Command or Action                                                                              | Purpose                                                                                                                       |
|---------------|------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configure terminal</b>                                                              | Enters global configuration mode.                                                                                             |
| <b>Step 2</b> | switch(config)# <b>interface</b> <i>interface-type slot/port</i>                               | Enters interface configuration mode.                                                                                          |
| <b>Step 3</b> | switch(config-if)# <b>vrf member</b> <i>vrf-name</i>                                           | Adds this interface to a VRF.                                                                                                 |
| <b>Step 4</b> | switch(config-if)# <b>ip address</b> <i>ip-prefix/length</i>                                   | Configures an IP address for this interface. You must do this step after you assign this interface to a VRF.                  |
| <b>Step 5</b> | (Optional) switch(config-vrf)# <b>show vrf</b> <i>vrf-name interface interface-type number</i> | Displays VRF information.                                                                                                     |
| <b>Step 6</b> | switch(config-vrf)# <b>exit</b>                                                                | Exits the current configuration mode.                                                                                         |
| <b>Step 7</b> | (Optional) switch(config)# <b>copy running-config startup-config</b>                           | Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration. |

### Example

This example shows how to add an interface to the VRF:

```

switch# configure terminal
switch(config)# interface ethernet 1/2
switch(config-if)# vrf member RemoteOfficeVRF
switch(config-if)# ip address 192.0.2.1/16
switch(config-if)# copy running-config startup-config

```

## Configuring VRF Parameters for a Routing Protocol

You can associate a routing protocol with one or more VRFs. See the appropriate chapter for information on how to configure VRFs for routing protocol. This section uses OSPFv2 as an example protocol for the detailed configuration steps.

### Procedure

|               | Command or Action                                                                      | Purpose                                                                                                                                   |
|---------------|----------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configure terminal</b>                                                      | Enters global configuration mode.                                                                                                         |
| <b>Step 2</b> | switch(config)# <b>router ospf</b> <i>instance tag</i>                                 | Creates a new OSPFv2 instance with the configured instance tag.                                                                           |
| <b>Step 3</b> | switch(config-router)# <b>vrf</b> <i>vrf-name</i>                                      | Enters VRF configuration mode.                                                                                                            |
| <b>Step 4</b> | switch(config-router-vrf)# <b>maximum-paths</b> <i>paths</i>                           | (Optional) Configures the maximum number of equal OSPFv2 paths to a destination in the route table for this VRF. Used for load balancing. |
| <b>Step 5</b> | switch(config)# <b>interface</b> <i>interface-type slot/port</i>                       | Enters interface configuration mode.                                                                                                      |
| <b>Step 6</b> | switch(config-if)# <b>ip address</b> <i>ip-prefix/length</i>                           | Assigns this interface to the OSPFv2 instance and area configured.                                                                        |
| <b>Step 7</b> | switch(config-if)# <b>ip address</b> <i>ip-prefix/length</i>                           | Configures an IP address for this interface. You must do this step after you assign this interface to a VRF.                              |
| <b>Step 8</b> | switch(config-if)# <b>ip router ospf area</b> <i>area-id instance-tag area area-id</i> | Assigns this interface to the OSPFv2 instance and area configured.                                                                        |
| <b>Step 9</b> | (Optional) switch(config)# <b>copy running-config startup-config</b>                   | Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.             |

### Example

This example shows how to add an interface to the VRF:

```
switch# configure terminal
switch(config)# vrf context RemoteOfficeVRF
switch(config-vrf)# exit
switch(config)# router ospf 201
switch(config-router)# vrf RemoteOfficeVRF
switch(config-router-vrf)# maximum-paths 4
switch(config-router-vrf)# interface ethernet 1/2
switch(config-if)# vrf member RemoteOfficeVRF
switch(config-if)# ip address 192.0.2.1/16
switch(config-if)# ip router ospf 201 area 0
switch(config-if)# exit
switch(config)# copy running-config startup-config
```

## Configuring VRF Aware Service

You can configure a VRF-aware service for reachability and filtering. See the “VRF-Aware Services” section for links to the appropriate chapter or configuration guide for information on how to configure the service for VRFs. This section uses SNMP and IP domain lists as example services for the detailed configuration steps.

### Procedure

|               | Command or Action                                                                                                                        | Purpose                                                                                                                                                                                          |
|---------------|------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configure terminal</b>                                                                                                        | Enters global configuration mode.                                                                                                                                                                |
| <b>Step 2</b> | switch(config)# <b>snmp-server host</b><br><i>ip-address</i> [ <b>filter-vrf</b> <i>vrf-name</i> ] [ <b>use-vrf</b><br><i>vrf-name</i> ] | Configures a global SNMP server and configures the VRF that Cisco NX-OS uses to reach the service. Use the <b>filter-vrf</b> keyword to filter information from the selected VRF to this server. |
| <b>Step 3</b> | switch(config)# <b>vrf context</b> <i>vrf-name</i>                                                                                       | Creates a new VRF.                                                                                                                                                                               |
| <b>Step 4</b> | switch(config-vrf)# <b>ip domain-list</b><br><i>domain-name</i> [ <b>all-vrfs</b> ] [ <b>use-vrf</b> <i>vrf-name</i> ]                   | Configures the domain list in the VRF and optionally configures the VRF that Cisco NX-OS uses to reach the domain name listed.                                                                   |
| <b>Step 5</b> | (Optional) switch(config)# <b>copy</b><br><b>running-config startup-config</b>                                                           | Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.                                                                    |

### Example

This example shows how to send SNMP information for all VRFs to SNMP host 192.0.2.1, reachable on VRF Red:

```
switch# configure terminal
switch(config)# snmp-server host 192.0.2.1 for-all-vrfs use-vrf Red
switch(config)# copy running-config startup-config
```

This example shows how to filter SNMP information for VRF Blue to SNMP host 192.0.2.12, reachable on VRF Red:

```
switch# configure terminal
switch(config)# vrf context Blue
switch(config-vrf)# snmp-server host 192.0.2.12 use-vrf Red
switch(config)# copy running-config startup-config
```

## Setting the VRF Scope

You can set the VRF scope for all EXEC commands (for example, show commands). This automatically restricts the scope of the output of EXEC commands to the configured VRF. You can override this scope by using the VRF keywords available for some EXEC commands.

To set the VRF scope, use the following command in EXEC mode:

**Procedure**

|               | Command or Action                                  | Purpose                                                                                                                                                                                                                                                                       |
|---------------|----------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>routing-context vrf</b> <i>vrf-name</i> | Sets the routing context for all EXEC commands. Default routing context is the default VRF.<br><br><b>Note</b> To return to the default VRF scope, use the following command in EXEC mode:<br><br><b>routing-context vrf default</b><br><br>Sets the default routing context. |

## Verifying the VRF Configuration

To display VRF configuration information, perform one of the following tasks:

**Procedure**

|               | Command or Action                                                                       | Purpose                                           |
|---------------|-----------------------------------------------------------------------------------------|---------------------------------------------------|
| <b>Step 1</b> | <b>show vrf</b> [ <i>vrf-name</i> ]                                                     | Displays the information for all or one VRF.      |
| <b>Step 2</b> | <b>show vrf</b> [ <i>vrf-name</i> ] <b>detail</b>                                       | Displays detailed information for all or one VRF. |
| <b>Step 3</b> | <b>show vrf</b> [ <i>vrf-name</i> ] [ <b>interface</b> <i>interface-typeslot/port</i> ] | Displays the VRF status for an interface.         |

## Configuration Examples for VRF

This example shows how to configure VRF Red, add an SNMP server to that VRF, and add an instance of OSPF to VRF Red:

```
configure terminal
vrf context Red
snmp-server host 192.0.2.12 use-vrf Red
router ospf 201
vrf Red
interface ethernet 1/2
vrf member Red
ip address 192.0.2.1/16
ip router ospf 201 area 0
```

This example shows how to configure VRF Red and Blue, add an instance of OSPF to each VRF, and create an SNMP context for each OSPF instance in each VRF:

```
configure terminal
!Create the VRFs
vrf context Red
vrf context Blue
```

```

vrf context Green
!Create the OSPF instances and associate them with a single VRF or multiple VRFs
(recommended)
feature ospf
router ospf Lab
vrf Red
!
router ospf Production
vrf Blue
router-id 1.1.1.1
vrf Green
router-id 2.2.2.2
!Configure one interface to use ospf Lab on VRF Red
interface ethernet 1/2
vrf member Red
ip address 192.0.2.1/16
ip router ospf Lab area 0
no shutdown
!Configure another interface to use ospf Production on VRF Blue
interface ethernet 10/2
vrf member Blue
ip address 192.0.2.1/16
ip router ospf Production area 0
no shutdown
!
interface ethernet 10/3
vrf member Green
ip address 192.0.2.1/16
ip router ospf Production area 0
no shutdown
!configure the SNMP server
snmp-server user admin network-admin auth md5 nbv-12345
snmp-server community public ro
!Create the SNMP contexts for each VRF
snmp-server context lab instance Lab vrf Red
snmp-server context production instance Production vrf Blue
!

```

Use the SNMP context lab to access the OSPF-MIB values for the OSPF instance Lab in VRF Red in this example. Use the SNMP context lab to access the OSPF-MIB values for the OSPF instance Lab in VRF Red in this example.

## Related Documents for VRF

| Related Topic | Document Title                                                                                                                                                                                                       |
|---------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| VRF CLI       | <i>Cisco Nexus 7000 Series NX-OS Unicast Routing Command Reference</i>                                                                                                                                               |
| VRFs          | <i>Cisco Nexus 7000 Series NX-OS Fundamentals Configuration Guide</i><br><i>Cisco Nexus 7000 Series NX-OS MPLS Configuration Guide</i><br><i>Cisco Nexus 7000 Series NX-OS System Management Configuration Guide</i> |

| Related Topic | Document Title                                                                  |
|---------------|---------------------------------------------------------------------------------|
| VDCs          | <i>Cisco Nexus 7000 Series NX-OS Virtual Device Context Configuration Guide</i> |

## Standards for VRF

No new or modified standards are supported by this feature, and support for existing standards has not been modified by this feature.

## Feature History for VRF

This table includes only the updates for those releases that have resulted in additions or changes to the feature.

*Table 28: Feature History for VRF*

| Feature Name | Release | Feature Information          |
|--------------|---------|------------------------------|
| VRF          | 4.0(1)  | This feature was introduced. |





## CHAPTER 17

# Managing the Unicast RIB and FIB

---

This chapter describes how to manage routes in the unicast Routing Information Base (RIB) and the Forwarding Information Base (FIB) on the Cisco NX-OS device.

- [Finding Feature Information, on page 383](#)
- [Information About the Unicast RIB and FIB, on page 383](#)
- [Default Settings for the Unicast RIB and FIB, on page 386](#)
- [Managing the Unicast RIB and FIB, on page 386](#)
- [Verifying the Unicast RIB and FIB, on page 395](#)
- [Related Documents for the Unicast RIB and FIB, on page 395](#)
- [Feature History for the Unicast RIB and FIB, on page 395](#)

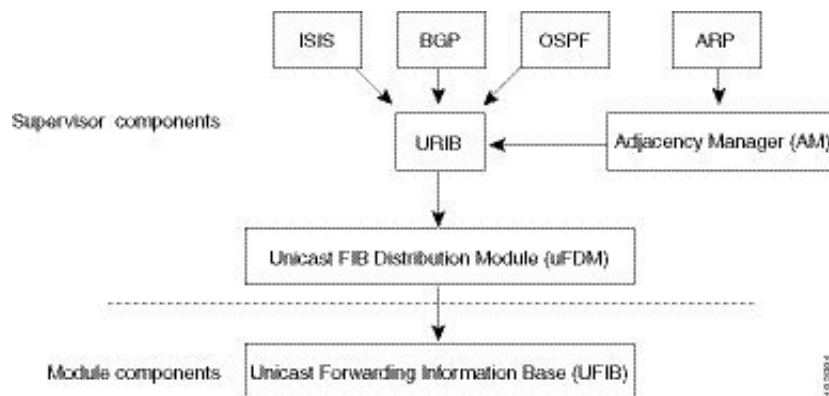
## Finding Feature Information

Your software release might not support all the features documented in this module. For the latest caveats and feature information, see the Bug Search Tool at <https://tools.cisco.com/bugsearch/> and the release notes for your software release. To find information about the features documented in this module, and to see a list of the releases in which each feature is supported, see the "New and Changed Information" chapter or the Feature History table in this chapter.

## Information About the Unicast RIB and FIB

The unicast RIB (IPv4 RIB and IPv6 RIB) and FIB are part of the Cisco NX-OS forwarding architecture,

Figure 44: Cisco NX-OS Forwarding Architecture



The unicast RIB exists on the active supervisor. It maintains the routing table with directly connected routes, static routes, and routes learned from dynamic unicast routing protocols. The unicast RIB also collects adjacency information from sources such as the Address Resolution Protocol (ARP). The unicast RIB determines the best next hop for a given route and populates the unicast forwarding information bases (FIBs) on the modules by using the services of the unicast FIB distribution module (FDM).

Each dynamic routing protocol must update the unicast RIB for any route that has timed out. The unicast RIB then deletes that route and recalculates the best next hop for that route (if an alternate path is available).

## Layer 3 Consistency Checker

In rare instances, an inconsistency can occur between the unicast RIB and the FIB on each module. Cisco NX-OS supports the Layer 3 consistency checker. This feature detects inconsistencies between the unicast IPv4 RIB on the supervisor module and the FIB on each interface module. Inconsistencies include the following:

- Missing prefix
- Extra prefix
- Wrong next-hop address
- Incorrect Layer 2 rewrite string in the ARP or neighbor discovery (ND) cache

The Layer 3 consistency checker compares the FIB entries to the latest adjacency information from the Adjacency Manager (AM) and logs any inconsistencies. The consistency checker then compares the unicast RIB prefixes to the module FIB and logs any inconsistencies.

## Dynamic TCAM Allocation

Dynamic TCAM allocation reallocates unused TCAM blocks on M1 Series non-XL modules to an adjacent region when all existing blocks in that region are full. Dynamic TCAM allocation allows more flexibility in the number of routes that the FIB can allocate for a route type.

Cisco NX-OS divides the FIB to support multiple address families. The FIB TCAM for M1 Series non-XL modules has 128K physical entries.

Table 29: Default FIB TCAM Allocation

| Region              | Default Number of Routes | Number of TCAM Blocks | Entry Size |
|---------------------|--------------------------|-----------------------|------------|
| IPv4 Unicast Routes | 56,000                   | 7                     | 72 bits    |

| Region                                       | Default Number of Routes | Number of TCAM Blocks | Entry Size |
|----------------------------------------------|--------------------------|-----------------------|------------|
| IPv4 Multicast Routes or IPv6 Unicast Routes | 32,000                   | 8                     | 144 bits   |
| IPv6 Multicast Routes                        | 2,000                    | 1                     | 288 bits   |

## Maximum TCAM Entries and FIB Scale Limits

The FIB TCAM entries are system wide resources that are shared across virtual device contexts (VDC) configured on the module. Table 16-2 describes the supported maximum FIB scale entries on the Nexus 7000 system configuration per route-type.

**Table 30: Maximum Supported TCAM Entries and FIB Scale Limits**

| Module Type in a VDC                         | Maximum TCAM Physical Entries in a VDC | Maximum Supported IPv4 Unicast Routes | Maximum Supported IPv4 Multicast Routes | Maximum Supported IPv6 Unicast Routes | Maximum Supported IPv6 Multicast Routes |
|----------------------------------------------|----------------------------------------|---------------------------------------|-----------------------------------------|---------------------------------------|-----------------------------------------|
| Only non-XL modules in a VDC                 | 128,000                                | 112,000                               | 32,000 mroutes                          | 56,000 routes                         | 2,000 routes                            |
| Only XL modules in a VDC                     | 900,000                                | 900,000                               | 32,000 mroutes                          | 350,000 routes                        | 2,000 routes                            |
| Mix of XL/non-XL modules in the same VDC     | 128,000                                | 112,000                               | 32,000 mroutes                          | 56,000 routes                         | 2,000 routes                            |
| Only F2 Series modules in a VDC <sup>2</sup> | 32,000                                 | 32,768                                | 16,384 mroutes                          | 16,384 routes                         | 8,192 routes                            |

<sup>2</sup> Utilization may vary based on the sequence of routes being added and on the mix of unicast and multicast routes.



**Note** The table above captures the scale limits in a VDC. In a Cisco Nexus 7000 system, the total memory on the supervisor module restricts the actual route scale limits across all VDCs in the system.



**Note** Do not exceed the maximum route limits for non-XL modules in a VDC that contains both XL modules and non-XL modules.



**Note** The actual FIB TCAM can scale to a higher scale number from a hardware perspective. The table above captures the currently supported FIB sizes.



**Note** The maximum routes are individual route-type maximum values and these values are not cumulative across each route-type.

Configure the higher shared memory sizes (see the Cisco Nexus 7000 Series NX-OS Virtual Device Context Configuration Guide, Release 5.x) for the routing table to enable the higher FIB scale on the XL modules. See the Cisco Nexus 7000 Series Hardware Installation and Reference Guide for more information on the XL modules.



**Note** The full IPv4 Internet route tables currently have more than 500,000 routes and require the XL modules.

The unicast RIB and FIB support virtual routing and forwarding (VRF) instances. VRF exists within VDCs. By default, Cisco NX-OS places you in the default VDC and default VRF unless you specifically configure another VDC and VRF. For more information, see the *Cisco Nexus 7000 Series NX-OS Virtual Device Context Configuration Guide*.

## Default Settings for the Unicast RIB and FIB

*Table 31: Default Unicast RIB and FIB Parameters*

| Parameters              | Default                                    |
|-------------------------|--------------------------------------------|
| Dynamic TCAM allocation | Enabled by default and cannot be disabled. |

## Managing the Unicast RIB and FIB



**Note** If you are unfamiliar with the Cisco IOS CLI, be aware that the Cisco NX-OS commands for this feature might differ from the Cisco IOS commands that you would use.

## Displaying Module FIB Information

The following show commands can be entered in any mode.

### Procedure

|               | Command or Action                                                                                   | Purpose                                              |
|---------------|-----------------------------------------------------------------------------------------------------|------------------------------------------------------|
| <b>Step 1</b> | switch# <b>show forwarding</b> { <b>ipv4</b>   <b>ipv6</b> }<br><b>adjacency module</b> <i>slot</i> | Displays the adjacency information for IPv4 or IPv6. |
| <b>Step 2</b> | switch# <b>show forwarding</b> { <b>ipv4</b>   <b>ipv6</b> } <b>route</b><br><b>moduleslot</b>      | Displays the route table for IPv4 or IPv6.           |

## Configuring Load Sharing in the Unicast FIB

Dynamic routing protocols such as Open Shortest Path First (OSPF) support load balancing with equal-cost multipath (ECMP). The routing protocol determines its best routes based on the metrics configured for the protocol and installs up to the protocol-configured maximum paths in the unicast RIB. The unicast RIB compares the administrative distances of all routing protocol paths in the RIB and selects a best path set from all of the path sets installed by the routing protocols. The unicast RIB installs this best path set into the FIB for use by the forwarding plane.

The forwarding plane uses a load-sharing algorithm to select one of the installed paths in the FIB to use for a given data packet.

You can globally configure the following load-sharing settings:

- **Load-share mode**—Selects the best path based on the destination address and port or the source and the destination address and port.
- **Universal ID**—Sets the random seed for the hash algorithm. You do not need to configure the Universal ID. Cisco NX-OS chooses the Universal ID if you do not configure it.



---

**Note** Load sharing uses the same path for all packets in a given flow. A flow is defined by the load-sharing method that you configure. For example, if you configure source-destination load sharing, then all packets with the same source IP address and destination IP address pair follow the same path.

---

To configure the unicast FIB load-sharing algorithm, use the following command in global configuration mode:

| Command                                                                                                                                                                                         | Purpose                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <pre>switch(config)# ip load-sharing address {destination port destination   source-destination [port source-destination]} [universal-id seed] [gtp-teid] [rotate rotate] [concatenation]</pre> | <p>Configures the unicast FIB load-sharing algorithm for data traffic.</p> <ul style="list-style-type: none"> <li>The <b>universal-id</b> option sets the random seed for the hash algorithm and shifts the flow from one link to another.</li> </ul> <p>You do not need to configure the universal ID. Cisco NX-OS chooses the Universal ID if you do not configure it. The universal-id range is from 1 to 4294967295.</p> <ul style="list-style-type: none"> <li>The <b>rotate</b> option causes the hash algorithm to rotate the link picking selection so that it does not continually choose the same link across all nodes in the network. It does so by influencing the bit pattern for the hash algorithm. This option shifts the flow from one link to another and load balances the already load-balanced (polarized) traffic from the first ECMP level across multiple links.</li> </ul> <p>If you specify a rotate value, the 64-bit stream is interpreted starting from that bit position in a cyclic rotation. The rotate range is from 1 to 63, and the default is 32.</p> <p><b>Note</b> With multi-tier Layer 3 topology, polarization is possible. To avoid polarization, use a different rotate bit at each tier of the topology.</p> <p><b>Note</b> To configure a rotation value for port channels, use the port-channel load-balance src-dst ip-l4port rotate rotate command. For more information on this command, see the <i>Cisco Nexus 7000 Series NX-OS Interfaces Configuration Guide</i>.</p> <ul style="list-style-type: none"> <li>For packets with GTP header, <b>gtp-teid</b> specifies that 32-bit TEID value has to be considered for the path calculation.</li> <li>The concatenation option ties together the hash tag values for ECMP and the hash tag values for port channels in order to use a stronger 64-bit hash. If you do not use this option, you can control ECMP load-balancing and port-channel load-balancing independently. The default is disabled.</li> </ul> |

To display the unicast FIB load-sharing algorithm, use the following command in any mode:

| Command                                     | Purpose                                                           |
|---------------------------------------------|-------------------------------------------------------------------|
| switch(config)# <b>show ip load-sharing</b> | Displays the unicast FIB load-sharing algorithm for data traffic. |

To display the route that the unicast RIB and FIB use for a particular source address and destination address, use the following command in any mode:

| Command                                                                                                                          | Purpose                                                                                                                                                                                                                                                                                                  |
|----------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| switch# <b>show routing hash</b> <i>source-addr dest-addr</i><br>[ <i>source-port dest-port</i> ] [ <b>vrf</b> <i>vrf-name</i> ] | Displays the route that the unicast RIB FIB use for a source and destination address pair. The source address and destination address format is x.x.x.x. The source port and destination port range is from 1 to 65535. The VRF name can be any case-sensitive, alphanumeric string up to 64 characters. |

This example shows the route selected for a source/destination pair:

```
switch# show routing hash 10.0.0.5 30.0.0.2
Load-share parameters used for software forwarding:
load-share mode: address source-destination port source-destination
Universal-id seed: 0xe05e2e85
Hash for VRF "default"
Hashing to path *20.0.0.2 (hash: 0x0e), for route:
```

## Configuring Per-Packet Load Sharing

You can use per-packet load sharing to evenly distribute data traffic in an IP network over multiple equal-cost connections. Per-packet load sharing allows the router to send successive data packets over paths on a packet-by-packet basis rather than on a per-flow basis.



**Note** Using per-packet load sharing can result in out-of-order packets. Packets for a given pair of source-destination hosts might take different paths and arrive at the destination out of order. Make sure you understand the implications of out-of-order packets to your network and applications. Per-packet load sharing is not appropriate for all networks. Per-flow load sharing ensures packets always arrive in the order that they were sent.

Per-packet load sharing uses the round-robin method to determine which path each packet takes to the destination. With per-packet load sharing enabled on interfaces, the router sends one packet for destination1 over the first path, the second packet for (the same) destination1 over the second path, and so on. Per-packet load sharing ensures balancing over multiple links.

Use per-packet load sharing to ensure that a path for a single source-destination pair does not get overloaded. If most of the traffic passing through parallel links is for a single pair, per-destination load sharing will overload a single link while other links will have very little traffic. Enabling per-packet load sharing allows you to use alternate paths to the same busy destination.



**Note** Per-packet load sharing on an interface overrides the global load-sharing configuration.

You configure per-packet load sharing on the input interface. This configuration determines the output interface that Cisco NX-OS chooses for the packet.

For example, if you have ECMP paths on two output interfaces, Cisco NX-OS uses the following load-sharing methods for input packets on Ethernet 1/1:

- Per-packet load sharing if you configure per-packet load sharing on Ethernet 1/1.
- Per-flow load sharing.

The configurations for the other interfaces have no effect on the load-sharing method used for Ethernet 1/1 in this example.

#### Procedure

|               | Command or Action                                    | Purpose                                             |
|---------------|------------------------------------------------------|-----------------------------------------------------|
| <b>Step 1</b> | switch(config-if)# <b>ip load-sharing per-packet</b> | Configures per-packet load sharing on an interface. |

## Displaying Routing and Adjacency Information

#### Procedure

|               | Command or Action                                                                                                                                                                                                                   | Purpose                                                                                                                                                                                                                                                                                                                                                  |
|---------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>show {ip   ipv6} route</b> [ <i>route-type</i>   <b>interface</b> <i>int-type number</i>   <b>next-hop</b> ]                                                                                                             | Displays the unicast route table. The route-type argument can be a single route prefix, direct, static, or a dynamic route protocol. Use the ? command to see the supported interfaces.                                                                                                                                                                  |
| <b>Step 2</b> | switch# <b>show {ip   ipv6} adjacency</b> [ <i>prefix</i>   <i>interface-type number</i> [ <b>summary</b> ]   <b>non-best</b> ] [ <b>detail</b> ] [ <b>vrf</b> <i>vrf-id</i> ]                                                      | Displays the adjacency table.<br>The argument ranges are as follows: <ul style="list-style-type: none"> <li>• <i>prefix</i>—Any IPv4 or IPv6 prefix address.</li> <li>• <i>interface-type number</i>—Use the ? command to see the supported interfaces.</li> <li>• <i>vrf-id</i>—Any case-sensitive, alphanumeric string up to 64 characters.</li> </ul> |
| <b>Step 3</b> | switch# <b>show {ip   ipv6} routing</b> [ <i>route-type</i>   <b>interface</b> <i>int-type number</i>   <b>next-hop</b>   <b>recursive-next-hop</b>   <b>summary</b>   <b>updated</b> { <b>since</b>   <b>until</b> } <i>time</i> ] | Displays the unicast route table. The route-type argument can be a single route prefix, direct, static, or a dynamic route protocol. Use the ? command to see the supported interfaces.                                                                                                                                                                  |

#### Example

The following example displays the unicast route table:



```

switch# show ip route

IP Route Table for Context "default"
 '*' denotes best ucast next-hop '*' denotes best mcast next-hop
 '[x/y]' denotes [preference/metric]

0.0.0.0/0, 1 ucast next-hops, 0 mcast next-hops
 *via 10.1.1.1, mgmt0, [1/0], 5d21h, static
0.0.0.0/32, 1 ucast next-hops, 0 mcast next-hops
 *via Null0, [220/0], 1w6d, local, discard
10.1.0.0/22, 1 ucast next-hops, 0 mcast next-hops, attached
 *via 10.1.1.55, mgmt0, [0/0], 5d21h, direct
10.1.0.0/32, 1 ucast next-hops, 0 mcast next-hops, attached
 *via 10.1.1.0.0, Null0, [0/0], 5d21h, local
10.1.1.1/32, 1 ucast next-hops, 0 mcast next-hops, attached
 *via 10.1.1.1, mgmt0, [2/0], 5d16h, am
10.1.1.55/32, 1 ucast next-hops, 0 mcast next-hops, attached
 *via 10.1.1.55, mgmt0, [0/0], 5d21h, local
10.1.1.253/32, 1 ucast next-hops, 0 mcast next-hops, attached
 *via 10.1.1.253, mgmt0, [2/0], 5d20h, am
10.1.3.255/32, 1 ucast next-hops, 0 mcast next-hops, attached
 *via 10.1.3.255, mgmt0, [0/0], 5d21h, local
255.255.255.255/32, 1 ucast next-hops, 0 mcast next-hops
 *via Eth Inband Port, [0/0], 1w6d, local

```

The following example shows the adjacency information:

```

switch# show ip adjacency

IP Adjacency Table for context default
Total number of entries: 2
Address Age MAC Address Pref Source Interface Best
10.1.1.1 02:20:54 00e0.b06a.71eb 50 arp mgmt0 Yes
10.1.1.253 00:06:27 0014.5e0b.81d1 50 arp mgmt0 Yes

```

## Triggering the Layer 3 Consistency Checker

You can manually trigger the Layer 3 consistency checker.

To manually trigger the Layer 3 consistency checker, use the following commands in global configuration mode:

| Command                                                                                          | Purpose                                                                                                                                                            |
|--------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>test forwarding [ipv4   ipv6] [unicast] inconsistency [vrf vrf-name] [module {slot  all}]</b> | Starts a Layer 3 consistency check. The <i>vrf-name</i> can be any case-sensitive, alphanumeric string up to 64 characters. The <i>slot</i> range is from 1 to 10. |

To stop the Layer 3 consistency checker, use the following commands in global configuration mode:

| Command                                                                                               | Purpose                                                                                                                                                           |
|-------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>test forwarding [ipv4   ipv6] [unicast] inconsistency [vrf vrf-name] [module {slot  all}] stop</b> | Stops a Layer 3 consistency check. The <i>vrf-name</i> can be any case-sensitive, alphanumeric string up to 64 characters. The <i>slot</i> range is from 1 to 10. |

To display the Layer 3 inconsistencies, use the following commands in any mode:

| Command                                                                                       | Purpose                                                                                                                                                                             |
|-----------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>show forwarding [ipv4   ipv6] inconsistency [vrf vrf-name] [module {slot   all}]</code> | Displays the results of a Layer 3 consistency check. The <i>vrf-name</i> can be any case-sensitive, alphanumeric string up to 64 characters. The <i>slot</i> range is from 1 to 10. |

## Clearing Forwarding Information in the FIB

You can clear one or more entries in the FIB. Clearing a FIB entry does not affect the unicast RIB.



**Caution** The `clear forwarding` command disrupts forwarding on the device.

### Procedure

|               | Command or Action                                                                                         | Purpose                                                                                                                                                                                                                                                                                                               |
|---------------|-----------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | <code>switch# clear forwarding {ipv4   ipv6} route {*   prefix} [vrf vrf-name] module {slot   all}</code> | Clears one or more entries from the FIB. The route options are: <ul style="list-style-type: none"> <li>• <i>*</i>—all routes.</li> <li>• <i>prefix</i>—Any IP or IPv6 prefix.</li> </ul> The <i>vrf-name</i> can be a case-sensitive, alphanumeric string up to 64 characters. The <i>slot</i> range is from 1 to 10. |
| <b>Step 2</b> | (Optional) <code>switch(config-if)# copy running-config startup-config</code>                             | Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.                                                                                                                                                                                         |

### Example

The following example shows how to clear one or more entries from the FIB:

```
switch(config)# clear forwarding ipv4 route * module 1
switch(config-if)# copy running-config startup-config
```

## Configuring Maximum Routes for the Unicast RIB

You can configure the maximum number of routes allowed in the routing table.

### Before you begin

Ensure that you are in the default VDC (or use the `switchto vdc` command).

**Procedure**

|               | <b>Command or Action</b>                                                                                                                               | <b>Purpose</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|---------------|--------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configure terminal</b>                                                                                                                      | Enters global configuration mode.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Step 2</b> | switch (config)# <b>vrf context</b> <i>vrf name</i>                                                                                                    | Creates a VRF and enters VRF configuration mode.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Step 3</b> | switch (config-vrf)# <b>ip4 unicast</b>                                                                                                                | Enters address family configuration mode.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Step 4</b> | switch (config vrf-af-ipv4)# <b>maximum routes</b> <i>max routes</i> [ <i>threshold</i> [ <b>reinstall</b> <i>threshold</i> ]   <b>warning -only</b> ] | Configures the maximum number of routes allowed in the routing table.<br><br>You can optionally specify the following: <ul style="list-style-type: none"> <li>• <i>threshold</i>—Percentage of maximum routes that triggers a warning message. The range is from 1 to 100.</li> <li>• <b>warning-only</b>—Logs a warning message when the maximum number of routes is exceeded.</li> <li>• <b>reinstall</b> <i>threshold</i>—Reinstalls routes that previously exceeded the maximum route limit and were rejected and specifies the threshold value at which to reinstall them. The threshold range is from 1 to 100.</li> </ul> |
| <b>Step 5</b> | (Optional) switch(config-if)# <b>copy running-config startup-config</b>                                                                                | Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |

**Example**

The following example shows how to configure maximum routes for the unicast RIB:

```
switch# configure terminal
switch(config)# vrf context Red
switch(config-vrf)# ip4 unicast
switch(config-vrf-af-ipv4)# maximum routes 250 90
switch(config-vrf-af-ipv4)# copy running-config startup-config
```

## Estimating Memory Requirements for Routes

You can estimate the memory that will be used by a number of routes and next-hop addresses.

**Procedure**

|               | <b>Command or Action</b>                                                                                                         | <b>Purpose</b>                                                                                                                                   |
|---------------|----------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>show routing</b> { <b>ipv6</b> } <b>memory estimate routes</b> <i>num-routes</i> <b>next-hops</b> <i>num-nexthops</i> | Displays the memory requirements for routes. The <i>num-routes</i> range is from 1000 to 1000000. The <i>num-nexthops</i> range is from 1 to 16. |

|               | Command or Action                                                             | Purpose                                                                                                                       |
|---------------|-------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 2</b> | (Optional) <code>switch(config-if)# copy running-config startup-config</code> | Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration. |

### Example

The following example shows how to estimate the memory requirements for routes:

```
switch# show routing memory estimate routes 5000 next-hops
switch(config-if)# copy running-config startup-config
```

## Clearing Routes in the Unicast RIB

You can estimate the memory to be used by a number of routes and next-hop addresses.



**Caution** The \* keyword is severely disruptive to routing.

### Procedure

|               | Command or Action                                                                                                                   | Purpose                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|---------------|-------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | <code>clear {ip   ip4   ipv6} route {*   prefix/length} [next hop interface] [vrf vrf-name] module {slot   all}</code>              | <p>Clears one or more routes from both the unicast RIB and all the module FIBs. The route options are:</p> <ul style="list-style-type: none"> <li>• <b>*</b>—All routes.</li> <li>• <b>route</b>— An individual IP or IPv6 route.</li> <li>• <i>prefix/length</i>— Any IP or IPv6 prefix.</li> <li>• <i>next-hop</i>—The next-hop address.</li> <li>• <i>interface</i>—The interface to reach the next-hop address.</li> </ul> <p>The <i>vrf-name</i> can be an case-sensitive, alphanumeric string up to 32 64characters.</p> |
| <b>Step 2</b> | <code>clear routing [multicast   unicast]{ip   ip4   ipv6} {*   {route   prefix/length} [next-hop interface]} [vrf vrf-name]</code> | <p>Clears one or more routes from both the unicast RIB and all the module FIBs. The route options are:</p> <ul style="list-style-type: none"> <li>• <b>*</b>—All routes.</li> <li>• <b>route</b>— An individual IP or IPv6 route.</li> <li>• <i>prefix/length</i>— Any IP or IPv6 prefix.</li> <li>• <i>next-hop</i>—The next-hop address.</li> <li>• <i>interface</i>—The interface to reach the next-hop address.</li> </ul>                                                                                                 |

|               | Command or Action                                                             | Purpose                                                                                                                       |
|---------------|-------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------|
|               |                                                                               | The <i>vrf-name</i> can be any case-sensitive, alphanumeric string up to 64 characters.                                       |
| <b>Step 3</b> | (Optional) <code>switch(config-if)# copy running-config startup-config</code> | Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration. |

## Verifying the Unicast RIB and FIB

To display advanced BGP statistics, use the following commands:

| Command                                                         | Purpose                                                |
|-----------------------------------------------------------------|--------------------------------------------------------|
| <code>show forwarding adjacency</code>                          | Displays the adjacency table on a module.              |
| <code>show forwarding distribution {clients   fib-state}</code> | Displays the FIB distribution information.             |
| <code>show forwarding interfaces module slot</code>             | Displays the FIB information for a module.             |
| <code>show forwarding {ip   ipv4   ipv6} route</code>           | Displays routes in the FIB.                            |
| <code>show {ip   ipv4   ipv6} adjacency</code>                  | Displays the adjacency table.                          |
| <code>show {ip   ipv6} route</code>                             | Displays the IPv4 or IPv6 routes from the unicast RIB. |
| <code>show routing</code>                                       | Displays routes from the unicast RIB.                  |

## Related Documents for the Unicast RIB and FIB

| Feature Name                     | Feature Information                                                    |
|----------------------------------|------------------------------------------------------------------------|
| Unicast RIB and FIB CLI commands | <i>Cisco Nexus 7000 Series NX-OS Unicast Routing Command Reference</i> |

## Feature History for the Unicast RIB and FIB

This table includes only the updates for those releases that have resulted in additions or changes to the feature.

Table 32: Feature History for the Unicast RIB and FIB

| Feature Name                          | Release     | Feature Information                                                                                                                               |
|---------------------------------------|-------------|---------------------------------------------------------------------------------------------------------------------------------------------------|
| Load Sharing in Unicast FIB           | 7.3(0)DX(1) | Added support for GTP headers.                                                                                                                    |
| Unicast FIB                           | 6.2(2)      | Added the ability to check for inconsistent, missing, or failed routes in the unicast FIB.                                                        |
| TCAM utilization                      | 6.2(2)      | Added the ability to monitor TCAM utilization on M1 Series modules.                                                                               |
| Unicast RIB                           | 6.2(2)      | Added the optional keyword <b>longer-prefixes [detail]</b> to the <b>show routing</b> command to display specific routes for a particular prefix. |
| Maximum routes                        | 5.2(1)      | Added support to configure the maximum number of routes allowed in the routing table.                                                             |
| TCAM Size for XL Modules              | 5.0(2)      | Added support for larger TCAM and FIB sizes with XL modules.                                                                                      |
| Dynamic TCAM allocation               | 5.0(2)      | Enabled by default and cannot be disabled.                                                                                                        |
| IPv6 forwarding inconsistency checker | 4.2(1)      | Added support to check for inconsistencies in the IPv6 forwarding table.                                                                          |
| Dynamic TCAM allocation               | 4.2(1)      | Added support for dynamically allocating TCAM blocks in the FIB.                                                                                  |
| Per-packet load sharing               | 4.1(2)      | Added support to load balance per packet on an interface.                                                                                         |
| Unicast RIB and FIB                   | 4.0(3)      | Added support to clear individual routes in unicast RIB and FIB.                                                                                  |
| Unicast RIB and FIB                   | 4.0(1)      | This feature was introduced.                                                                                                                      |



## CHAPTER 18

# Configuring Route Policy Manager

This chapter contains the following sections:

- [Finding Feature Information, on page 397](#)
- [Information About Route Policy Manager, on page 397](#)
- [Prerequisites for Route Policy Manager , on page 405](#)
- [Guidelines and Limitations, on page 405](#)
- [Default Settings for Route Policy Manager Parameters, on page 405](#)
- [Configuring Route Policy Manager, on page 406](#)
- [Configuration Examples for Route Policy Manager, on page 416](#)
- [Related Documents for Route Policy Manager, on page 416](#)
- [Standards for Route Policy Manager, on page 416](#)
- [Feature History for Route Policy Manager, on page 416](#)

## Finding Feature Information

Your software release might not support all the features documented in this module. For the latest caveats and feature information, see the Bug Search Tool at <https://tools.cisco.com/bugsearch/> and the release notes for your software release. To find information about the features documented in this module, and to see a list of the releases in which each feature is supported, see the "New and Changed Information" chapter or the Feature History table in this chapter.

## Information About Route Policy Manager

Route Policy Manager supports route maps and IP prefix lists. These features are used for route redistribution. A prefix list contains one or more IPv4 network prefixes and the associated prefix length values. You can use a prefix list by itself in features such as Border Gateway Protocol (BGP) templates, route filtering, or redistribution of routes that are exchanged between routing domains.

Route maps can apply to both routes and IP packets. Route filtering and redistribution pass a route through a route map.

## Prefix Lists

You can use prefix lists to permit or deny an address or range of addresses. Filtering by a prefix list involves matching the prefixes of routes or packets with the prefixes listed in the prefix list. An implicit deny is assumed if a given prefix does not match any entries in a prefix list.

You can configure multiple entries in a prefix list and permit or deny the prefixes that match the entry. Each entry has an associated sequence number that you can configure. If you do not configure a sequence number, Cisco NX-OS assigns a sequence number automatically. Cisco NX-OS evaluates prefix lists starting with the lowest sequence number. Cisco NX-OS processes the first successful match for a given prefix. Once a match occurs, Cisco NX-OS processes the permit or deny statement and does not evaluate the rest of the prefix list.

Prefix Lists in Cisco NX-OS support only one of the following addresses at a time:

- source
- destination
- group address



---

**Note** An empty prefix list permits all routes.

---

## MAC Lists

You can use MAC lists to permit or deny a MAC address or range of addresses. A MAC list consists of a list of MAC addresses and optional MAC masks. A MAC mask is a wild-card mask that is logically AND-ed with the MAC address when the route map matches on the MAC list entry. Filtering by a MAC list involves matching the MAC address of packets with the MAC addresses listed in the MAC list. An implicit deny is assumed if a given MAC address does not match any entries in a MAC list.

You can configure multiple entries in a MAC list and permit or deny the MAC addresses that match the entry. Each entry has an associated sequence number that you can configure. If you do not configure a sequence number, Cisco NX-OS assigns a sequence number automatically. Cisco NX-OS evaluates MAC lists starting with the lowest sequence number. Cisco NX-OS processes the first successful match for a given MAC address. Once a match occurs, Cisco NX-OS processes the permit or deny statement and does not evaluate the rest of the MAC list.

## Route Maps

You can use route maps for route redistribution or policy-based routing. Route map entries consist of a list of match and set criteria. The match criteria specify match conditions for incoming routes or packets, and the set criteria specify the action taken if the match criteria are met.

You can configure multiple entries in the same route map. These entries contain the same route map name and are differentiated by a sequence number.

You create a route map with one or more route map entries arranged by the sequence number under a unique route map name. The route map entry has the following parameters:

- Sequence number
- Permission—permit or deny



- Match criteria
- Set changes

The IPv4 and the IPv6-based matches on the same route map sequence number is not supported in Cisco Nexus 7000 Series.

By default, a route map processes routes or IP packets in a linear fashion, that is, starting from the lowest sequence number. You can configure the route map to process in a different order using the **continue** statement, which allows you to determine which route map entry to process next.

## Match Criteria

You can use a variety of criteria to match a route or IP packet in a route map. Some criteria, such as BGP community lists, are applicable only to a specific routing protocol, while other criteria, such as the IP source or the destination address, can be used for any route or IP packet.

When Cisco NX-OS processes a route or packet through a route map, it compares the route or packet to each of the match statements configured. If the route or packet matches the configured criteria, Cisco NX-OS processes it based on the permit or deny configuration for that match entry in the route map and any set criteria configured.

The match categories and parameters are as follows:

- IP access lists—(For policy-based routing only). Match based on source or destination IP address, protocol, or QoS parameters.
- BGP parameters—Match based on AS numbers, AS-path, community attributes, or extended community attributes.
- Prefix lists—Match based on an address or range of addresses.
- Multicast parameters—Match based on rendezvous point, groups, or sources.
- Other parameters—Match based on IP next-hop address or packet length.

## Set Changes

Once a route or packet matches an entry in a route map, the route or packet can be changed based on one or more configured set statements.

The set changes are as follows:

- BGP parameters—Change the AS-path, tag, community, extended community, dampening, local preference, origin, or weight attributes.
- Metrics—Change the route-metric, the route-tag, or the route-type.
- Policy-based routing only—Change the interface or the default next-hop address.
- Other parameters—Change the forwarding address or the IP next-hop address.

## Access Lists

IP access lists can match the packet to a number of IP packet fields such as the following:

- Source or destination IPv4 or IPv6 address

- Protocol
- Precedence
- ToS

You can use ACLs in a route map for policy-based routing only. See the *Cisco Nexus 7000 Series NX-OS Security Configuration Guide* for more information on ACLs.

## AS Numbers for BGP

You can configure a list of AS numbers to match against BGP peers. If a BGP peer matches an AS number in the list and matches the other BGP peer configuration, BGP creates a session. If the BGP peer does not match an AS number in the list, BGP ignores the peer. You can configure the AS numbers as a list, a range of AS numbers, or you can use an AS-path list to compare the AS numbers against a regular expression.

## AS-path Lists for BGP

You can configure an AS-path list to filter inbound or outbound BGP route updates. If the route update contains an AS-path attribute that matches an entry in the AS-path list, the router processes the route based on the permit or deny condition configured. You can configure AS-path lists within a route map.

You can configure multiple AS-path entries in an AS-path list by using the same AS-path list name. The router processes the first entry that matches.

## Community Lists for BGP

You can filter BGP route updates based on the BGP community attribute by using community lists in a route map. You can match the community attribute based on a community list, and you can set the community attribute using a route map.

A community list contains one or more community attributes. If you configure more than one community attribute in the same community list entry, the BGP route must match all community attributes listed to be considered a match.

You can also configure multiple community attributes as individual entries in the community list by using the same community list name. In this case, the router processes the first community attribute that matches the BGP route, using the permit or deny configuration for that entry.

You can configure community attributes in the community list in one of the following formats:

- A named community attribute, such as **internet** or **no-export**.
- In *aa:nn* format, where the first two bytes represent the two-byte AS number and the last two bytes represent a user-defined network number.
- A regular expression.

See the *Cisco Nexus 7000 Series NX-OS Unicast Routing Command Reference* for more information on regular expressions.

## Extended Community Lists for BGP

Extended community lists support 4-byte AS numbers. You can configure community attributes in the extended community list in one of the following formats:

- In *aa4:nn* format, where the first four bytes represent the four-byte AS number and the last two bytes represent a user-defined network number.
- A regular expression.

See the *Cisco Nexus 7000 Series NX-OS Unicast Routing Command Reference* for more information on regular expressions.

Cisco NX-OS supports generic specific extended community lists, which provide similar functionality to regular community lists for four-byte AS numbers. You can configure generic specific extended community lists with the following properties:

- Transitive—BGP propagates the community attributes across autonomous systems.
- Nontransitive—BGP removes community attributes before propagating the route to another autonomous system.

## Route Redistribution and Route Maps

You can use route maps to control the redistribution of routes between routing domains. Route maps match on the attributes of the routes to redistribute only those routes that pass the match criteria. The route map can also modify the route attributes during this redistribution using the set changes.

The router matches redistributed routes against each route map entry. If there are multiple match statements, the route must pass all of the match criteria. If a route passes the match criteria defined in a route map entry, the actions defined in the entry are executed. If the route does not match the criteria, the router compares the route against subsequent route map entries. Route processing continues until a match is made or the route is processed by all entries in the route map with no match.

Each ACL ends with an implicit deny statement, by design convention; there is no similar convention for route-maps. If the end of a route-map is reached during matching attempts, the result depends on the specific application of the route-map. Fortunately, route-maps that are applied to redistribution behave the same way as ACLs: if the route does not match any clause in a route-map, then the route redistribution is denied, as if the route-map contained a deny statement at the end.



---

**Note** When you redistribute BGP to IGP, iBGP is redistributed as well. To override this behavior, you must insert an additional deny statement into the route map.

---

## Route Map Support Matrix for Routing Protocols

The following tables include the configurable match and set statements for routing protocols on Cisco Nexus 70xx and 77xx Series switches running the latest shipping release. For specific release information, see the [Feature History for Route Policy Manager, on page 416](#).

The following legend applies to the tables:

- Yes—The statement is supported for the protocol.
- No—The statement is not supported for the protocol.
- If a statement does not apply for the protocol, there is an em dash (—) in the column next to the statement.

- Where clarification is required, information is added in the appropriate row/column.

Table 33: SET Route Map Statements by Protocol

| SET Route Map Statement      | OSPF Redistribution | EIGRP Redistribution | ISIS Redistribution | RIP Redistribution | BGP Redistribution      |
|------------------------------|---------------------|----------------------|---------------------|--------------------|-------------------------|
| Forwarding-address           | Yes                 | —                    | —                   | —                  | —                       |
| Standard/Extended Community  | —                   | —                    | —                   | —                  | Standard community only |
| Site of Origin (SOO)         | —                   | —                    | —                   | —                  | No                      |
| Routing Protocol Metric      | Yes                 | Yes                  | Yes                 | Yes                | Yes                     |
| Routing Protocol Metric Type | Yes                 | —                    | No                  | —                  | —                       |
| Route Tag                    | Yes                 | Yes                  | No                  | Yes                | —                       |
| NSSA Only                    | Yes                 | —                    | —                   | —                  | —                       |
| Origin                       | —                   | —                    | —                   | —                  | Yes                     |
| Level                        | —                   | —                    | Yes                 | —                  | —                       |
| Weight                       | —                   | —                    | —                   | —                  | Yes                     |

Table 34: SET Route Map Statements by Protocol

| SET Route Map Statement                   | BGP Neighbor | BGP Table Map | OSPF Table Map | EIGRP Table Map | ISIS Table Map | EIGRP Distribute List |
|-------------------------------------------|--------------|---------------|----------------|-----------------|----------------|-----------------------|
| Standard/Extended Community               | Yes          | No            | —              | —               | —              | —                     |
| Standard/Extended Community-List Deletion | Yes          | No            | —              | —               | —              | —                     |
| Site of Origin (SOO)                      | No           | —             | —              | —               | —              | —                     |
| Routing Protocol Metric                   | No           | No            | —              | —               | —              | Yes                   |
| Routing Protocol Metric Type              | Yes          | No            | —              | —               | —              | —                     |
| IPv4 Next Hop                             | Yes          | —             | —              | —               | —              | —                     |
| IPv6 Next Hop                             | Yes          | —             | —              | —               | —              | —                     |

| SET Route Map Statement | BGP Neighbor | BGP Table Map | OSPF Table Map | EIGRP Table Map | ISIS Table Map | EIGRP Distribute List |
|-------------------------|--------------|---------------|----------------|-----------------|----------------|-----------------------|
| IPv4 Prefix list        | Yes          | —             | —              | —               | —              | —                     |
| IPv6 Prefix list        | Yes          | —             | —              | —               | —              | —                     |
| Interface               | No           | —             | —              | —               | —              | —                     |
| Route Tag               | —            | —             | —              | —               | —              | Yes                   |
| AS PATH                 | Yes          | No            | —              | —               | —              | —                     |
| Orgin                   | Yes          | No            | —              | —               | —              | —                     |
| All Path Advertisement  | Yes          | No            | —              | —               | —              | —                     |
| Distance                | —            | Yes           | Yes            | Yes             | Yes            | —                     |
| Dampening               | No           | No            | —              | —               | —              | —                     |
| Level                   | No           | No            | —              | —               | —              | —                     |
| Weight                  | Yes          | Yes           | No             | —               | —              | —                     |

Table 35: MATCH Route Map Statements by Protocol

| MATCH Route Map Statement | OSPF Redistribution | EIGRP Redistribution | ISIS Redistribution | RIP Redistribution | BGP Redistribution |
|---------------------------|---------------------|----------------------|---------------------|--------------------|--------------------|
| Community List            | OSPFv2 only         | Yes                  | yes                 | Yes                | —                  |
| Ext Community List        | OSPFv2 only         | Yes                  | —                   | Yes                | —                  |
| Interface                 | Yes                 | Yes                  | Yes                 | Yes                | Yes                |
| IPv4 Next Hop             | Yes                 | Yes                  | Yes                 | Yes                | Yes                |
| IPv6 Next Hop             | Yes                 | Yes                  | Yes                 | No                 | Yes                |
| Metric                    | Yes                 | Yes                  | Yes                 | Yes                | Yes                |
| Route Type                | Yes                 | Yes                  | Yes                 | Yes                | Yes                |
| Tag                       | Yes                 | Yes                  | Yes                 | Yes                | Yes                |
| IPv6 Prefix List          | Yes                 | Yes                  | Yes                 | No                 | Yes                |
| IPv4 Prefix list          | Yes                 | Yes                  | Yes                 | Yes                | Yes                |
| IP ACL                    | No                  | No                   | No                  | No                 | No                 |
| Source Protocol           | Yes                 | Yes                  | Yes                 | Yes                | —                  |

| <b>MATCH Route Map Statement</b> | <b>OSPF Redistribution</b> | <b>EIGRP Redistribution</b> | <b>ISIS Redistribution</b> | <b>RIP Redistribution</b> | <b>BGP Redistribution</b> |
|----------------------------------|----------------------------|-----------------------------|----------------------------|---------------------------|---------------------------|
| AS Path                          | No                         | No                          | No                         | No                        | —                         |
| AS Number                        | No                         | No                          | No                         | No                        | —                         |

**Table 36: MATCH Route Map Statements by Protocol**

| <b>MATCH Route Map Statement</b> | <b>BGP Neighbor</b> | <b>BGP Table Map</b> | <b>OSPF Table Map</b> | <b>EIGRP Table Map</b> | <b>ISIS Table Map</b> | <b>EIGRP Distribute List</b> |
|----------------------------------|---------------------|----------------------|-----------------------|------------------------|-----------------------|------------------------------|
| Community List                   | Yes                 | Yes                  | —                     | —                      | —                     | —                            |
| Ext Community List               | Yes                 | Yes                  | —                     | —                      | —                     | —                            |
| Interface                        | —                   | No                   | Yes                   | Yes                    | Yes                   | —                            |
| IPv4 Next Hop                    | Yes                 | Yes                  | Yes                   | Yes                    | Yes                   | Yes                          |
| IPv6 Next Hop                    | Yes                 | Yes                  | Yes                   | Yes                    | Yes                   | Yes                          |
| Metric                           | Yes                 | Yes                  | Yes                   | No                     | No                    | No                           |
| Route Type                       | Yes                 | Yes                  | Yes                   | Yes                    | Yes                   | No                           |
| Tag                              | —                   | Yes                  | Yes                   | Yes                    | No                    | Yes                          |
| IPv6 Prefix List                 | Yes                 | Yes                  | Yes                   | Yes                    | Yes                   | Yes                          |
| IPv4 Prefix list                 | Yes                 | Yes                  | Yes                   | Yes                    | Yes                   | Yes                          |
| IP ACL                           | No                  | No                   | No                    | No                     | No                    | No                           |
| AS Path                          | Yes                 | Yes                  | —                     | —                      | —                     | —                            |
| AS Number                        | Yes                 | No                   | —                     | —                      | —                     | —                            |
| IPv4 Route Source                | —                   | —                    | Yes                   | —                      | —                     | —                            |

## Policy-Based Routing

You can use policy-based routing to forward a packet to a specified next-hop address based on the source of the packet or other fields in the packet header.

## Prerequisites for Route Policy Manager

If you configure VDCs, install the appropriate license and enter the desired VDC (see the *Cisco Nexus 7000 Series NX-OS Virtual Device Context Configuration Guide* for configuration information and the *Cisco NX-OS Licensing Guide* for licensing information).

## Guidelines and Limitations

- An empty route map denies all the routes.
- An empty prefix list permits all the routes.
- Without any match statement in a route-map entry, the permission (permit or deny) of the route-map entry decides the result for all the routes or packets.
- If referred policies (for example, prefix lists) within a match statement of a route-map entry return either a no-match or a deny-match, Cisco NX-OS fails the match statement and processes the next route-map entry.
- When you change a route map, Cisco NX-OS holds all the changes until you exit from the route-map configuration submode. Cisco NX-OS then sends all the changes to the protocol clients to take effect.
- Because you can use a route map before you define it, verify that all your route maps exist when you finish a configuration change.
- You can view the route-map usage for redistribution and filtering. Each individual routing protocol provides a way to display these statistics.
- When you redistribute BGP to IGP, iBGP is redistributed as well. To override this behavior, you must insert an additional deny statement into the route map.
- If you are familiar with the Cisco IOS CLI, be aware that the Cisco NX-OS commands for this feature might differ from the Cisco IOS commands that you would use.

## Default Settings for Route Policy Manager Parameters

**Default Route Policy Manager Parameters**

| Parameters              | Default |
|-------------------------|---------|
| Route Policy Manager    | Enabled |
| Administrative distance | 115     |

# Configuring Route Policy Manager

## Configuring IP Prefix Lists

IP prefix lists match the IP packet or route against a list of prefixes and prefix lengths. You can create an IP prefix list for IPv4 and create an IPv6 prefix list for IPv6.

You can configure the prefix list entry to match the prefix length exactly, or to match any prefix with a length that matches the configured range of prefix lengths.

Use the **ge** and **le** keywords to create a range of possible prefix lengths. The incoming packet or route matches the prefix list if the prefix matches and if the prefix length is greater than or equal to the **ge** keyword value (if configured) and less than or equal to the **le** keyword value (if configured).



**Note** When you use keyword **eq**, you must set the value greater than the mask length for the prefix.

### Procedure

|               | Command or Action                                                                                                                                 | Purpose                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|---------------|---------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configure terminal</b>                                                                                                                 | Enters global configuration mode.                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Step 2</b> | (Optional) switch(config)# <b>ip ipv6</b><br><b>prefix-list name description string</b>                                                           | Adds an information string about the prefix list.                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Step 3</b> | switch(config)# <b>ip prefix-list name [seq number] [{permit   deny} prefix {[eq prefix-length]   [ge prefix-length] [le prefix-length]}</b> }}   | Creates an IPv4 prefix list or adds a prefix to an existing prefix list. The prefix length is matched as follows: <ul style="list-style-type: none"> <li>• <b>eq</b>—Matches the exact <i>prefix length</i>.</li> <li>• <b>ge</b>—Matches a prefix length that is equal to or greater than the configured <i>prefix length</i>.</li> <li>• <b>le</b>—Matches a prefix length that is equal to or less than the configured <i>prefix length</i>.</li> </ul> |
| <b>Step 4</b> | switch(config)# <b>ipv6 prefix-list name [seq number] [{permit   deny} prefix {[eq prefix-length]   [ge prefix-length] [le prefix-length]}</b> }} | Creates an IPv6 prefix list or adds a prefix to an existing prefix list. The prefix length is matched as follows: <ul style="list-style-type: none"> <li>• <b>eq</b>—Matches the exact <i>prefix length</i>.</li> <li>• <b>ge</b>—Matches a prefix length that is equal to or greater than the configured <i>prefix length</i>.</li> <li>• <b>le</b>—Matches a prefix length that is equal to or less than the configured <i>prefix length</i>.</li> </ul> |



|               | Command or Action                                                   | Purpose                                                                                                                       |
|---------------|---------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 5</b> | (Optional) switch(config)# <b>show {ip   ipv6} prefix-list name</b> | Displays information about prefix lists.                                                                                      |
| <b>Step 6</b> | (Optional) switch# <b>copy running-config startup-config</b>        | Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration. |

### Example

This example shows how to create an IPv4 prefix list with two entries and apply the prefix list to a BGP neighbor:

```
switch# configure terminal
switch(config)# ip prefix-list allowprefix seq 10 permit 192.0.2.0/23 eq 24
switch(config)# ip prefix-list allowprefix seq 20 permit 209.165.201.0/27 eq 28
switch(config)# router bgp 65536:20
switch(config-router)# neighbor 192.0.2.1/16 remote-as 65535:20
switch(config-router-neighbor)# address-family ipv4 unicast
switch(config-router-neighbor-af)# prefix-list allowprefix in
```

## Configuring MAC Lists

You can configure a MAC list to permit or deny a range of MAC addresses.

### Procedure

|               | Command or Action                                                                        | Purpose                                                                                                                                                                                                                     |
|---------------|------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configure terminal</b>                                                        | Enters global configuration mode.                                                                                                                                                                                           |
| <b>Step 2</b> | switch(config)# <b>mac-list name [seq number] {permit   deny} mac-address [mac-mask]</b> | Creates a MAC list or adds a MAC address to an existing MAC list. The <i>seq</i> range is from 1 to 4294967294. The <i>mac-mask</i> specifies the portion of the MAC address to match against and is in MAC address format. |
| <b>Step 3</b> | (Optional) switch(config)# <b>show mac-list name</b>                                     | Displays information about MAC lists.                                                                                                                                                                                       |
| <b>Step 4</b> | (Optional) switch# <b>copy running-config startup-config</b>                             | Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.                                                                                               |

### Example

This example shows how to create a MAC list and copy the running configuration to the startup configuration:

```
switch# configure terminal
switch(config)# mac-list AllowMac seq 1 permit 0022.5579.a4c1 ffff.ffff.0000
switch# copy running-config startup-config
```

## Configuring AS-path Lists

You can specify an AS-path list filter on both inbound and outbound BGP routes. Each filter is an access list based on regular expressions. If the regular expression matches the representation of the AS-path attribute of the route as an ASCII string, the permit or deny condition applies.

### Procedure

|               | Command or Action                                                                           | Purpose                                                                                                                       |
|---------------|---------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configure terminal</b>                                                           | Enters global configuration mode.                                                                                             |
| <b>Step 2</b> | switch(config)# <b>ip as-path access-list</b> <i>name</i> {deny   permit} <i>expression</i> | Creates a BGP AS-path list using a regular expression.                                                                        |
| <b>Step 3</b> | (Optional) switch(config)# <b>show {ip   ipv6} as-path-access-list</b> <i>name</i>          | Displays information about as-path access lists.                                                                              |
| <b>Step 4</b> | (Optional) switch# <b>copy running-config startup-config</b>                                | Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration. |

### Example

This example shows how to create an AS-path list with two entries and apply the AS path list to a BGP neighbor:

```
switch# configure terminal
switch(config)# ip as-path access-list AllowAS permit 64510
switch(config)# ip as-path access-list AllowAS permit 64496
switch(config)# copy running-config startup-config
switch(config)# router bgp 65536:20
switch(config-router)# neighbor 192.0.2.1/16 remote-as 65535:20
switch(config-router-neighbor)# address-family ipv4 unicast
switch(config-router-neighbor-af)# filter-list AllowAS in
```

## Configuring Community Lists

You can use community lists to filter BGP routes based on the community attribute. The community number consists of a 4-byte value in the *aa:nn* format. The first two bytes represent the autonomous system number, and the last two bytes represent a user-defined network number.

When you configure multiple values in the same community list statement, all community values must match to satisfy the community list filter. When you configure multiple values in separate community list statements, the first list that matches a condition is processed.

Use community lists in a match statement to filter BGP routes based on the community attribute.

### Procedure

|               | Command or Action                 | Purpose                           |
|---------------|-----------------------------------|-----------------------------------|
| <b>Step 1</b> | switch# <b>configure terminal</b> | Enters global configuration mode. |

|               | Command or Action                                                                                                                                             | Purpose                                                                                                                                                                                                                                                                                                    |
|---------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 2</b> | switch(config)# <b>ip community-list standard</b> <i>list-name</i> {deny   permit} [ <i>community-list</i> ] [internet] [local-AS] [no-advertise] [no-export] | Creates a standard BGP community list. The <i>list-name</i> can be any case-sensitive, alphanumeric string up to 63 characters. The <i>community-list</i> can be one or more communities in the <i>aa:nn</i> format.<br><br>Do not perform this step if you need to create an expanded BGP community list. |
| <b>Step 3</b> | switch(config)# <b>ip community-list expanded</b> <i>list-name</i> {deny   permit} <i>expression</i><br><br><b>Example:</b>                                   | Creates an expanded BGP community list using a regular expression.<br><br>Do not perform this step if you need to create a standard BGP community list.                                                                                                                                                    |
| <b>Step 4</b> | (Optional) switch(config)# <b>show ip community list</b> <i>name</i>                                                                                          | Displays information about community lists.                                                                                                                                                                                                                                                                |
| <b>Step 5</b> | (Optional) switch# <b>copy running-config startup-config</b>                                                                                                  | Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.                                                                                                                                                                              |

### Example

This example shows how to create a standard community list with two entries:

```
switch# configure terminal
switch(config)# ip community-list standard BGPCommunity permit no-advertise 65536:20
switch(config)# ip community-list standard BGPCommunity permit local-AS no-export
switch(config)# copy running-config startup-config
```

## Configuring Extended Community Lists

You can use extended community lists to filter BGP routes based on the community attribute. The community number consists of a 6-byte value in the *aa4:nn* format. The first four bytes represent the autonomous system number, and the last two bytes represent a user-defined network number.

When you configure multiple values in the same extended community list statement, all extended community values must match to satisfy the extended community list filter. When you configure multiple values in separate extended community list statements, the first list that matches a condition is processed.

Use extended community lists in a match statement to filter BGP routes based on the extended community attribute.

### Procedure

|               | Command or Action                 | Purpose                           |
|---------------|-----------------------------------|-----------------------------------|
| <b>Step 1</b> | switch# <b>configure terminal</b> | Enters global configuration mode. |

|               | Command or Action                                                                                                                                                                 | Purpose                                                                                                                                                                                                                                  |
|---------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 2</b> | switch(config)# <b>ip extcommunity-list standard</b> <i>list-name</i> {deny   permit} <b>4bytegeneric</b> {transitive   nontransitive} <i>community1</i> [ <i>community2...</i> ] | Creates a standard BGP extended community list. The <i>community</i> can be one or more extended communities in the <i>aa4:nn</i> format.<br><br>Do not perform this step if you need to create an expanded BGP extended community list. |
| <b>Step 3</b> | switch(config)# <b>ip extcommunity-list expanded</b> <i>list-name</i> {deny   permit} <i>expression</i>                                                                           | Creates an expanded BGP extended community list using a regular expression.<br><br>Do not perform this step if you need to create a standard BGP extended community list.                                                                |
| <b>Step 4</b> | (Optional) switch(config)# <b>show ip extcommunity list</b> <i>name</i>                                                                                                           | Displays information about community lists.                                                                                                                                                                                              |
| <b>Step 5</b> | (Optional) switch(config)# <b>copy running-config startup-config</b>                                                                                                              | Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.                                                                                                            |

### Example

This example shows how to create a generic specific extended community list:

```
switch# configure terminal
switch(config)# ip extcommunity-list standard test1 permit 4bytegeneric transitive 65536:40
65536:60
switch(config)# copy running-config startup-config
```

## Optional Match Parameters for Route Maps

You can configure the following optional match parameters for route maps in route-map configuration mode:



**Note** The **default-information originate** command ignores **match** statements in the optional route map.

| Command                                                                                                                                      | Purpose                                                                                                                                                                                                                                                         |
|----------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| switch(config-route-map)# <b>match as-path</b> <i>name</i> [ <i>name...</i> ]                                                                | Matches against one or more AS-path lists. Create the AS-path list with the <b>ip as-path access-list</b> command.                                                                                                                                              |
| switch(config-route-map)# <b>match as-number</b> { <i>number</i> [, <i>number...</i> ]   <b>as-path-list</b> <i>name</i> [ <i>name...</i> ]} | Matches against one or more AS numbers or AS-path lists. Create the AS-path list with the <b>ip as-path access-list</b> command. The number range is from 1 to 65535. The AS-path list name can be any case-sensitive, alphanumeric string up to 63 characters. |

| Command                                                                                                                                      | Purpose                                                                                                                                                                                                                                                 |
|----------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| switch(config-route-map)# <b>match community</b> <i>name</i> [ <i>name...</i> ][ <b>exact-match</b> ]                                        | Matches against one or more community lists. Create the community list with the <b>ip community-list</b> command.                                                                                                                                       |
| switch(config-route-map)# <b>match extcommunity</b> <i>name</i> [ <i>name...</i> ][ <b>exact-match</b> ]                                     | Matches against one or more extended community lists. Create the community list with the <b>ip extcommunity-list</b> command.                                                                                                                           |
| switch(config-route-map)# <b>match interface</b> <i>interface-type number</i> [ <i>interface-type number...</i> ]                            | Matches any routes that have their next hop out one of the configured interfaces. Use ? to find a list of supported interface types.                                                                                                                    |
| switch(config-route-map)# <b>match ip address prefix-list</b> <i>name</i> [ <i>name...</i> ]                                                 | Matches against one or more IPv4 prefix lists. Use the <b>ip prefix-list</b> command to create the prefix list.                                                                                                                                         |
| switch(config-route-map)# <b>match ipv6 address prefix-list</b> <i>name</i> [ <i>name...</i> ]                                               | Matches against one or more IPv6 prefix lists. Use the <b>ipv6 prefix-list</b> command to create the prefix list.                                                                                                                                       |
| switch(config-route-map)# <b>match ip multicast</b> [ <i>source ipsource</i> ] [[ <b>group</b> <i>ipgroup</i> ] [ <b>rp</b> <i>iprp</i> ]]   | Matches an IPv4 multicast packet based on the multicast source, group, or rendezvous point.                                                                                                                                                             |
| switch(config-route-map)# <b>match ipv6 multicast</b> [ <i>source ipsource</i> ] [[ <b>group</b> <i>ipgroup</i> ] [ <b>rp</b> <i>iprp</i> ]] | Matches an IPv6 multicast packet based on the multicast source, group, or rendezvous point.                                                                                                                                                             |
| switch(config-route-map)# <b>match ip next-hop prefix-list</b> <i>name</i> [ <i>name...</i> ]                                                | Matches the IPv4 next-hop address of a route to one or more IP prefix lists. Use the <b>ip prefix-list</b> command to create the prefix list.                                                                                                           |
| switch(config-route-map)# <b>match ipv6 next-hop prefix-list</b> <i>name</i> [ <i>name...</i> ]                                              | Matches the IPv6 next-hop address of a route to one or more IP prefix lists. Use the <b>ipv6 prefix-list</b> command to create the prefix list.                                                                                                         |
| switch(config-route-map)# <b>match ip route-source prefix-list</b> <i>name</i> [ <i>name...</i> ]                                            | Matches the IPv4 route source address of a route to one or more IP prefix lists. Use the <b>ip prefix-list</b> command to create the prefix list.                                                                                                       |
| switch(config-route-map)# <b>match ipv6 route-source prefix-list</b> <i>name</i> [ <i>name...</i> ]                                          | Matches the IPv6 route-source address of a route to one or more IP prefix lists. Use the <b>ipv6 prefix-list</b> command to create the prefix list.                                                                                                     |
| switch(config-route-map)# <b>match mac-list</b> <i>name</i> [ <i>name...</i> ]                                                               | Matches against one or more MAC lists. Use the <b>mac-list</b> command to create the MAC list.                                                                                                                                                          |
| switch(config-route-map)# <b>match metric</b> <i>value</i> [+ <i>deviation</i> .] [ <i>value..</i> ]                                         | Matches the route metric against one or more metric values or value ranges. Use +- deviation argument to set a metric range. The route map matches any route metric that falls the range:<br><br><i>value - deviation</i> to <i>value + deviation</i> . |

| Command                                                                          | Purpose                                                                                                                                                                                                                                                                                                                                      |
|----------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| switch(config-route-map)# <b>match route-type</b> <i>route-type</i>              | Matches against a type of route. The <i>route-type</i> can be one or more of the following: <ul style="list-style-type: none"> <li>• external</li> <li>• inter-area</li> <li>• internal</li> <li>• intra-area</li> <li>• level-1</li> <li>• level-2</li> <li>• local</li> <li>• nssa-external</li> <li>• type-1</li> <li>• type-2</li> </ul> |
| switch(config-route-map)# <b>match tag</b> <i>tagid</i> [ <i>tagid...</i> ]      | Matches a route against one or more tags for filtering or redistribution.                                                                                                                                                                                                                                                                    |
| switch(config-route-map)# <b>match vlan</b> <i>vlan-id</i> [ <i>vlan-range</i> ] | Matches against a VLAN.                                                                                                                                                                                                                                                                                                                      |

## Optional Set Parameters for Route Maps

You can configure the following optional set parameters for route maps in route-map configuration mode:

| Command                                                                                                                                       | Purpose                                                                                                                                                                       |
|-----------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| switch(config-route-map)# <b>set as-path</b> { <b>tag</b>   <b>prepend</b> { <b>last-as</b> <i>number</i>   <i>as-1</i> [ <i>as-2...</i> ]} } | Modifies an AS-path attribute for a BGP route. You can prepend the configured number of last AS numbers or a string of particular AS-path values ( <i>as-1 as-2...as-n</i> ). |
| switch(config-route-map)# <b>set comm-list</b> <i>name</i> <b>delete</b>                                                                      | Removes communities from the community attribute of an inbound or outbound BGP route update. Use the <b>ip community-list</b> command to create the community list.           |

| Command                                                                                                                                                                                                   | Purpose                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <pre>switch(config-route-map)# <b>set community</b> {none   <b>additive</b>   <b>local-AS</b>   <b>no-advertise</b>   <b>no-export</b>   community-1 [community-2...]}</pre>                              | <p>Sets the community attribute for a BGP route update.</p> <p><b>Note</b> When you use both the <b>set community</b> and <b>set comm-list delete</b> commands in the same sequence of a route map attribute, the deletion operation is performed before the set operation.</p> <p><b>Note</b> Use the <b>send-community</b> command in BGP neighbor address family configuration mode to propagate BGP community attributes to BGP peers.</p>                         |
| <pre>switch(config-route-map)# <b>set dampening</b> <i>halflife</i> <i>reuse suppress duration</i></pre>                                                                                                  | <p>Sets the following BGP route dampening parameters:</p> <ul style="list-style-type: none"> <li>• <i>halflife</i>—The range is from 1 to 45 minutes. The default is 15.</li> <li>• <i>reuse</i>—The range is from is 1 to 20000 seconds. The default is 750.</li> <li>• <i>suppress</i>—The range is from is 1 to 20000. The default is 2000.</li> <li>• <i>duration</i>—The range is from is 1 to 255 minutes. The default is 60.</li> </ul>                         |
| <pre>switch(config-route-map)# <b>set distance</b> <i>value</i></pre>                                                                                                                                     | <p>Sets the administrative distance of routes for OSPFv2 or OSPFv3. The range is from 1 to 255.</p>                                                                                                                                                                                                                                                                                                                                                                    |
| <pre>switch(config-route-map)# <b>set extcomm-list</b> <i>name</i> <b>delete</b></pre>                                                                                                                    | <p>Removes communities from the extended community attribute of an inbound or outbound BGP route update. Use the <b>ip extcommunity-list</b> command to create the extended community list.</p>                                                                                                                                                                                                                                                                        |
| <pre>switch(config-route-map)# <b>set extcommunity</b> <b>4byteas-generic</b> {<b>transitive</b>   <b>nontransitive</b>} {<b>none</b>   <b>additive</b>} <i>community-1</i> [<i>community-2...</i>]</pre> | <p>Sets the extended community attribute for a BGP route update.</p> <p><b>Note</b> When you use both the <b>set extcommunity</b> and <b>set extcomm-list delete</b> commands in the same sequence of a route map attribute, the deletion operation is performed before the set operation.</p> <p><b>Note</b> Use the <b>send-community</b> command in BGP neighbor address family configuration mode to propagate BGP extended community attributes to BGP peers.</p> |

| Command                                                                                                                                           | Purpose                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|---------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| switch(config-route-map)# <b>set extcommunity cost</b> <i>community-id1 cost</i> [ <b>igp</b>   <b>pre-bestpath</b> ] [ <b>community-id2...</b> ] | <p>Sets the cost community attribute for a BGP route update. This attribute allows you to customize the BGP best path selection process for a local autonomous system or confederation. The <i>community-id</i> range is from 0 to 255. The <i>cost</i> range is from 0 to 4294967295. The path with the lowest cost is preferred. For paths with equal cost, the path with the lowest community ID is preferred.</p> <p>The <b>igp</b> keyword compares the cost after the IGP cost comparison. The <b>pre-bestpath</b> keyword compares before all other steps in the bestpath algorithm.</p> |
| switch(config-route-map)# <b>set extcommunity rt</b> <i>community-1</i> [ <b>additive</b> ] [ <i>community-2...</i> ]                             | <p>Sets the extended community route target attribute for a BGP route update. The <i>community</i> value can be a 2-byte AS number:4-byte network number, a 4-byte AS number:2-byte network number, or an IP address:2-byte network number.</p> <p>Use the <b>additive</b> keyword to add a route target to an existing extended community route target attribute.</p>                                                                                                                                                                                                                          |
| switch(config-route-map)# <b>set forwarding-address</b>                                                                                           | Sets the forwarding address for OSPF.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| switch(config-route-map)# <b>set level</b> { <b>backbone</b>   <b>level-1</b>   <b>level-1-2</b>   <b>level-2</b> }                               | Sets what area to import routes to for IS-IS. The options for IS-IS are level-1, level-1-2, or level-2. The default is level-1.                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| switch(config-route-map)# <b>set local-preference</b> <i>value</i>                                                                                | Sets the BGP local preference value. The range is from 0 to 4294967295.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| switch(config-route-map)# <b>set metric</b> [+   -] <i>bandwidth-metric</i>                                                                       | Adds or subtracts from the existing metric value. The metric is in Kb/s. The range is from 0 to 4294967295.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| switch(config-route-map)# <b>set metric</b> <i>bandwidth</i> [ <i>delay reliability load mtu</i> ]                                                | <p>Sets the route metric values. Metrics are as follows:</p> <ul style="list-style-type: none"> <li>• <i>metric0</i>—Bandwidth in Kb/s. The range is from 0 to 4294967295.</li> <li>• <i>metric1</i>—Delay in 10-microsecond units.</li> <li>• <i>metric2</i>—Reliability. The range is from 0 to 255 (100 percent reliable).</li> <li>• <i>metric3</i>—Loading. The range is from 1 to 200 (100 percent loaded).</li> <li>• <i>metric4</i>—MTU of the path. The range is from 1 to 4294967295.</li> </ul>                                                                                      |



| Command                                                                                                                | Purpose                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| switch(config-route-map)# <b>set metric-type</b> { <b>external</b>   <b>internal</b>   <b>type-1</b>   <b>type-2</b> } | <p>Sets the metric type for the destination routing protocol. The options are as follows:</p> <ul style="list-style-type: none"> <li>• <b>external</b>—IS-IS external metric</li> <li>• <b>internal</b>— IGP metric as the MED for BGP</li> <li>• <b>type-1</b>—OSPF external type 1 metric</li> <li>• <b>type-2</b>—OSPF external type 2 metric</li> </ul> <p>The <b>set metric-type internal</b> command affects an outgoing policy and an eBGP neighbor only. If you configure both the <b>metric</b> and <b>metric-type internal</b> commands in the same BGP peer outgoing policy, then Cisco NX-OS ignores the <b>metric-type internal</b> command.</p> |
| switch(config-route-map)# <b>set nssa-only</b>                                                                         | Sets Type-7 LSA generated on ASBR with no P bit set. This prevents Type-7 to Type-5 LSA translation in OSPF.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| switch(config-route-map)# <b>set origin</b> { <b>egp as-number</b>   <b>igp</b>   <b>incomplete</b> }                  | Sets the BGP origin attribute. The EGP <i>as-number</i> range is from 0 to 65535.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| switch(config-route-map)# <b>set tag name</b>                                                                          | Sets the tag value for the destination routing protocol. The <i>name</i> parameter is an unsigned integer.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| switch(config-route-map)# <b>set weight count</b>                                                                      | Sets the weight for the BGP route. The range is from 0 to 65535.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |

## Verifying the Route Policy Manager Configuration

Use one of the following commands to verify the configuration:

| Command                                             | Purpose                                                 |
|-----------------------------------------------------|---------------------------------------------------------|
| <b>show ip community-list</b> [ <i>name</i> ]       | Displays information about a community list.            |
| <b>show ip extcommunity-list</b> [ <i>name</i> ]    | Displays information about an extended community list.  |
| <b>show [ip   ipv6] prefix-list</b> [ <i>name</i> ] | Displays information about an IPv4 or IPv6 prefix list. |
| <b>show route-map</b> [ <i>name</i> ]               | Displays information about a route map.                 |

## Configuration Examples for Route Policy Manager

This example shows how to use an address family to configure Route Policy Manager so that any unicast and multicast routes from neighbor 209.0.2.1 are accepted if they match prefix-list AllowPrefix:

```
router bgp 64496

neighbor 209.0.2.1 remote-as 64497
 address-family ipv4 unicast
 route-map filterBGP in

route-map filterBGP
 match ip address prefix-list AllowPrefix

ip prefix-list AllowPrefix 10 permit 192.0.2.0/24
ip prefix-list AllowPrefix 20 permit 209.165.201.0/27
```

## Related Documents for Route Policy Manager

| Related Topic                     | Document Title                                                                  |
|-----------------------------------|---------------------------------------------------------------------------------|
| Route Policy Manager CLI commands | <i>Cisco Nexus 7000 Series NX-OS Unicast Routing Command Reference</i>          |
| VDCs and VRFs                     | <i>Cisco Nexus 7000 Series NX-OS Virtual Device Context Configuration Guide</i> |

## Standards for Route Policy Manager

No new or modified standards are supported by this feature, and support for existing standards has not been modified by this feature.

## Feature History for Route Policy Manager

This table includes only the updates for those releases that have resulted in additions or changes to the feature.

| Feature Name                              | Releases | Feature Information                                          |
|-------------------------------------------|----------|--------------------------------------------------------------|
| Multiple match statements under table-map | 6.2(14)  | Added support for multiple match statements under table-map. |
| Route map support matrix                  | 6.2(2)   | Added the route map support matrix for routing protocols.    |

| Feature Name                  | Releases | Feature Information                                                                                                                                                                                                                                                                                                         |
|-------------------------------|----------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Match interfaces              | 6.2(2)   | <p>Added support for null interfaces to the <b>match interface</b> command.</p> <p>Added support for the following set and match statements in a route map for the EIGRP distribute list:</p> <ul style="list-style-type: none"> <li>• Set routing protocol metric</li> <li>• Set route tag</li> <li>• Match tag</li> </ul> |
| Route policy manager          | 6.1(1)   | Added support for the <b>set distance</b> command and for the <b>inter-area</b> and <b>intra-area</b> options for the <b>match route-type</b> command.                                                                                                                                                                      |
| MPLS set clauses              | 5.2(1)   | Added support for <b>set extcommunity cost</b> , <b>set extcommunity rt</b> , and <b>set nssa-only</b> commands.                                                                                                                                                                                                            |
| MAC lists , metric, and VLANs | 5.0(2)   | Added support for the <b>match mac-list</b> , <b>match metric</b> , and <b>match vlan</b> commands.                                                                                                                                                                                                                         |
| Extended community lists      | 4.2(1)   | Added support for generic specific extended community lists.                                                                                                                                                                                                                                                                |
| Match interfaces              | 4.1(2)   | Added support to match a list of interfaces in a route map.                                                                                                                                                                                                                                                                 |
| Match AS numbers              | 4.1(2)   | Added support to match a range of AS numbers in a route map.                                                                                                                                                                                                                                                                |
| Route policy manager          | 4.0(1)   | This feature was introduced.                                                                                                                                                                                                                                                                                                |





## CHAPTER 19

# Configuring Policy-Based Routing

This chapter contains the following sections:

- [Finding Feature Information, on page 419](#)
- [Information About Policy Based Routing, on page 419](#)
- [Prerequisites for Policy-Based Routing, on page 422](#)
- [Guidelines and Limitations for Policy-Based Routing, on page 422](#)
- [Default Settings for Policy-Based Routing, on page 424](#)
- [Configuring Policy-Based Routing, on page 424](#)
- [Verifying the Policy-Based Routing Configuration, on page 430](#)
- [Configuration Examples for Policy Based-Routing, on page 430](#)
- [Related Documents for Policy-Based Routing, on page 431](#)
- [Standards for Policy-Based Routing, on page 431](#)
- [Feature History for Policy-Based Routing, on page 431](#)

## Finding Feature Information

Your software release might not support all the features documented in this module. For the latest caveats and feature information, see the Bug Search Tool at <https://tools.cisco.com/bugsearch/> and the release notes for your software release. To find information about the features documented in this module, and to see a list of the releases in which each feature is supported, see the "New and Changed Information" chapter or the Feature History table in this chapter.

## Information About Policy Based Routing

Policy-based routing allows you to configure a defined policy for IPv4 and IPv6 traffic flows, lessening reliance on routes derived from routing protocols. All packets received on an interface with policy-based routing enabled are passed through enhanced packet filters or route maps. The route maps dictate the policy, determining where to forward packets.

Route maps are composed of match and set statements that you can mark as permit or deny. You can interpret the statements as follows:

- If the packets match any route map statements, all the set statements are applied. One of these actions involves choosing the next hop.

- If a statement is marked as deny, the packets that meet the match criteria are sent back through the normal forwarding channels and destination-based routing is performed.
- If the statement is marked as permit and the packets do not match any route-map statements, the packets are sent back through the normal forwarding channels and destination-based routing is performed.

Policy-based routing includes the following features:

- Source-based routing—Routes traffic that originates from different sets of users through different connections across the policy routers.
- Quality of Service (QoS)—Differentiates traffic by setting the precedence or type of service (ToS) values in the IP packet headers at the periphery of the network and leveraging queuing mechanisms to prioritize traffic in the core or backbone of the network (see the *Cisco Nexus 7000 Series NX-OS Quality of Service Configuration Guide*).
- Load sharing—Distributes traffic among multiple paths based on the traffic characteristics.

## Policy Route Maps

Each entry in a route map contains a combination of match and set statements. The match statements define the criteria for whether appropriate packets meet the particular policy (that is, the conditions to be met). The set clauses explain how the packets should be routed once they have met the match criteria.

You can mark the route-map statements as permit or deny. If the statement is marked as a deny, the packets that meet the match criteria are sent back through the normal forwarding channels (destination-based routing is performed). If the statement is marked as permit and the packets meet the match criteria, all the set clauses are applied. If the statement is marked as permit and the packets do not meet the match criteria, those packets are also forwarded through the normal routing channel.



---

**Note** Policy routing is specified on the interface that receives the packets, not on the interface from which the packets are sent.

---

## Set Criteria for Policy-Based Routing

The set criteria in a route map is evaluated in the order listed in the route map. Set criteria specific to route maps used for policy-based routing are as follows:

1. List of interfaces through which the packets can be routed—If more than one interface is specified, the first interface that is found to be up is used for forwarding the packets.
2. List of specified IP addresses—The IP address can specify the adjacent next-hop router in the path toward the destination to which the packets should be forwarded. The first IP address associated with a currently up connected interface is used to route the packets.



---

**Note** You can optionally configure the set criteria for next-hop addresses to load balance traffic across up to 16 IP addresses. In this case, Cisco NX-OS sends all traffic for each IP flow to a particular IP next-hop address.

---

3. List of default interfaces—If there is no explicit route available to the destination address of the packet being considered for policy routing, the route map routes it to the first up interface in the list of specified default interfaces.
4. List of default next-hop IP addresses—Route to the interface or the next-hop address specified by this set statement only if there is no explicit route for the destination address of the packet in the routing table.



**Note** You can optionally configure the set criteria for the default next-hop addresses to load balance traffic across a maximum of 16 IP addresses. In this case, Cisco NX-OS sends all traffic for each IP flow to a particular IP next-hop address.

If the packets do not meet any of the defined match criteria, those packets are routed through the normal destination-based routing process.

## Local Policy Routing

Local policy routing allows you to apply a route map to local (device-generated) traffic. All packets originating on the device that are not normally policy routed are subject to local policy routing.

## Route Map Support Matrix for Policy-Based Routing

The following tables include the configurable match and set statements for policy-based routing on Cisco Nexus 70xx and 77xx Series switches running the latest shipping release. For specific release information, see the [Feature History for Policy-Based Routing, on page 431](#).

The following legend applies to the tables:

- Yes—The statement is supported for policy-based routing.
- No—The statement is not supported for policy-based routing.
- If a statement does not apply for policy-based routing, there is an em dash (—) in the column next to the statement.
- Where clarification is required, information is added in the appropriate row/column.

**Table 37: SET Route Map Statements for Policy-Based Routing**

| SET Route Map Statement           | Policy-Based Routing (PBR) |
|-----------------------------------|----------------------------|
| IPv4 Next Hop                     | Yes                        |
| IPv6 Next Hop                     | Yes                        |
| Default IPv4 Next Hop             | Yes                        |
| Default IPv6 Next Hop             | Yes                        |
| IPv4 Next Hop Verify Availability | Yes                        |
| IPv6 Next Hop Verify Availability | Yes                        |

| SET Route Map Statement                   | Policy-Based Routing (PBR) |
|-------------------------------------------|----------------------------|
| Default IPv4 Next Hop Verify Availability | Yes                        |
| Default IPv6 Next Hop Verify Availability | Yes                        |
| Interface null0                           | Yes                        |
| VRF                                       | Yes                        |
| IPv4 Precedence                           | Yes                        |
| IPv6 Precedence                           | Yes                        |
| Interface, GRE Ethernet                   | No                         |

**Table 38: MATCH Route Map Statements for Policy-Based Routing**

| MATCH Route Map Statement | Policy-Based Routing (PBR) |
|---------------------------|----------------------------|
| Tag                       | Yes                        |
| Packet Length             | Yes                        |
| VLAN ID                   | Yes                        |
| MAC ACL                   | Yes                        |
| IPv4 Prefix List          | No                         |
| IPv6 Prefix List          | No                         |
| IP ACL                    | Yes                        |

## Prerequisites for Policy-Based Routing

Policy-based routing has the following prerequisites:

- Install the correct license.
- You must enable policy-based routing.
- Assign an IP address on the interface and bring the interface up before you apply a route map on the interface for policy-based routing.
- If you configure VDCs, install the Advanced Services license and enter the desired VDC (see the *Cisco Nexus 7000 Series NX-OS Virtual Device Context Configuration Guide*).

## Guidelines and Limitations for Policy-Based Routing

Policy-based routing has the following configuration guidelines and limitations:



- If incoming traffic has an mpls label, then PRB will not work for such traffic.
- Cisco NX-OS uses recursive next hops. You do not need to enter any commands for recursive next hops like you do for Cisco IOS.
- A policy-based routing route map can have only one match or set statement per route-map statement.
- A **match** command cannot refer to more than one ACL in a route map used for policy-based routing.
- The same route map can be shared among different interfaces for policy-based routing as long as the interfaces belong to the same virtual routing and forwarding (VRF) instance.
- Prior to Cisco NX-OS Release 8.0(1) setting a tunnel interface or an IP address via a tunnel interface as a next hop in a policy-based routing policy is not supported. Applying policy-based routing or **ip policy route-map** on tunnel interfaces is also not supported. From Cisco NX-OS Release 8.0(1) onwards GRE next hop is supported on policy-based routing.
- Policy-based routing is not supported with inbound traffic on FEX ports.
- Using a prefix-list as a match criteria is not supported. Do not use a prefix-list in a policy-based routing route-map.
- Beginning with Cisco NX-OS Release 5.2(4), policy-based routing and WCCPv2 are supported on the same interface. However, policy-based routing with statistics and WCCPv2 is supported on the same interface only if bank chaining is disabled.
- Beginning with Cisco NX-OS Release 6.1(3), you can configure the device to support deny access control entries (ACEs) in a sequence for the following sequence-based features: VACLs and QoS. For more information, see the “Configuring VLAN ACLs” chapter in the *Cisco Nexus 7000 Series NX-OS Security Configuration Guide*.
- If you are familiar with the Cisco IOS CLI, be aware that the Cisco NX-OS commands for this feature might differ from the Cisco IOS commands that you would use.
- PBR marks the next-hop as down even when the next hop and the corresponding tracks are up. This issue is due to the RPM that does not effectively process the tracks.

Currently the object tracking manager (OTM) does not support forward referencing for track objects. Track objects must be created in the OTM before they are used in any configuration.

Perform the following steps to configure the track objects with RPM or PBR so that the PBR next-hop issue does not occur:

1. Create the track object in OTM using the **track <object id>** command.
2. Use the configured track object in a route map using the **set ip next-hop verify-availability <ip1> track <object id>** command.
3. Apply the route map to an interface using the **ip policy route-map <map-name>** command.

# Default Settings for Policy-Based Routing

Table 39: Default Policy-Based Routing Parameters

| Parameters           | Default  |
|----------------------|----------|
| Policy-based routing | Disabled |

## Configuring Policy-Based Routing

### Enabling the Policy-Based Routing

You must enable the policy-based routing feature before you can configure a route policy.

#### Before you begin

Ensure that you are in the correct VDC (or use the **switchto vdc** command).

#### Procedure

|               | Command or Action                                                    | Purpose                                                                                                                                                              |
|---------------|----------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configure terminal</b>                                    | Enters global configuration mode.                                                                                                                                    |
| <b>Step 2</b> | switch(config)# <b>[no] feature pbr</b>                              | Enables the policy-based routing feature. Use the <b>no feature pbr</b> command to disable the policy-based routing feature and remove all associated configuration. |
| <b>Step 3</b> | (Optional) switch(config)# <b>show feature</b>                       | Displays enabled and disabled features.                                                                                                                              |
| <b>Step 4</b> | (Optional) switch(config)# <b>copy running-config startup-config</b> | Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.                                        |

### Configuring a Route Policy

You can use route maps in policy-based routing to assign routing policies to the inbound interface.

Cisco NX-OS routes the packet as soon as it finds a next hop and an interface.

#### Procedure

|               | Command or Action                               | Purpose                              |
|---------------|-------------------------------------------------|--------------------------------------|
| <b>Step 1</b> | switch# <b>configure terminal</b>               | Enters global configuration mode.    |
| <b>Step 2</b> | switch(config)# <b>interface type slot/port</b> | Enters interface configuration mode. |

|               | Command or Action                                                    | Purpose                                                                                                                       |
|---------------|----------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 3</b> | switch(config-if)# <b>ip policy route-map</b><br><i>map-name</i>     | Assigns a route map for IPv4 policy-based routing to the interface.                                                           |
| <b>Step 4</b> | switch(config-if)# <b>ipv6 policy route-map</b><br><i>map-name</i>   | Assigns a route map for IPv6 policy-based routing to the interface.                                                           |
| <b>Step 5</b> | (Optional) switch(config-route-map)# <b>end</b>                      | Exits route-map configuration mode and enters the privileged executive mode.                                                  |
| <b>Step 6</b> | (Optional) switch(config)# <b>copy running-config startup-config</b> | Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration. |

### Example

This example shows how to add a route map to an interface:

```
switch# configure terminal
switch(config)# interface ethernet 1/2
switch(config-if)# ip policy route-map Testmap
switch(config)# exit
switch(config)# copy running-config startup-config
```

You can configure the following optional match parameters for route maps in route-map configuration mode:

| Command                                                                                                                                                                          | Purpose                                                                                                                                                                        |
|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>match ip address access-list-name</b> <i>name [name...]</i><br>Example:<br><pre>switch(config-route-map)#<br/><b>match ip address access-list-name ACL1</b></pre>             | Matches an IPv4 address against one or more IP access control lists (ACLs). This command is used for policy-based routing and is ignored by route filtering or redistribution. |
| <b>match ipv6 address access-list-name</b> <i>name [name...]</i><br>Example:<br><pre>switch(config-route-map)# <b>match</b><br/><b>ipv6 address access-list-name ACLv6</b></pre> | Matches an IPv6 address against one or more IPv6 ACLs. This command is used for policy-based routing and is ignored by route filtering or redistribution.                      |
| <b>match length</b> <i>min max</i><br>Example:<br><pre>switch(config-route-map)# <b>match length 64 1500</b></pre>                                                               | Matches against the length of the packet. This command is used for policy-based routing.                                                                                       |
| <b>match mac-list</b> <i>maclist [...maclist]</i>                                                                                                                                | Matches against a list of MAC addresses. This command is used for policy-based routing.                                                                                        |

| Command                                                                                                                                                                                                                                                                                                       | Purpose                                                                                     |
|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------|
| <p><b>match metric</b> <i>metric-value</i> [<i>+ - deviation-number</i>]<br/>[...<i>metric-value</i> [<i>+ - deviation-number</i>]] [<i>+ - deviation-number</i>]<br/>[... <i>metric-value</i> [<i>+ - deviation-number</i>]]</p> <p>Example:</p> <pre>switch(config-route-map)# <b>match metric</b> 10</pre> | Matches against the routing protocol metric. This command is used for policy-based routing. |
| <p><b>match vlan</b> <i>vlan-range</i></p> <pre>switch(config-route-map)# <b>match vlan</b> 64</pre>                                                                                                                                                                                                          | Matches against the VLAN ID of the packet. This command is used for policy-based routing.   |

You can configure the following optional set parameters for route maps in route-map configuration mode:

| Command                                                                                                                                                                                                                                                | Purpose                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <p><b>set ip next-hop</b> <i>address2</i> [<i>address2...</i>] [<b>load-share</b>   <b>peer-address</b>   <b>unchanged</b>   <b>verify-availability</b>]</p> <p>Example:</p> <pre>switch(config-route-map)#<br/><b>set ip next-hop</b> 192.0.2.1</pre> | <p>Sets the IPv4 next-hop address for policy-based routing. This command uses the first valid next-hop address if multiple addresses are configured. This can be done with next-hop tracking only.</p> <ul style="list-style-type: none"> <li>• Use the optional <b>load-share</b> keyword to load balance traffic across a maximum of 16 next-hop addresses.</li> <li>• Use the optional <b>peer-address</b> keyword to the next hop to be the Border Gateway Protocol (BGP) peering address.</li> <li>• Use the optional <b>unchanged</b> keyword to specify that the next-hop attribute in the BGP update to the eBGP peer is unmodified.</li> <li>• Use the optional <b>verify-availability</b> keyword to verify the reachability of the tracked object.</li> </ul> |

| Command                                                                                                                                                                                                                                             | Purpose                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <p><b>set ip default next-hop</b> <i>address2</i> [<i>address2...</i>]<br/>[<b>load-share</b>   <b>verify-availability</b>]</p> <p>Example:</p> <pre>switch(config-route-map) # set ip default next-hop 192.0.2.2</pre>                             | <p>Sets the IPv4 next-hop address for policy-based routing when there is no explicit route to a destination. This command uses the first valid next-hop address if multiple addresses are configured. This can done with next-hop tracking only.</p> <ul style="list-style-type: none"> <li>• Use the optional <b>load-share</b> keyword to load balance traffic across a maximum of 16 next-hop addresses.</li> <li>• Use the optional <b>verify-availability</b> keyword to verify the reachability of the tracked object.</li> </ul> <p><b>Note</b> For software-forwarded traffic, the route that is present in the unicast routing table (of the VRF in which packet was received) for packet-specified destination takes preference over what is specified in <b>set ip default next-hop</b> command, when there is condition match. Even if there is a default route present in the VRF, that default route overrides what is set in the command. This applies to software-forwarded traffic only.</p> |
| <p><b>set ipv6 next-hop</b> <i>address2</i> [<i>address2...</i>] [<b>load-share</b>   <b>peer-address</b>   <b>unchanged</b>   <b>verify-availability</b>]</p> <p>Example:</p> <pre>switch(config-route-map) # set ipv6 next-hop 2001:0DB8::1</pre> | <p>Sets the IPv6 next-hop address for policy-based routing. This command uses the first valid next-hop address if multiple addresses are configured. This can done with next-hop tracking only.</p> <ul style="list-style-type: none"> <li>• Use the optional <b>load-share</b> keyword to load balance traffic across a maximum of 16 next-hop addresses.</li> <li>• Use the optional <b>peer-address</b> keyword to the next hop to be the Border Gateway Protocol (BGP) peering address.</li> <li>• Use the optional <b>unchanged</b> keyword to specify that the next-hop attribute in the BGP update to the eBGP peer is unmodified.</li> <li>• Use the optional <b>verify-availability</b> keyword to verify the reachability of the tracked object.</li> </ul>                                                                                                                                                                                                                                         |

| Command                                                                                                                                                                                                                       | Purpose                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <p><b>set ipv6 default next-hop</b> <i>address2</i> [<i>address2...</i>]<br/>[<b>load-share</b>   <b>verify-availability</b>]</p> <p>Example:</p> <pre>switch(config-route-map)# set ipv6 default next-hop 2001:0DB8::2</pre> | <p>Sets the IPv6 next-hop address for policy-based routing when there is no explicit route to a destination. This command uses the first valid next-hop address if multiple addresses are configured. This can be done with next-hop tracking only.</p> <ul style="list-style-type: none"> <li>• Use the optional <b>load-share</b> keyword to load balance traffic across a maximum of 16 next-hop addresses.</li> <li>• Use the optional <b>verify-availability</b> keyword to verify the reachability of the tracked object.</li> </ul> |
| <p><b>set ip precedence</b> <i>precedence-value</i></p> <p>Example:</p> <pre>switch(config-route-map)# set ip precedence highv4</pre>                                                                                         | Sets the precedence value in the IPv4 packet header.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <p><b>set ipv6 precedence</b> <i>precedence-value</i></p> <p>Example:</p> <pre>switch(config-route-map)# set ipv6 precedence highv6</pre>                                                                                     | Sets the precedence value in the IPv6 packet header.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <p><b>set ipv6 precedence address prefix-list</b><br/><i>prefix-list-name</i></p> <p>Example:</p> <pre>switch(config-route-map)# set ipv6 precedence address prefix-list acl1</pre>                                           | Sets the IPv6 map routes to be injected.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <p><b>set interface</b> {<i>null10</i>   <i>tunnel-te</i>}</p> <p>Example:</p> <pre>switch(config-route-map)# set interface null10</pre>                                                                                      | Sets the interface used for routing. Use the <b>null0</b> interface to drop packets. Use the <b>tunnel-te</b> interface to forward packets on the MPLS TE tunnel.                                                                                                                                                                                                                                                                                                                                                                          |
| <p><b>set vrf</b> <i>vrf-name</i></p> <p>Example:</p> <pre>switch(config-route-map)# set vrf MainVRF</pre>                                                                                                                    | Sets the VRF for next-hop resolution.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |

## Configuring Local Policy Routing

You can enable local policy routing for packets generated by the device and specify which route map the device should use.

**Procedure**

|               | <b>Command or Action</b>                                             | <b>Purpose</b>                                                                                                                |
|---------------|----------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configure terminal</b>                                    | Enters global configuration mode.                                                                                             |
| <b>Step 2</b> | switch(config)# <b>{ip   ipv6} local policy route-map map-name</b>   | Configures IPv4 or IPv6 local policy route maps for packets generated by the device.                                          |
| <b>Step 3</b> | (Optional) <b>show {ip   ipv6} local policy</b>                      | Displays the route map used for IPv4 or IPv6 local policy routing.                                                            |
| <b>Step 4</b> | (Optional) switch(config)# <b>copy running-config startup-config</b> | Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration. |

## Configuring a Deny ACE

Beginning with Cisco NX-OS Release 6.1(3), you can configure the device to support deny access control entries (ACEs) in a sequence for the following sequence-based features: VACL and Quality of service (QoS).

When deny ACEs are enabled, the traffic that matches a deny ACE (an ACL rule with the **deny** keyword) in a class-map-acl is recursively matched against subsequent class-map-acls until it hits a permit ACE.



**Note** In earlier releases, an ACL used in a policy-based routing route map cannot include a deny statement.

**Before you begin**

Ensure that you are in the default or admin VDC.

**Procedure**

|               | <b>Command or Action</b>                                             | <b>Purpose</b>                                                                                                                |
|---------------|----------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configure terminal</b>                                    | Enters global configuration mode.                                                                                             |
| <b>Step 2</b> | switch(config)# <b>[no] hardware access-list allow deny ace</b>      | Enables deny ACEs in a sequence.<br>The <b>no</b> form of the command disables deny ACEs.                                     |
| <b>Step 3</b> | (Optional) <b>show running-config aclmgr</b>                         | Displays the ACL configuration.                                                                                               |
| <b>Step 4</b> | (Optional) switch(config)# <b>copy running-config startup-config</b> | Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration. |

## Verifying the Policy-Based Routing Configuration

To display policy-based routing configuration information, perform one of the following tasks:

| Command                                         | Purpose                                                            |
|-------------------------------------------------|--------------------------------------------------------------------|
| <b>show [ip   ipv6] policy [name]</b>           | Displays information about an IPv4 or IPv6 policy.                 |
| <b>show {ip   ipv6} local policy [vrf name]</b> | Displays the route map used for IPv4 or IPv6 local policy routing. |
| <b>show route-map [name] pbr-statistics</b>     | Displays policy statistics.                                        |

Use the **route-map map-name pbr-statistics** to enable policy statistics. Use the **clear route-map map-name pbr-statistics** to clear these policy statistics.

## Configuration Examples for Policy Based-Routing

This example shows how to configure a simple route policy on an interface:

```
feature pbr
ip access-list pbr-sample
 permit tcp host 10.1.1.1 host 192.168.2.1 eq 80
!
route-map pbr-sample
 match ip address pbr-sample
 set ip next-hop 192.168.1.1
!
route-map pbr-sample pbr-statistics

interface ethernet 1/2
 ip policy route-map pbr-sample
```

The following output verifies this configuration:

```
n7000# show route-map pbr-sample

route-map pbr-sample, permit, sequence 10
Match clauses:
 ip address (access-lists): pbr-sample
Set clauses:
 ip next-hop 192.168.1.1

n7000# show route-map pbr-sample pbr-statistics

route-map pbr-sample, permit, sequence 10
Policy routing matches: 84 packets

Default routing: 233 packets
```

## Configuration Example for Local Policy Routing

The following example sends packets with a destination IP address matching that allowed by extended access list 131 to the router at IP address 172.30.3.20:



```

ip local policy route-map xyz
!
route-map xyz
match ip address 131
set ip next-hop 172.30.3.20

```

## Related Documents for Policy-Based Routing

| Related Topic                     | Document Title                                                                |
|-----------------------------------|-------------------------------------------------------------------------------|
| Policy-based routing CLI commands | <i>Cisco Nexus 7000 Series NX-OS Unicast Routing Command Reference</i>        |
| VDCs and VRFs                     | <i>Cisco Nexus 7000 Series NX-OS Virtual Device Context Command Reference</i> |

## Standards for Policy-Based Routing

No new or modified standards are supported by this feature, and support for existing standards has not been modified by this feature.

## Feature History for Policy-Based Routing

This table includes only the updates for those releases that have resulted in additions or changes to the feature.

**Table 40: Feature History for Policy-Based Routing**

| Feature Name             | Release | Feature Information                                                                                                                                 |
|--------------------------|---------|-----------------------------------------------------------------------------------------------------------------------------------------------------|
| Route map support matrix | 6.2(2)  | Added the route map support matrix for policy-based routing.                                                                                        |
| Policy-based routing     | 6.1(3)  | Added support for deny access control entries (ACEs) in a sequence for the following sequence-based features: VACLs, policy-based routing, and QoS. |
| Policy-based routing     | 5.2(4)  | Added support for policy-based routing and WCCPv2 on the same interface if bank chaining is disabled.                                               |
| Interfaces               | 5.2(1)  | Added support for <b>set interface route-map</b> command.                                                                                           |
| IPv6 policies            | 4.2(1)  | Added support for IPv6 policies.                                                                                                                    |
| Policy-based routing     | 4.0(1)  | This feature was introduced.                                                                                                                        |





## PART **III**

# First-Hop Redundancy Protocols

- [Configuring GLBP, on page 435](#)
- [Configuring HSRP, on page 451](#)
- [Configuring VRRP, on page 483](#)
- [Configuring Object Tracking, on page 507](#)





## CHAPTER 20

# Configuring GLBP

---

This chapter contains the following sections:

- [Finding Feature Information, on page 435](#)
- [Information About GLBP, on page 435](#)
- [Prerequisites for GLBP, on page 440](#)
- [Guidelines and Limitations for GLBP, on page 440](#)
- [Default Settings for GLBP, on page 440](#)
- [Configuring GLBP, on page 441](#)
- [Verifying the GLBP Configuration, on page 448](#)
- [Configuration Examples for GLBP, on page 449](#)
- [Related Documents for GLBP, on page 449](#)
- [Standards for GLBP, on page 449](#)
- [Feature History for GLBP, on page 450](#)

## Finding Feature Information

Your software release might not support all the features documented in this module. For the latest caveats and feature information, see the Bug Search Tool at <https://tools.cisco.com/bugsearch/> and the release notes for your software release. To find information about the features documented in this module, and to see a list of the releases in which each feature is supported, see the "New and Changed Information" chapter or the Feature History table in this chapter.

## Information About GLBP

Gateway Load Balancing Protocol (GLBP) provides path redundancy for IP by sharing protocol and Media Access Control (MAC) addresses between redundant gateways. Additionally, GLBP allows a group of Layer 3 routers to share the load of the default gateway on a LAN. A GLBP router can automatically assume the forwarding function of another router in the group if the other router fails.

GLBP provides automatic gateway backup for IP hosts configured with a single default gateway on an IEEE 802.3 LAN. Multiple routers on the LAN combine to offer a single virtual first-hop IP gateway while sharing the IP packet forwarding load. Other routers on the LAN might act as redundant GLBP gateways that become active if any of the existing forwarding gateways fail.

GLBP performs a similar function to the Hot Standby Redundancy Protocol (HSRP) and the Virtual Router Redundancy Protocol (VRRP). HSRP and VRRP allow multiple routers to participate in a virtual group configured with a virtual IP address. These protocols elect one member as the active router to forward packets to the virtual IP address for the group. The other routers in the group are redundant until the active router fails.

GLBP performs an additional load balancing function that the other protocols do not provide. GLBP load balances over multiple routers (gateways) using a single virtual IP address and multiple virtual MAC addresses. GLBP shares the forwarding load among all routers in a GLBP group instead of allowing a single router to handle the whole load while the other routers remain idle. You configure each host with the same virtual IP address, and all routers in the virtual group participate in forwarding packets. GLBP members communicate between each other using periodic hello messages.

## GLBP Active Virtual Gateway

GLBP prioritizes gateways to elect an active virtual gateway (AVG). If multiple gateways have the same priority, the gateway with the highest real IP address becomes the AVG. The AVG assigns a virtual MAC address to each member of the GLBP group. Each member is the active virtual forwarder (AVF) for its assigned virtual MAC address, forwarding packets sent to its assigned virtual MAC address.

The AVG also answers Address Resolution Protocol (ARP) requests for the virtual IP address. Load sharing is achieved when the AVG replies to the ARP requests with different virtual MAC addresses.



---

**Note** Packets received on a routed port destined for the GLBP virtual IP address terminate on the local router, regardless of whether that router is the active GLBP router or a redundant GLBP router. This termination includes ping and Telnet traffic. Packets received on a Layer 2 (VLAN) interface destined for the GLBP virtual IP address terminate on the active router.

---

## GLBP Virtual MAC Address Assignment

The AVG assigns the virtual MAC addresses to each member of the group. The group members request a virtual MAC address after they discover the AVG through hello messages. The AVG assigns the next MAC address based on the load-balancing algorithm selected. A gateway that is assigned with a virtual MAC address by the AVG is the primary virtual forwarder. The other members of the GLBP group that learn the virtual MAC addresses from hello messages are secondary virtual forwarders.

## GLBP Virtual Gateway Redundancy

GLBP provides virtual gateway redundancy. A member in a group can be in the active, standby, or listen state. GLBP uses a priority algorithm to elect one gateway as the AVG and elect another gateway as the standby virtual gateway. The remaining gateways go into the listen state. You can configure the GLBP priority on each gateway. If the GLBP priority is identical on multiple gateways, GLBP uses the gateway with the highest IP address as the AVG.

If an AVG fails, the standby virtual gateway assumes responsibility for the virtual IP address. GLBP elects a new standby virtual gateway from the gateways in the listen state.

## GLBP Virtual Forwarder Redundancy


GLBP provides virtual forwarder redundancy. Virtual forwarder redundancy is similar to virtual gateway redundancy with an active virtual forwarder (AVF). If the AVF fails, a secondary virtual forwarder in the listen state assumes responsibility for the virtual MAC address. This secondary virtual forwarder is also a primary virtual forwarder for a different virtual MAC address. GLBP migrates hosts away from the old virtual MAC address of the failed AVF, using the following two timers:

- Redirect timer—Specifies the interval during which the AVG continues to redirect hosts to the old virtual MAC address. When the redirect time expires, the AVG stops using the old virtual MAC address in ARP replies, although the secondary virtual forwarder continues to forward packets that were sent to the old virtual MAC address.
- Secondary hold timer—Specifies the interval during which the virtual MAC address is valid. When the secondary hold time expires, GLBP removes the virtual MAC address from all gateways in the GLBP group and load balances the traffic over the remaining AVFs. The expired virtual MAC address becomes eligible for reassignment by the AVG.

GLBP uses hello messages to communicate the current state of the timers.

In the figure, router A is the AVG for a GLBP group and is responsible for the virtual IP address 192.0.2.1. Router A is also an AVF for the virtual MAC address 0007.b400.0101. Router B is a member of the same GLBP group and is designated as the AVF for the virtual MAC address 0007.b400.0102. Client 1 has a default gateway IP address of 192.0.2.1, the virtual IP address, and a gateway MAC address of 0007.b400.0101 that points to router A. Client 2 shares the same default gateway IP address but receives the gateway MAC address 0007.b400.0102 because router B is sharing the traffic load with router A.

**Figure 45: GLBP Topology**



If router A becomes unavailable, client 1 does not lose access to the WAN because router B assumes responsibility for forwarding packets sent to the virtual MAC address of router A and for responding to packets

sent to its own virtual MAC address. Router B also assumes the role of the AVG for the entire GLBP group. Communication for the GLBP members continues despite the failure of a router in the GLBP group.

## GLBP Authentication

GLBP has three authentication types:

- MD5 authentication
- Plain text authentication
- No authentication

MD5 authentication provides greater security than plain text authentication. MD5 authentication allows each GLBP group member to use a secret key to generate a keyed MD5 hash that is part of the outgoing packet. At the receiving end, a keyed hash of an incoming packet is generated. If the hash within the incoming packet does not match the generated hash, the packet is ignored. The key for the MD5 hash can either be given directly in the configuration using a key string or supplied indirectly through a key chain.

You can also choose to use a simple password in plain text to authenticate GLBP packets, or choose no authentication for GLBP.

GLBP rejects packets in any of the following cases:

- The authentication schemes differ on the router and in the incoming packet.
- MD5 digests differ on the router and in the incoming packet.
- Text authentication strings differ on the router and in the incoming packet.

## GLBP Load Balancing and Tracking

You can configure the following load-balancing methods for GLBP:


- Round-robin—GLBP cycles through the virtual MAC addresses sent in ARP replies, load balancing the traffic across all the AVFs.
- Weighted—AVG uses the advertised weight for an AVF to decide the load directed to the AVF. A higher weight means that the AVG directs more traffic to the AVF.
- Host dependent—GLBP uses the MAC address of the host to determine which virtual MAC address to direct the host to use. This algorithm guarantees that a host gets the same virtual MAC address if the number of virtual forwarders does not change.

The default for IPv4 networks is round-robin. You can disable all load balancing for GLBP on an interface. If you do not configure load balancing, the AVG handles all traffic for the hosts while the other GLBP group members are in standby or listen mode.

You can configure GLBP to track an interface or routes and enable the secondary virtual forwarder to take over if a tracked link goes down. GLBP tracking uses weighted load-balancing to determine whether a GLBP group member acts as an AVF. You must configure the initial weighting values and optional thresholds to enable or disable this group member as an AVF. You can also configure the interface to track and the value that reduces the interface's weighting if the interface goes down. When the GLBP group weighting drops below the lower threshold, the member is no longer an AVF and a secondary virtual forwarder takes over. When the weighting rises above the upper threshold, the member can resume its role as an AVF.



Figure 46: GLBP Object Tracking and Weighting



In the figure, the Ethernet 1/2 interface on router 1 is the gateway for host 1 (the AVF for virtual MAC address, vMAC1), while Ethernet 2/2 on router 2 acts as a secondary virtual forwarder for Host 1. Ethernet 1/2 tracks Ethernet 3/1, which is the network connection for router 1. If Ethernet 3/1 goes down, the weighting for Ethernet 1/2 drops to 90. Ethernet 2/2 on router 2 preempts Ethernet 1/2 and takes over as AVF because it has the default weighting of 100 and is configured to preempt the AVF.

## High Availability and Extended Nonstop Forwarding

GLBP supports high availability through stateful restarts and stateful switchovers. A stateful restart occurs when the GLBP process fails and is restarted. A stateful switchover occurs when the active supervisor switches to the standby supervisor. Cisco NX-OS applies the run-time configuration after the switchover.

If GLBP hold timers are configured for short time periods, these timers might expire during a controlled switchover or in-service software upgrade (ISSU). GLBP supports extended non-stop forwarding (NSF) to temporarily extend these GLBP hold timers during a controlled switchover or ISSU.

With extended NSF configured, GLBP sends hello messages with the extended timers. GLBP peers update their hold timers with these new values. The extended timers prevent unnecessary GLBP state changes during the switchover or ISSU. After the switchover or ISSU event, GLBP restores the hold timers to their original configured values. If the switchover fails, GLBP restores the hold timers after the extended hold timer values expire.

## Virtualization Support

GLBP supports virtual routing and forwarding (VRF) instances. VRFs exist within virtual device contexts (VDCs). By default, Cisco NX-OS places you in the default VDC and default VRF unless you specifically configure another VDC and VRF.

If you change the VRF membership of an interface, Cisco NX-OS removes all Layer 3 configuration, including GLBP.

For more information, see the *Cisco Nexus 7000 Series NX-OS Virtual Device Context Configuration Guide*.

## Prerequisites for GLBP

GLBP has the following prerequisites:

- Globally enable the GLBP feature.
- You can only configure GLBP on Layer 3 interfaces (see the *Cisco Nexus 7000 Series NX-OS Interfaces Configuration Guide*, and the *Interfaces Configuration Guide, Cisco DCNM for LAN*).
- If you configure VDCs, install the Advanced Services license and enter the desired VDC (see the *Cisco Nexus 7000 Series NX-OS Virtual Device Context Configuration Guide*).

## Guidelines and Limitations for GLBP

GLBP has the following configuration guidelines and limitations:

- You should configure all customization options for GLBP on all GLBP member gateways before enabling a GLBP group by configuring a virtual IP address.
- You must configure an IP address for the interface that you configure GLBP on and enable that interface before GLBP becomes active.
- The GLBP virtual IP address must be in the same subnet as the interface IP address.
- We recommend that you do not configure more than one first-hop redundancy protocol on the same interface.
- Cisco NX-OS removes all Layer 3 configuration on an interface when you change the VDC, interface VRF membership, port channel membership, or when you change the port mode to Layer 2.
- Cisco NX-OS does not support GLBP group configuration on interface secondary subnets.
- Cisco NX-OS does not support GLBP for IPv6.
- The GLBP does not support gratuitous ARP by design.
- If you are familiar with the Cisco IOS CLI, be aware that the Cisco NX-OS commands for this feature might differ from the Cisco IOS commands that you would use.

## Default Settings for GLBP

*Table 41: Default GLBP Parameters*

| Parameters     | Default           |
|----------------|-------------------|
| Authentication | No authentication |

| Parameters                 | Default       |
|----------------------------|---------------|
| Extended hold timer        | 10 seconds    |
| Forwarder preemption delay | 30 seconds    |
| Forwarder timeout          | 14400 seconds |
| Hello timer                | 3 seconds     |
| Hold timer                 | 10 seconds    |
| GLBP feature               | Disabled      |
| Load balancing             | Round robin   |
| Preemption                 | Disabled      |
| Priority                   | 100           |
| Redirect timer             | 600 seconds   |
| Weighting                  | 100           |

## Configuring GLBP

### Enabling GLBP

You must enable GLBP before you can configure and enable any GLBP groups.

#### Before you begin

Ensure that you are in the correct VDC (or use the **switchto vdc** command).

#### Procedure

|               | Command or Action                        | Purpose                           |
|---------------|------------------------------------------|-----------------------------------|
| <b>Step 1</b> | switch# <b>configure terminal</b>        | Enters global configuration mode. |
| <b>Step 2</b> | switch(config)# <b>[no] feature glbp</b> | Enables GLBP.                     |

#### Example

```
switch# configure terminal
switch(config)# feature glbp
switch(config)# copy running-config startup-config
```

## Configuring GLBP Authentication

You can configure GLBP to authenticate the protocol using cleartext or an MD5 digest. MD5 authentication uses a key chain (see the *Cisco Nexus 7000 Series NX-OS Security Configuration Guide*).

### Before you begin

Ensure that you are in the correct VDC (or use the **switchto vdc** command).

Enable GLBP.



**Note** You must configure the same authentication and keys on all members of the GLBP group.

### Procedure

|               | Command or Action                                                                                                                                              | Purpose                                                                                                                                                                                                                                                                                                                                                                                    |
|---------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configure terminal</b>                                                                                                                              | Enters global configuration mode.                                                                                                                                                                                                                                                                                                                                                          |
| <b>Step 2</b> | switch(config)# <b>interface</b> <i>interface-type</i><br><i>slot/port</i>                                                                                     | Enters interface configuration mode.                                                                                                                                                                                                                                                                                                                                                       |
| <b>Step 3</b> | switch(config-if)# <b>ip</b> <i>ip-address/length</i>                                                                                                          | Configures the IPv4 address for the interface.                                                                                                                                                                                                                                                                                                                                             |
| <b>Step 4</b> | switch(config-if)# <b>glbp</b> <i>group-number</i>                                                                                                             | Creates a GLBP group and enters GLBP configuration mode. The range is from 0 to 1024.                                                                                                                                                                                                                                                                                                      |
| <b>Step 5</b> | switch(config-if-glb)# <b>authentication text</b><br><i>string</i>                                                                                             | Configures cleartext authentication for GLBP on this interface.                                                                                                                                                                                                                                                                                                                            |
| <b>Step 6</b> | switch(config-if-glb)# <b>authentication md5</b><br>{ <b>key-chain</b> <i>key-chain</i>   <b>key-string</b> { <i>text</i>  <br><b>encrypted</b> <i>text</i> }} | Configures MD5 authentication for GLBP on this interface.                                                                                                                                                                                                                                                                                                                                  |
| <b>Step 7</b> | switch(config-if-glb)# <b>ip</b> [ <i>ip-address</i><br>[ <b>secondary</b> ]]                                                                                  | Enables GLBP on an interface and identifies the primary IP address of the virtual gateway.<br><br>After you identify a primary IP address, you can use the <b>glbp group ip</b> command again with the <b>secondary</b> keyword to indicate additional IP addresses supported by this group. If you only use the <b>ip</b> keyword, GLBP learns the virtual IP address from the neighbors. |
| <b>Step 8</b> | (Optional) switch(config-if-glb)# <b>show glbp</b><br>[ <b>group</b> <i>group-number</i> ]                                                                     | Displays GLBP information.                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Step 9</b> | (Optional) switch(config)# <b>copy</b><br><b>running-config startup-config</b>                                                                                 | Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.                                                                                                                                                                                                                                                              |

### Example

This example shows how to configure MD5 authentication for GLBP on Ethernet 1/2 after creating the key chain:

```
switch# configure terminal
switch(config)# key chain glbp-keys
switch(config-keychain)# key 0
switch(config-keychain-key)# key-string 7 zqdest
switch(config-keychain-key) accept-lifetime 00:00:00 Jun 01 2008 23:59:59 Sep 12 2008
switch(config-keychain-key) send-lifetime 00:00:00 Jun 01 2008 23:59:59 Aug 12 2008
switch(config-keychain-key) key 1
switch(config-keychain-key) key-string 7 uaeqdyito
switch(config-keychain-key) accept-lifetime 00:00:00 Aug 12 2008 23:59:59 Dec 12 2008
switch(config-keychain-key) send-lifetime 00:00:00 Sep 12 2008 23:59:59 Nov 12 2008
switch(config)# interface ethernet 1/2
switch(config-if)# glbp 1
switch(config-if-glbp)# authenticate md5 key-chain glbp-keys
switch(config-if-glbp)# copy running-config startup-config
```

## Configuring GLBP Load Balancing

You can configure GLBP to use load balancing based on round-robin, weighted, or host-dependent methods.

### Procedure

|               | Command or Action                                                                                              | Purpose                                                                               |
|---------------|----------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch(config-if)# <b>glbp</b> <i>group-number</i>                                                             | Creates a GLBP group and enters GLBP configuration mode. The range is from 0 to 1024. |
| <b>Step 2</b> | switch(config-if-glbp)# <b>load-balancing</b> [ <b>host-dependent</b>   <b>round-robin</b>   <b>weighted</b> ] | Sets the GLBP load-balancing method. The default is round-robin.                      |

### Example

This example shows how to configure load balancing for GLBP:

```
switch(config-if)# glbp 1
switch(config-if-glbp)# load-balancing weighted
```

## Configuring GLBP Weighting and Tracking

You can configure GLBP weighting values and object tracking to work with the GLBP weighted load-balancing method.

You can optionally configure the interface to preempt an AVF if the interface was originally assigned with the virtual MAC address or if this interface has a higher weight than the AVF.

**Before you begin**

Ensure that you are in the correct VDC (or use the **switchto vdc** command).

Enable GLBP.

**Procedure**

|               | <b>Command or Action</b>                                                                                                                 | <b>Purpose</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|---------------|------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configure terminal</b>                                                                                                        | Enters global configuration mode.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Step 2</b> | switch(config)# <b>track</b> <i>object-id</i> <b>interface</b> <i>interface-type number</i> { <b>ip routing</b>   <b>line-protocol</b> } | Configures the interface that this GLBP interface tracks. Changes in the state of the interface affect the priority of this GLBP interface as follows: <ul style="list-style-type: none"> <li>You configure the interface and corresponding object number that you use with the <b>track</b> command in GLBP configuration mode.</li> <li>The <b>line-protocol</b> keyword tracks whether the interface is up. The <b>ip</b> keyword also checks that IP routing is enabled on the interface and an IP address is configured.</li> </ul> |
| <b>Step 3</b> | switch(config)# <b>track</b> <i>object-id</i> <b>ip route</b> <i>ip-prefix/length</i> <b>reachability</b>                                | Creates a tracked object for a route and enters tracking configuration mode. The <i>object-id</i> range is from 1 to 500.                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Step 4</b> | switch(config)# <b>interface</b> <i>interface-type slot/por</i>                                                                          | Enters interface configuration mode.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Step 5</b> | switch(config-if)# <b>ip</b> <i>ip-address/length</i>                                                                                    | Configures the IPv4 address for the interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Step 6</b> | switch(config-if)# <b>glbp</b> <i>group-number</i>                                                                                       | Creates a GLBP group and enters GLBP configuration mode.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Step 7</b> | switch(config-if-glb)# <b>weighting</b> <i>maximum</i> [ <b>lower</b> <i>lower</i> ] [ <b>upper</b> <i>upper</i> ]                       | Specifies the initial weighting value and the upper and lower thresholds for a GLBP gateway. The maximum range is from 1 to 254. The default weighting value is 100. The lower range is from 1 to 253. The upper range is from 1 to 254.                                                                                                                                                                                                                                                                                                 |
| <b>Step 8</b> | switch(config-if-glb)# <b>weighting-track</b> <i>object-number</i> [ <b>decrement</b> <i>value</i> ]                                     | Specifies an object to be tracked that affects the weighting of a GLBP gateway. The value argument specifies a reduction in the weighting of a GLBP gateway when a tracked object fails. The range is from 1 to 255.                                                                                                                                                                                                                                                                                                                     |
| <b>Step 9</b> | (Optional) switch(config-if-glb)# <b>forwarder preempt</b> [ <b>delay</b> <i>minimum seconds</i> ]                                       | Configures the router to take over as AVF for a GLBP group if the current AVF for a GLBP                                                                                                                                                                                                                                                                                                                                                                                                                                                 |

|                | Command or Action                                                               | Purpose                                                                                                                                                                                                                                                                                                                                                                                    |
|----------------|---------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|                |                                                                                 | group falls below its low weighting threshold. The range is from 0 to 3600 seconds.<br><br>This command is enabled by default with a delay of 30 seconds.                                                                                                                                                                                                                                  |
| <b>Step 10</b> | switch(config-if-glb)# <b>ip</b> [ <i>ip-address</i> [ <b>secondary</b> ]]      | Enables GLBP on an interface and identifies the primary IP address of the virtual gateway.<br><br>After you identify a primary IP address, you can use the <b>glbp group ip</b> command again with the <b>secondary</b> keyword to indicate additional IP addresses supported by this group. If you only use the <b>ip</b> keyword, GLBP learns the virtual IP address from the neighbors. |
| <b>Step 11</b> | (Optional) switch(config-if-glb)# <b>show glbp</b> <i>interface-type number</i> | Displays GLBP information for an interface.                                                                                                                                                                                                                                                                                                                                                |
| <b>Step 12</b> | (Optional) switch(config)# <b>copy running-config startup-config</b>            | Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.                                                                                                                                                                                                                                                              |

### Example

The following example shows how to configure GLBP weighting and tracking on Ethernet 1/2:

```
switch# configure terminal
switch(config)# track 2 interface ethernet 2/2 ip routing
switch(config)# interface ethernet 1/2
switch(config-if)# glbp 1
switch(config-if-glb) weighting 110 lower 95 upper 105
switch(config-if-glb) weighting track 2 decrement 20
switch(config-if-glb)# copy running-config startup-config
```

## Customizing GLBP

Customizing the behavior of GLBP is optional. Be aware that as soon as you enable a GLBP group by configuring a virtual IP address, that group is operational. If you enable a GLBP group before you customize GLBP, the router could take over control of the group and become the AVG before you finish customizing the feature. If you plan to customize GLBP, you should do so before enabling GLBP.

### Procedure

|               | Command or Action                                                                                     | Purpose                                                             |
|---------------|-------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------|
| <b>Step 1</b> | switch(config-if-glb)# <b>glbp</b> <i>group-number</i>                                                | Creates a GLBP group and enters GLBP configuration mode.            |
| <b>Step 2</b> | switch(config-if-glb)# <b>timers</b> [ <b>msec</b> ] <i>hellotime</i> [ <b>msec</b> ] <i>holdtime</i> | Configures the following hello and hold times for this GLBP member: |

|               | Command or Action                                                              | Purpose                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|---------------|--------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|               |                                                                                | <ul style="list-style-type: none"> <li>• <i>hellotime</i>—The interval between successive hello packets sent by the AVG in a GLBP group. The range is from 1 to 60 seconds or from 250 to 60000 milliseconds. The default value is 3 seconds.</li> <li>• <i>holdtime</i>—The interval before the virtual gateway and virtual forwarder information in the hello packet is considered invalid. The range is from 2 to 180 seconds or from 1020 to 180000 milliseconds. The default is 10 seconds.</li> </ul> <p>The optional <b>msec</b> keyword specifies that the argument is expressed in milliseconds, instead of the default seconds.</p> |
| <b>Step 3</b> | switch(config-if-glbp)# <b>timers redirect</b> <i>redirect timeout</i>         | <p>Configures the following timers:</p> <ul style="list-style-type: none"> <li>• <i>redirect</i>—The time interval in seconds during which the AVG continues to redirect clients to an AVF. The range is from 0 to 3600 seconds. The default is 600 seconds.</li> <li>• <i>timeout</i>—The interval in seconds before a secondary virtual forwarder becomes invalid. The range is from 610 to 64800 seconds. The default is 14,440 seconds.</li> </ul>                                                                                                                                                                                        |
| <b>Step 4</b> | switch(config-if-glbp)# <b>priority</b> <i>level</i>                           | <p>Sets the priority level used to select the AVG in a GLBP group. The range is from 1 to 255. The default is 100.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Step 5</b> | switch(config-if-glbp)# <b>preempt</b> [ <b>delay minimum</b> <i>seconds</i> ] | <p>Configures the router to take over as AVG for a GLBP group if it has a higher priority than the current AVG. This command is disabled by default.</p> <p>Use the optional <b>delay minimum</b> keywords and the <i>seconds</i> argument to specify a minimum delay interval in seconds before preemption of the AVG takes place.</p> <p>The seconds range is from 0 to 3600 seconds. The minimum delay default is 3600 seconds.</p>                                                                                                                                                                                                        |

### Example

The following example shows how to customize GLBP:



```
switch(config-if)# glbp 1
switch(config-if-glb) # timers 5 18
switch(config-if-glb) # timers redirect 600 7200
switch(config-if-glb) # priority 254
switch(config-if-glb) # preempt delay minimum 60
```

## Configuring Extended Hold Timers for GLBP

You can configure GLBP to use extended hold timers to support extended NSF during a controlled (graceful) switchover or ISSU, including software upgrades and supervisor switchovers. You should configure extended hold timers on all GLBP gateways.



**Note** You must configure extended hold timers on all GLBP gateways if you configure non-default extended hold timers. You can configure different extended hold timer values on each GLBP gateway, based on the expected system switchover delays.

Use the **show glbp** command to display the extended hold time.

### Procedure

|               | Command or Action                                                     | Purpose                                                                                              |
|---------------|-----------------------------------------------------------------------|------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch(config-if)# <b>glbp</b> <i>group-number</i><br><b>Example:</b> | Creates a GLBP group and enters GLBP configuration mode.                                             |
| <b>Step 2</b> | switch(config)# <b>glbp timers extended-hold</b> [ <i>timer</i> ]     | Sets the GLBP extended hold timer, in seconds. The timer range is from 10 to 255. The default is 10. |

### Example

The following example shows how to configure extended hold timers for GLBP:

```
switch(config-if)# glbp 1
switch(config)# glbp timers extended-hold 30
```

## Enabling a GLBP Group

You can configure the virtual IP address on an interface to enable the GLBP group. You must configure each gateway in the GLBP group with the same group number. The GLBP member can learn all other required parameters from another GLBP member.

### Before you begin

Ensure that you are in the correct VDC (or use the **switchto vdc** command).

Enable GLBP.

**Procedure**

|               | <b>Command or Action</b>                                                                       | <b>Purpose</b>                                                                                                                                                                                                                                                                                                                                                                                                                               |
|---------------|------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configure terminal</b>                                                              | Enters global configuration mode.                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Step 2</b> | switch(config)# <b>interface</b> <i>interface-type slot/por</i>                                | Enters interface configuration mode.                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Step 3</b> | switch(config-if)# <b>ip</b> <i>ip-address/length</i>                                          | Configures the IPv4 address for the interface.                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Step 4</b> | switch(config-if)# <b>glbp</b> <i>group-number</i>                                             | Creates a GLBP group and enters GLBP configuration mode.                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Step 5</b> | switch(config-if-glbp)# <b>ip</b> [ <i>ip-address [secondary]</i> ]                            | Enables GLBP on an interface and identifies the virtual IP address. The virtual IP should be in the same subnet as the interface IP address.<br><br>After you identify a virtual IP address, you can use the <b>glbp group ip</b> command again with the <b>secondary</b> keyword to indicate additional IP addresses supported by this group. If you only use the <b>ip</b> keyword, GLBP learns the virtual IP address from the neighbors. |
| <b>Step 6</b> | switch(config-if-glbp)# <b>show glbp</b> [ <b>group</b> <i>group-number</i> ] [ <b>brief</b> ] | Displays a brief summary of GLBP information.                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Step 7</b> | (Optional) switch(config)# <b>copy running-config startup-config</b>                           | Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.                                                                                                                                                                                                                                                                                                                |

**Example**

This example shows how to enable GLBP on Ethernet 1/2:

```
switch# configure terminal
switch(config)# interface ethernet 1/2
switch(config-if)# glbp 1
switch(config-if-glbp)# ip 192.0.2.10
```

## Verifying the GLBP Configuration

To display GLBP configuration information, perform one of the following tasks:

| <b>Command</b>                                             | <b>Purpose</b>                                     |
|------------------------------------------------------------|----------------------------------------------------|
| <b>show glbp</b> [ <b>group</b> <i>group-number</i> ]      | Displays the GLBP status for all or one group.     |
| <b>show glbp capability</b>                                | Displays the GLBP capability for all or one group. |
| <b>show glbp interface</b> <i>interface-type slot/port</i> | Displays the GLBP status for an interface.         |

| Command                                                                                                       | Purpose                                                                                                            |
|---------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------|
| <code>show glbp interface interface-type slot/port [active] [disabled] [init] [listen] [standby]</code>       | Displays the GLBP status for a group or interface for virtual forwarders in the selected state.                    |
| <code>show glbp interface interface-type slot/port [active] [disabled] [init] [listen] [standby] brief</code> | Displays a brief summary of the GLBP status for a group or interface for virtual forwarders in the selected state. |

## Configuration Examples for GLBP

The following example shows how to enable GLBP on an interface, with MD5 authentication, interface tracking, and weighted load balancing:

```
key chain glbp-keys
key 0
 key-string 7 zqdest
 accept-lifetime 00:00:00 Jun 01 2008 23:59:59 Sep 12 2008
 send-lifetime 00:00:00 Jun 01 2008 23:59:59 Aug 12 2008
key 1
 key-string 7 uaeqdyito
 accept-lifetime 00:00:00 Aug 12 2008 23:59:59 Dec 12 2008
 send-lifetime 00:00:00 Sep 12 2008 23:59:59 Nov 12 2008
feature glbp
track 2 interface ethernet 2/2 ip
interface ethernet 1/2
ip address 192.0.2.2/8
glbp 1
 authentication md5 key-chain glbp-keys
 weighting 110 lower 95 upper 105
 weighting track 2 decrement 20
 ip 192.0.2.10
no shutdown
```

## Related Documents for GLBP

| Related Topic      | Document Title                                                              |
|--------------------|-----------------------------------------------------------------------------|
| IS-IS CLI commands | <i>Cisco Nexus 7000 Series NX-OS Unicast Routing Command Reference</i>      |
| VDCs and VRFs      | <i>Cisco Nexus 7000 Series NX-OS High Availability and Redundancy Guide</i> |

## Standards for GLBP

No new or modified standards are supported by this feature, and support for existing standards has not been modified by this feature.

## Feature History for GLBP

This table includes only the updates for those releases that have resulted in additions or changes to the feature.

*Table 42: Feature History for GLBP*

| Feature Name | Release | Feature Information          |
|--------------|---------|------------------------------|
| GLBP         | 4.0(1)  | This feature was introduced. |



# CHAPTER 21

## Configuring HSRP

---

This chapter contains the following sections:

- [Finding Feature Information, on page 451](#)
- [Information About HSRP, on page 451](#)
- [Prerequisites for HSRP, on page 459](#)
- [Guidelines and Limitations for HSRP, on page 459](#)
- [Default Settings for HSRP Parameters, on page 461](#)
- [Configuring HSRP, on page 461](#)
- [Verifying the HSRP Configuration, on page 479](#)
- [Configuration Examples for HSRP, on page 480](#)
- [Related Documents for HSRP, on page 480](#)
- [MIBs, on page 481](#)
- [Feature History for HSRP, on page 481](#)

### Finding Feature Information

Your software release might not support all the features documented in this module. For the latest caveats and feature information, see the Bug Search Tool at <https://tools.cisco.com/bugsearch/> and the release notes for your software release. To find information about the features documented in this module, and to see a list of the releases in which each feature is supported, see the "New and Changed Information" chapter or the Feature History table in this chapter.

### Information About HSRP

HSRP is a first-hop redundancy protocol (FHRP) that allows a transparent failover of the first-hop IP router. HSRP provides first-hop routing redundancy for IP hosts on Ethernet networks configured with a default router IP address. You use HSRP in a group of routers for selecting an active router and a standby router. In a group of routers, the active router is the router that routes packets; the standby router is the router that takes over when the active router fails or when preset conditions are met.

Many host implementations do not support any dynamic router discovery mechanisms but can be configured with a default router. Running a dynamic router discovery mechanism on every host is not practical for many reasons, including administrative overhead, processing overhead, and security issues. HSRP provides failover services to these hosts.

When you use HSRP, you configure the HSRP virtual IP address as the host default router (instead of the IP address of the actual router). The virtual IP address is an IPv4 or IPv6 address that is shared among a group of routers that run HSRP.

When you configure HSRP on a network segment, you provide a virtual MAC address and a virtual IP address for the HSRP group. You configure the same virtual address on each HSRP-enabled interface in the group. You also configure a unique IP address and MAC address on each interface that acts as the real address. HSRP selects one of these interfaces to be the active router. The active router receives and routes packets destined for the virtual MAC address of the group.

HSRP detects when the designated active router fails. At that point, a selected standby router assumes control of the virtual MAC and IP addresses of the HSRP group. HSRP also selects a new standby router at that time.

HSRP uses a priority designator to determine which HSRP-configured interface becomes the default active router. To configure an interface as the active router, you assign it with a priority that is higher than the priority of all the other HSRP-configured interfaces in the group. The default priority is 100, so if you configure just one interface with a higher priority, that interface becomes the default active router.


Interfaces that run HSRP send and receive multicast User Datagram Protocol (UDP)-based hello messages to detect a failure and to designate active and standby routers. When the active router fails to send a hello message within a configurable period of time, the standby router with the highest priority becomes the active router. The transition of packet forwarding functions between the active and standby router is completely transparent to all hosts on the network.

You can configure multiple HSRP groups on an interface.

A network configured for HSRP. By sharing a virtual MAC address and a virtual IP address, two or more interfaces can act as a single virtual router.

**Figure 47: HSRP Topology with Two Enabled Routers**

The following figure shows a network configured for HSRP. By sharing a virtual MAC address and a virtual IP address, two or more interfaces can act as a single virtual router.



The virtual router does not physically exist but represents the common default router for interfaces that are configured to provide backup to each other. You do not need to configure the hosts on the LAN with the IP address of the active router. Instead, you configure them with the IP address of the virtual router (virtual IP address) as their default router. If the active router fails to send a hello message within the configurable period of time, the standby router takes over, responds to the virtual addresses, and becomes the active router, assuming the active router duties. From the host perspective, the virtual router remains the same.

Packets received on a routed port destined for the HSRP virtual IP address terminate on the local router, regardless of whether that router is the active HSRP router or the standby HSRP router. This includes ping and Telnet traffic. Packets received on a Layer 2 (VLAN) interface destined for the HSRP virtual IP address terminate on the active router.

## HSRP for IPv4

HSRP routers communicate with each other by exchanging HSRP hello packets. These packets are sent to the destination IP multicast address 224.0.0.2 (reserved multicast address used to communicate to all routers) on UDP port 1985. The active router sources hello packets from its configured IP address and the HSRP virtual MAC address while the standby router sources hellos from its configured IP address and the interface MAC address, which might be the burned-in address (BIA). The BIA is the last six bytes of the MAC address that is assigned by the manufacturer of the network interface card (NIC).

Because hosts are configured with their default router as the HSRP virtual IP address, hosts must communicate with the MAC address associated with the HSRP virtual IP address. This MAC address is a virtual MAC address, 0000.0C07.ACxy, where xy is the HSRP group number in hexadecimal based on the respective interface. For example, HSRP group 1 uses the HSRP virtual MAC address of 0000.0C07.AC01. Hosts on the adjoining LAN segment use the normal Address Resolution Protocol (ARP) process to resolve the associated MAC addresses.

HSRP version 2 uses the new IP multicast address 224.0.0.102 to send hello packets instead of the multicast address of 224.0.0.2, which is used by version 1. HSRP version 2 permits an expanded group number range of 0 to 4095 and uses a new MAC address range of 0000.0C9F.F000 to 0000.0C9F.FFFF.



---

**Note** On the HSRP Standby, HSRP adds the HSRP virtual IP address with a cookie "deadbeef".

---

## HSRP for IPv6

IPv6 hosts learn of available IPv6 routers through IPv6 neighbor discovery (ND) router advertisement (RA) messages. These messages are multicast periodically, or might be solicited by hosts, but the time delay for detecting when a default route is down might be 30 seconds or more. HSRP for IPv6 provides a much faster switchover to an alternate default router than the IPv6 ND protocol provides, less than a second if the milliseconds timers are used. HSRP for IPv6 provides a virtual first hop for IPv6 hosts.

When you configure an IPv6 interface for HSRP, the periodic RAs for the interface link-local address stop after IPv6 ND sends a final RA with a router lifetime of zero. No restrictions occur for the interface IPv6 link-local address. Other protocols continue to receive and send packets to this address.

IPv6 ND sends periodic RAs for the HSRP virtual IPv6 link-local address when the HSRP group is active. These RAs stop after a final RA is sent with a router lifetime of 0 when the HSRP group leaves the active state. HSRP uses the virtual MAC address for active HSRP group messages only (hello, coup, and redesign).

HSRP for IPv6 uses the following parameters:

- HSRP version 2
- UDP port 2029
- Virtual MAC address range from 0005.73A0.0000 through 0005.73A0.0FFF
- Multicast link-local IP destination address of FF02::66

- Hop limit set to 255

## HSRP for IPv6 Addresses

An HSRP IPv6 group has a virtual MAC address that is derived from the HSRP group number and a virtual IPv6 link-local address that is derived, by default, from the HSRP virtual MAC address. The default virtual MAC address for an HSRP IPv6 group is always used to form the virtual IPv6 link-local address, regardless of the actual virtual MAC address used by the group.

The following table shows the MAC and IP addresses used for IPv6 neighbor discovery packets and HSRP packets.

**Table 43: HSRP and IPv6 ND Addresses**

| Packet                      | MAC Source Address    | IPv6 source Address    | IPv6 Destination Address | Link-layer Address Option |
|-----------------------------|-----------------------|------------------------|--------------------------|---------------------------|
| Neighbor solicitation (NS)  | Interface MAC address | Interface IPv6 address | —                        | Interface MAC address     |
| Router solicitation (RS)    | Interface MAC address | Interface IPv6 address | —                        | Interface MAC address     |
| Neighbor advertisement (NA) | Interface MAC address | Interface IPv6 address | Virtual IPv6 address     | HSRP virtual MAC address  |
| Route advertisement (RA)    | Interface MAC address | Interface IPv6 address | —                        | HSRP virtual MAC address  |
| HSRP (inactive)             | Interface MAC address | Interface IPv6 address | —                        | —                         |
| HSRP (active)               | Virtual MAC address   | Interface IPv6 address | —                        | —                         |

HSRP does not add IPv6 link-local addresses to the Unicast Routing Information Base (URIB). There are also no secondary virtual IP addresses for link-local addresses.

For global unicast addresses, HSRP adds the virtual IPv6 address to the URIB and IPv6, but does not register the virtual IPv6 addresses to ICMPv6. ICMPv6 redirects are not supported for HSRP IPv6 groups.

## Multiple Group Optimization for HSRP

Beginning with Cisco NX-OS Release 6.2(2), HSRP supports multiple group optimization (MGO). MGO optimizes performance and bandwidth when multiple HSRP groups are configured on many subinterfaces. MGO requires only one HSRP group, known as the master group, on the physical interface for the purpose of electing active and standby routers.

You can create other HSRP groups on subinterfaces of the physical interface or a different interface, such as an SVI interface, and link these to the master HSRP group. These groups are known as slave groups. Slave groups follow their master group state so that they do not participate in any HSRP election mechanisms. Master groups send hello messages at their configured rates. Slave groups send hello messages at a reduced rate, which is called the mac-refresh interval rate. This process is required so that the slave groups can send out periodic messages in order to refresh MAC addresses in switches and learning bridges.



## HSRP Versions

- Cisco NX-OS supports HSRP version 1 by default. You can configure an interface to use HSRP version 2.
- HSRP version 2 has the following enhancements to HSRP version 1:
  - Expands the group number range. HSRP version 1 supports group numbers from 0 to 255. HSRP version 2 supports group numbers from 0 to 4095.
  - For IPv4, uses the IPv4 multicast address 224.0.0.102 or the IPv6 multicast address FF02::66 to send hello packets instead of the multicast address of 224.0.0.2, which is used by HSRP version 1.
  - Uses the MAC address range from 0000.0C9F.F000 to 0000.0C9F.FFFF for IPv4 and 0005.73A0.0000 through 0005.73A0.0FFF for IPv6 addresses. HSRP version 1 uses the MAC address range 0000.0C07.AC00 to 0000.0C07.ACFE.
  - Adds support for MD5 authentication.
- When you change the HSRP version, Cisco NX-OS reinitializes the group because it now has a new virtual MAC address.
- HSRP version 2 has a different packet format than HSRP version 1. The packet format uses a type-length-value (TLV) format. HSRP version 2 packets received by an HSRP version 1 router are ignored.

## HSRP Authentication

HSRP message digest 5 (MD5) algorithm authentication protects against HSRP-spoofing software and uses the industry-standard MD5 algorithm for improved reliability and security. HSRP includes the IPv4 or IPv6 address in the authentication TLVs.

## HSRP Messages

Routers that are configured with HSRP exchange the following three types of multicast messages:


- Hello—The hello message conveys the HSRP priority and state information of the router to other HSRP routers.
- Coup—When a standby router wants to assume the function of the active router, it sends a coup message.
- Resign—A router that is the active router sends this message when it is about to shut down or when a router that has a higher priority sends a hello or coup message.

## HSRP Load Sharing

HSRP allows you to configure multiple groups on an interface. You can configure two overlapping IPv4 HSRP groups to load share traffic from the connected hosts while providing the default router redundancy expected from HSRP. The following figure shows an example of a load-sharing HSRP IPv4 configuration.

**Figure 48: HSRP Load Sharing**

The figure shows two routers A and B and two HSRP groups. Router A is the active router for group A but is the standby router for group B. Similarly, router B is the active router for group B and the standby router for group A. If both routers remain active, HSRP load balances the traffic from the hosts across both routers. If either router fails, the remaining router continues to process traffic for both hosts.



**Note** HSRP for IPv6 load-balances by default. If there are two HSRP IPv6 groups on the subnet, then hosts learn of both groups from their router advertisements and choose to use one so that the load is shared between the advertised routers.

## Object Tracking and HSRP

You can use object tracking to modify the priority of an HSRP interface based on the operational state of another interface. Object tracking allows you to route to a standby router if the interface to the main network fails.

Two objects that you can track are the line protocol state of an interface or the reachability of an IP route. If the specified object goes down, Cisco NX-OS reduces the HSRP priority by the configured amount.

## vPC and HSRP

HSRP interoperates with virtual port channels (vPCs). vPCs allow links that are physically connected to two different Cisco Nexus devices to appear as a single port channel by a third device. See the *Cisco Nexus 7000 Series NX-OS Layer 2 Switching Configuration Guide*, for more information on vPCs.

vPC forwards traffic through both the active HSRP router and the standby HSRP router.



---

**Note** You should configure HSRP on the primary vPC peer device as active and HSRP on the vPC secondary device as standby.

---

## vPC Peer Gateway and HSRP

Some third-party devices can ignore the HSRP virtual MAC address and instead use the source MAC address of an HSRP router. In a vPC environment, the packets using this source MAC address may be sent across the vPC peer link, causing a potential dropped packet. Configure the vPC peer gateway to enable the HSRP routers to directly handle packets sent to the local vPC peer MAC address and the remote vPC peer MAC address, as well as the HSRP virtual MAC address. See the *Cisco Nexus 7000 Series NX-OS Layer 2 Switching Configuration Guide*, for more information on the vPC peer gateway.

For mixed-chassis configurations where the vPC peer link is configured on an F-series module, configure the vPC peer gateway exclude option to exclude the Layer 3 backup route that traverses the vPC peer link. See the *Cisco Nexus 7000 Series NX-OS Layer 2 Switching Configuration Guide*, for more information on the vPC peer gateway exclude option.

## FabricPath Anycast HSRP

Cisco NX-OS Release 6.2(2) and later releases facilitate further scalability at the spine layer by providing support for more than two nodes. You can create an anycast bundle, which is an association between a set of VLANs and an anycast switch ID. The set of VLANs or the HSRP group elects an active router and a standby router. The remaining routers in the group are in listen state.

All of the HSRP routers that have a configured anycast switch ID advertise the ID through FabricPath IS-IS. The active HSRP router is the only router that uses the anycast switch ID in its hello packets. The leaf switches learn that the anycast switch ID is reachable by all of the routers in the group.

All of the first hop gateways at the spine layer need to function in active-active forwarding mode. IP packets are received by any of the spine switches with the destination set as the gateway MAC address, and these packets are terminated and locally forwarded.

## BFD

This feature supports bidirectional forwarding detection (BFD). BFD is a detection protocol that provides fast-forwarding and path-failure detection times. BFD provides subsecond failure detection between two adjacent devices and can be less CPU-intensive than protocol hello messages because some of the BFD load can be distributed onto the data plane on supported modules. See the *Cisco Nexus 7000 Series NX-OS Interfaces Configuration Guide*, for more information.

## High Availability and Extended Nonstop Forwarding

HSRP supports stateful restarts and stateful switchovers. A stateful restart occurs when the HSRP process fails and is restarted. A stateful switchover occurs when the active supervisor switches to the standby supervisor. Cisco NX-OS applies the run-time configuration after the switchover.

If HSRP hold timers are configured for short time periods, these timers might expire during a controlled switchover or in-service software upgrade (ISSU). Ping to a virtual IP is also unreachable during this timer

expiry period. HSRP supports extended non-stop forwarding (NSF) to temporarily extend these HSRP hold timers during a controlled switchover or ISSU.

With extended NSF configured, HSRP sends hello messages with the extended timers. HSRP peers update their hold timers with these new values. The extended timers prevent unnecessary HSRP state changes during the switchover or ISSU. After the switchover or ISSU event, HSRP restores the hold timers to their original configured values. If the switchover fails, HSRP restores the hold timers after the extended hold timer values expire.

## Virtualization Support

HSRP supports virtual routing and forwarding (VRF) instances. VRFs exist within virtual device contexts (VDCs). By default, Cisco NX-OS places you in the default VDC and default VRF unless you specifically configure another VDC and VRF.

If you change the VRF membership of an interface, Cisco NX-OS removes all Layer 3 configuration, including HSRP.

## HSRP VIP

Starting with Cisco NX-OS Release 7.2(0)D1(1), the Hot Standby Router Protocol (HSRP) Virtual IP (VIP) feature provides support for an HSRP Virtual IP configuration to be in a different subnet than that of the interface subnet. This feature is supported only for IPv4 address and not for IPv6. The following are the enhancements:

- Enhance ARP to source with VIP from Supervisor Engine (SUP) for hosts, when the hosts in VIP subnet are referenced by static route to VLAN configuration.
- Support periodic ARP synchronization to vPC peer if the HSRP VIP feature is enabled.
- Allow VIP address as the Layer 3 source address and gateway address for all communications with a Dynamic Host Configuration Protocol (DHCP) server.
- Enhance DHCP relay agent to relay DHCP packets with VIP address as source address instead of SVI IP address.




---

**Note** HSRP subnet VIP should be configured in the virtual port channel (vPC) topology. The HSRP VIP feature works only on HSRP with vPC topologies.

---




---

**Note** In a subnet VIP configuration, the VIP address must be in a different subnet than the interface IP subnet. Without the subnet VIP configuration, the VIP address must be in the same subnet of the interface IP subnet.

---

The following is an example for VIP subnet address configuration wherein the VIP address is not configured in the same subnet of the interface IP subnet.

```
switch# configure terminal
switch(config)# feature hsrp
switch(config)# feature interface-vlan
switch(config)# interface vlan 2
```

```
switch(config-if)# ip address 192.0.2.1/24
switch(config-if)# hsrp 2
switch(config-if-hsrp)# ip 209.165.201.1/24
```

The following is an example for VIP address mismatch. Here the VIP address is not in the same subnet of the interface IP subnet.

```
switch# configure terminal
switch(config)# feature hsrp
switch(config)# feature interface-vlan
switch(config)# interface vlan 2
switch(config-if)# ip address 192.0.2.1/24
switch(config-if)# hsrp 2
switch(config-if-hsrp)# ip 209.165.201.1

!ERROR: Invalid IP address(Mismatch with IP subnet)!
```

The following is an example for VIP address mismatch. Here the VIP subnet address is configured along with VIP address in the same subnet of the interface IP subnet.

```
switch# configure terminal
switch(config)# feature hsrp
switch(config)# feature interface-vlan
switch(config)# interface vlan 2
switch(config-if)# ip address 192.0.2.1/24
switch(config-if)# hsrp 2
switch(config-if-hsrp)# ip 192.0.2.10/24

!ERROR: Invalid IP address(Mismatch with IP subnet)!
```

## Prerequisites for HSRP

- You must enable the HSRP feature in a device before you can configure and enable any HSRP groups.

## Guidelines and Limitations for HSRP

HSRP has the following configuration guidelines and limitations:

- When host connected to HSRP standby sends ping to HSRP Virtual IP, HSRP Active responds to the ping however the ping packets (transient traffic) get punted to the SUP on HSRP standby before reaching HSRP Active.
- You must configure an IP address for the interface on which you configure HSRP and enable that interface before HSRP becomes active.
- You must configure HSRP version 2 when you configure an IPv6 interface for HSRP.
- For IPv4, the virtual IP address must be in the same subnet as the interface IP address.
- The value of the first 2 digits of a type 7 key string configured by using the **key-string 7 text-string** command has to be between 0 and 15. For example, you can configure 07372b557e2c1a as the key string value in which case the sum value of the first 2 digits will be 7. But, you cannot configure 85782916342021 as the key string value because the value of the first 2 digits will be 85. We recommend unconfiguring any type 7 key strings that do not adhere to this value or to configure a type 0 string.

- We recommend that you do not configure more than one first-hop redundancy protocol on the same interface.
- HSRP version 2 does not interoperate with HSRP version 1. An interface cannot operate both version 1 and version 2 because both versions are mutually exclusive. However, the different versions can be run on different physical interfaces of the same router.
- You cannot change from version 2 to version 1 if you have configured groups above the group number range allowed for version 1 (0 to 255).
- HSRP for IPv4 is supported with BFD. HSRP for IPv6 is not supported with BFD.
- Cisco NX-OS removes all Layer 3 configurations on an interface when you change the interface VRF membership, port channel membership, or when you change the port mode to Layer 2.
- If you configure virtual MAC addresses with vPC, you must configure the same virtual MAC address on both vPC peers.
- For mixed-chassis configurations where the vPC peer link is configured on an F-series module, configure the vPC peer gateway exclude option to exclude the Layer 3 backup route that traverses the vPC peer link.
- You cannot use the HSRP MAC address burned-in option on a VLAN interface that is a vPC member.
- If you have not configured authentication, the **show hsrp** command displays the following string:  

```
Authentication text "cisco"
```
- The following is the default behavior of HSRP as defined in RFC 2281: If no authentication data is configured, the RECOMMENDED default value is 0x63 0x69 0x73 0x63 0x6F 0x00 0x00 0x00.
- Anycast HSRP does not support BFD.
- HSRP for MGO has the following limitations:
  - Master groups and slave groups are not restricted to the same interface.
  - HSRP for MGO supports only HSRP version 2.
  - Master and slave groups must have the same address types.
  - Configuring an HSRP group as a slave group clears the group's other configurations, such as its virtual IP address, without notification, so you must enter the **follow** command before you enter the **ipip-address** command.
  - Bidirectional forwarding (BFD) is not applicable to slave groups.
  - HSRP for MGO supports both IPv4 and IPv6 interfaces and works for all Layer 3 interfaces on which a regular HSRP group works.
  - An HSRP group cannot be configured as both a master and slave group at the same time.

# Default Settings for HSRP Parameters

## Default HSRP Parameters

| Parameters          | Default                                                   |
|---------------------|-----------------------------------------------------------|
| HSRP                | Disabled                                                  |
| Authentication      | Enabled as text for version 1, with cisco as the password |
| HSRP version        | Version 1                                                 |
| Preemption          | Disabled                                                  |
| Priority            | 100                                                       |
| Virtual MAC address | Derived from HSRP group number                            |

## Configuring HSRP

### Enabling HSRP

You must globally enable HSRP before you can configure and enable any HSRP groups.

#### Before you begin

Confirm that you are in the correct VDC. To change the VDC, use the **switchto vdc** command.

#### Procedure

|               | Command or Action                                                    | Purpose                                                                                                                                                                                                                                                    |
|---------------|----------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configure terminal</b>                                    | Enters global configuration mode.                                                                                                                                                                                                                          |
| <b>Step 2</b> | switch(config)# <b>[no] feature hsrp</b>                             | Enables the HSRP feature.<br><br>Use the <b>no</b> form of this command to disable this feature.<br><br>You can use this command to enable or disable the HSRP feature and remove all associated configurations in a VDC in the global configuration mode. |
| <b>Step 3</b> | (Optional) switch(config)# <b>copy running-config startup-config</b> | Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.                                                                                                                              |

**Example**

The following example enables HSRP:

```
switch # configure terminal
switch(config)# feature hsrp
switch(config)# copy running-config startup-config
```

## Configuring the HSRP Version

You can configure the HSRP version. If you change the version for existing groups, Cisco NX-OS reinitializes HSRP for those groups because the virtual MAC address changes. The HSRP version applies to all groups on the interface.



**Note** IPv6 HSRP groups must be configured as HSRP version 2.

**Procedure**

|               | Command or Action                                                       | Purpose                                                                                                                       |
|---------------|-------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configure terminal</b>                                       | Enters global configuration mode.                                                                                             |
| <b>Step 2</b> | switch(config)# <b>interface</b> <i>type number</i>                     | Enters interface configuration mode.                                                                                          |
| <b>Step 3</b> | switch(config-if)# <b>hsrp version</b> {1   2}                          | Confirms the HSRP version. Version 1 is the default.                                                                          |
| <b>Step 4</b> | (Optional) switch(config-if)# <b>copy running-config startup-config</b> | Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration. |

**Example**

The following example shows how to configure an HSRP version:

```
switch# configure terminal
switch(config)# interface ethernet 1/2
switch(config-if)# hsrp version 2
switch(config-if)# copy running-config startup-config
```

## Configuring an HSRP Group for IPv4

You can configure an HSRP group on an IPv4 interface and configure the virtual IP address and virtual MAC address for the HSRP group.

**Before you begin**

- You must enable HSRP.



- Cisco NX-OS enables an HSRP group once you configure the virtual IP address on any member interface in the group. You must configure HSRP attributes such as authentication, timers, and priority before you enable the HSRP group.
- Confirm that you are in the correct VDC. To change the VDC, use the **switchto vdc** command.

### Procedure

|               | Command or Action                                                                                   | Purpose                                                                                                                                                                             |
|---------------|-----------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configure terminal</b>                                                                   | Enters global configuration mode.                                                                                                                                                   |
| <b>Step 2</b> | switch(config)# <b>interface</b> <i>type number</i>                                                 | Enters interface configuration mode.                                                                                                                                                |
| <b>Step 3</b> | switch(config-if)# <b>ip</b> <i>ip-address/length</i>                                               | Configures the IPv4 address of the interface.                                                                                                                                       |
| <b>Step 4</b> | switch(config-if)# <b>hsrp</b> <i>group-number</i> [ <b>ipv4</b> ]                                  | Creates an HSRP group and enters hsrp configuration mode. The range for HSRP version 1 is from 0 to 255. The range is for HSRP version 2 is from 0 to 4095. The default value is 0. |
| <b>Step 5</b> | switch(config-if-hsrp)# <b>ip</b> [ <i>ip-address</i> [ <b>secondary</b> ]]                         | Configures the virtual IP address for the HSRP group and enables the group. This address should be in the same subnet as the IPv4 address of the interface.                         |
| <b>Step 6</b> | switch(config-if-hsrp)# <b>exit</b>                                                                 | Exits HSRP configuration mode.                                                                                                                                                      |
| <b>Step 7</b> | switch(config-if)# <b>no shutdown</b>                                                               | Enables the interface.                                                                                                                                                              |
| <b>Step 8</b> | (Optional) switch(config-if)# <b>copy running-config startup-config</b>                             | Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.                                                       |
| <b>Step 9</b> | (Optional) switch(config-if)# <b>show hsrp</b> [ <b>group</b> <i>group-number</i> ] [ <b>ipv4</b> ] | Displays HSRP information.                                                                                                                                                          |

### Example

The following example shows how to configure an HSRP group on Ethernet 1/2:

```
switch# configure terminal
switch(config)# interface ethernet 1/2

switch(config-if)# ip 192.0.2.2/8
switch(config-if)# hsrp 2
switch(config-if-hsrp)# ip 192.0.2.1
switch(config-if-hsrp)# exit
switch(config-if)# no shutdown
switch(config-if)# copy running-config startup-config
```

## Configuring an HSRP Group for IPv6

You can configure an HSRP group on an IPv6 interface and configure the virtual IP address and virtual MAC address for the HSRP group. When you configure an HSRP group for IPv6, HSRP generates a link-local address from the link-local prefix. HSRP also generates a modified EUI-64 format interface identifier in which the EUI-64 interface identifier is created from the relevant HSRP virtual MAC address. There are no HSRP IPv6 secondary addresses.

### Before you begin

- You must enable HSRP.
- Ensure that you have enabled HSRP version 2 on the interface where you want to configure an IPv6 HSRP group.
- Ensure that you have configured HSRP attributes such as authentication, timers, and priority before you enable the HSRP group.
- Confirm that you are in the correct VDC. To change the VDC, use the **switchto vdc** command.

### Procedure

|                | Command or Action                                                    | Purpose                                                                                                                                                                                   |
|----------------|----------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b>  |                                                                      |                                                                                                                                                                                           |
| <b>Step 2</b>  | switch# <b>configure terminal</b>                                    | Enters global configuration mode.                                                                                                                                                         |
| <b>Step 3</b>  | switch(config)# <b>interface</b> <i>type number</i>                  | Enters interface configuration mode.                                                                                                                                                      |
| <b>Step 4</b>  | switch(config-if)# <b>ipv6 address</b><br><i>ipv6-address/length</i> | Configures the IPv6 address of the interface.                                                                                                                                             |
| <b>Step 5</b>  | switch(config-if)# <b>hsrp version 2</b>                             | Configures the group for HSRP version 2.                                                                                                                                                  |
| <b>Step 6</b>  | switch(config-if)# <b>hsrp group-number [ipv6]</b>                   | Creates an IPv6 HSRP group and enters hsrp configuration mode. The range for HSRP version 2 is from 0 to 4095. The range is for HSRP version 2 is from 0 to 4095. The default value is 0. |
| <b>Step 7</b>  | switch(config-if-hsrp)# <b>ip [ipv6-address [secondary]]</b>         | Configures the virtual IPv6 address for the HSRP group and enables the group.                                                                                                             |
| <b>Step 8</b>  | switch(config-if-hsrp)# <b>ip autoconfig</b>                         | Autoconfigures the virtual IPv6 address for the HSRP group from the calculated link-local virtual IPv6 address and enables the group.                                                     |
| <b>Step 9</b>  | switch(config-if-hsrp)# <b>no shutdown</b>                           | Enables the interface.                                                                                                                                                                    |
| <b>Step 10</b> | switch(config-if-hsrp)# <b>copy running-config startup-config</b>    | Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.                                                             |

|                | Command or Action                                                                                    | Purpose                    |
|----------------|------------------------------------------------------------------------------------------------------|----------------------------|
| <b>Step 11</b> | (Optional) switch(config-if-hsrp)# <b>show hsrp</b><br>[ <i>group group-number</i> ] [ <i>ipv6</i> ] | Displays HSRP information. |

### Example

The following example shows how to configure an HSRP group on Ethernet 1/2:

```
switch# configure terminal
switch(config)# interface ethernet 3/2
switch(config-if)# ipv6 address 2001:0DB8:0001:0001:/64
switch(config-if)# hsrp version 2
switch(config-if)# hsrp 2 ipv6
switch(config)# ip 2001:DB8::1
switch(config-if-hsrp)# exit
switch(config-if-hsrp)# no shutdown
switch(config-if-hsrp)# copy running-config startup-config
```

## Configuring an HSRP Master Group Task

You can configure HSRP for MGO to optimize performance when scaling by configuring master and slave groups. Slave groups follow the master group state, which minimizes the number of hello messages that are sent. Cisco NX-OS enables an HSRP group once you configure its virtual IP address.

We recommend that you configure master groups on the same parent interface as their slave groups to allow the slave groups to have the same redundancy requirements as the master group. If a failure occurs on the master link, all the slave groups are brought down as well, even if the links on which they are configured remain up.

### Before you begin

- Ensure that you have enabled the HSRP feature.
- Configure HSRP attributes such as authentication, timers, and priority before you enable an HSRP group as a master group.
- Ensure that you are in the correct VDC (or use the **switchto vdc** command).

### Procedure

|               | Command or Action                                             | Purpose                                                                                                                                      |
|---------------|---------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configure terminal</b>                             | Enters global configuration mode.                                                                                                            |
| <b>Step 2</b> | switch(config)# <b>interface</b> <i>type/number</i>           | Enters interface configuration mode and configures an interface type.                                                                        |
| <b>Step 3</b> | switch(config-if)# <b>ip address</b> <i>ip-address/length</i> | Configures the IP address of the interface.                                                                                                  |
| <b>Step 4</b> | switch(config-if)# <b>hsrp version 2</b>                      | Configures the HSRP version. Because MGO supports only HSRP version 2, you must set the HSRP version to version 2. Version 1 is the default. |

|                | Command or Action                                                                                                         | Purpose                                                                                                                                                                                                                                                                                           |
|----------------|---------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 5</b>  | switch(config-if)# <b>hsrp</b> <i>group-number</i> [ <b>ipv6</b> ]                                                        | Creates an HSRP group and enters HSRP configuration mode. The range for the HSRP group number is from 0 to 4095. The no form of this command removes the group.                                                                                                                                   |
| <b>Step 6</b>  | switch(config-if-hsrp)# <b>name</b> [ <i>master-group-name</i> ]                                                          | Specifies a master group name. The name command changes a regular HSRP group into a master group. If you do not specify a name, a unique name is automatically generated. The no form of this command returns the master group to a regular HSRP group.                                           |
| <b>Step 7</b>  | switch(config-if-hsrp)# <b>ip</b> [ <i>ip-address</i> [ <b>secondary</b> ]]                                               | Configures the virtual IP address for the HSRP group and enables the master group.                                                                                                                                                                                                                |
| <b>Step 8</b>  | switch(config-if-hsrp)# <b>exit</b>                                                                                       | Exits the HSRP configuration mode.                                                                                                                                                                                                                                                                |
| <b>Step 9</b>  | switch(config-if)# <b>no shutdown</b>                                                                                     | Enables the interface.                                                                                                                                                                                                                                                                            |
| <b>Step 10</b> | switch(config-if)# <b>show hsrp</b> [ <b>brief</b> ] [ <b>group</b> <i>group-number</i> ] [ <b>ipv4</b> ] [ <b>ipv6</b> ] | (Optional)<br>Displays HSRP information.                                                                                                                                                                                                                                                          |
| <b>Step 11</b> | switch(config-if)# <b>show hsrp mgo</b> [ <b>name</b> <i>name</i> ] [ <b>brief</b> ]                                      | (Optional)<br>Displays the relationships between HSRP groups that are in use for MGO and their slave sessions. The name keyword restricts the output to the session with a matching configured name. The brief keyword provides a summary of each MGO session with the associated slave sessions. |

### Example

The following example shows how to configure an HSRP master group on Ethernet interface 1/1:

```
switch# configure terminal
switch(config)# interface ethernet 1/1
switch(config-if)# ip address 11.0.0.1/24
switch(config-if)# hsrp version 2
switch(config-if)# hsrp 11
switch(config-if-hsrp)# name master1
switch(config-if-hsrp)# ip 11.0.0.100
switch(config-if-hsrp)# exit
switch(config-if)# no shutdown
switch(config-if)# show hsrp group 11
switch(config-if)# show hsrp mgo name master1
```

## Configuring an HSRP Slave Group

If a failure occurs in a slave link that belongs to a different interface than the master group, the slave group is brought down, regardless of the state of the group it is following.

You can configure HSRP for MGO to optimize performance when scaling by configuring master and slave groups. Slave groups follow the master group state, which minimizes the number of hello messages that are sent. Cisco NX-OS enables an HSRP group once you configure its virtual IP address.

We recommend that you configure master groups on the same parent interface as their slave groups to allow the slave groups to have the same redundancy requirements as the master group. If a failure occurs on the master link, all the slave groups are brought down as well, even if the links on which they are configured remain up.

### Before you begin

- Ensure that you have enabled the HSRP feature.
- Configure HSRP attributes such as authentication, timers, and priority before you enable an HSRP group as a master group.
- Ensure that you are in the correct VDC (or use the `switchto vdc` command).

### Procedure

|               | Command or Action                                             | Purpose                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|---------------|---------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configure terminal</b>                             | Enters global configuration mode.                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Step 2</b> | switch(config)# <b>interface</b> <i>type/number</i>           | Enters interface configuration mode and configures an interface type.                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Step 3</b> | switch(config-if)# <b>ip address</b> <i>ip-address/length</i> | Configures the IP address of the interface.                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Step 4</b> | switch(config-if)# <b>hsrp version 2</b>                      | Configures the HSRP version. Because MGO supports only HSRP version 2, you must set the HSRP version to version 2. Version 1 is the default.                                                                                                                                                                                                                                                                                                  |
| <b>Step 5</b> | switch(config-if)# <b>hsrp mac refresh</b> <i>seconds</i>     | (Optional)<br>Configures the MAC refresh interval for the HSRP slave group. You can use this command to minimize the number of hello messages that are sent out and reduce HSRP protocol overheads and CPU utilization when multiple subinterfaces are configured.<br><br>This command is not available for individual subinterfaces. It applies to all groups on all subinterfaces. The default is 60 seconds. The range is from 0 to 10000. |
| <b>Step 6</b> | switch(config-if)# <b>hsrp group-number</b> [ <b>ipv6</b> ]   | Creates an HSRP group and enters HSRP configuration mode. The range for the HSRP                                                                                                                                                                                                                                                                                                                                                              |

|                | Command or Action                                                                                                            | Purpose                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|----------------|------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|                |                                                                                                                              | group number is from 0 to 4095. The no form of this command removes the group.                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Step 7</b>  | switch(config-if-hsrp)# <b>follow</b><br>[ <i>master-group-name</i> ]                                                        | Configures a regular HSRP group as a slave group.<br><br>Configuring an HSRP group as a slave group clears the group's other configurations, such as its virtual IP address without notification, so you must enter the follow command before you enter the ip ip-address command.<br><br>Slave groups may forward reference master group names that are undefined.<br><br>The no form of this command returns the slave group to a regular HSRP group. |
| <b>Step 8</b>  | switch(config-if-hsrp)# <b>ip</b> [ <i>ip-address</i> ]                                                                      | Configures the virtual IP address for the HSRP group and enables the slave group.                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Step 9</b>  | switch(config-if-hsrp)# <b>exit</b>                                                                                          | Exits the HSRP configuration mode.                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Step 10</b> | switch(config-if)# <b>no shutdown</b>                                                                                        | Enables the interface.                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Step 11</b> | switch(config-if)# <b>show hsrp</b> [ <b>brief</b> ] [ <b>group</b><br><i>group-number</i> ] [ <b>ipv4</b> ] [ <b>ipv6</b> ] | (Optional)<br>Displays HSRP information.                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Step 12</b> | switch(config-if)# <b>show hsrp mgo</b> [ <b>name</b><br><i>name</i> ] [ <b>brief</b> ]                                      | (Optional)<br>Displays the relationships between HSRP groups that are in use for MGO and their slave sessions. The name keyword restricts the output to the session with a matching configured name. The brief keyword provides a summary of each MGO session with the associated slave sessions.                                                                                                                                                       |

### Example

The following example shows how to configure an HSRP slave group on Ethernet interface 1/2:

```
switch# configure terminal
switch(config)# interface ethernet 1/2
switch(config-if)# ip address 12.0.0.1/24
switch(config-if)# hsrp version 2
switch(config-if)# hsrp 12
switch(config-if-hsrp)# follow master1
switch(config-if-hsrp)# ip 12.0.0.100
switch(config-if-hsrp)# exit
switch(config-if)# no shutdown
switch(config-if)# show hsrp group 11
switch(config-if)# show hsrp mgo name master1
```

## Configuring the HSRP Virtual MAC Address Manually

You can override the default virtual MAC address that HSRP derives from the configured group number. You must configure the same virtual MAC address on both vPC peers of a vPC link.

### Procedure

|               | Command or Action                                                  | Purpose                                                                                                                                                                             |
|---------------|--------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configure terminal</b>                                  | Enters global configuration mode.                                                                                                                                                   |
| <b>Step 2</b> | switch(config)# <b>interface</b> <i>type number</i>                | Enters interface configuration mode.                                                                                                                                                |
| <b>Step 3</b> | switch(config-if)# <b>hsrp</b> <i>group-number</i> [ <b>ipv4</b> ] | Creates an HSRP group and enters hsrp configuration mode. The range for HSRP version 1 is from 0 to 255. The range is for HSRP version 2 is from 0 to 4095. The default value is 0. |
| <b>Step 4</b> | switch(config-if-hsrp)# <b>mac-address</b> <i>string</i>           | Configures the virtual MAC address for an HSRP group. The string uses the standard MAC address format (xxxx.xxxx.xxxx).                                                             |
| <b>Step 5</b> | switch(config-if-hsrp)# <b>copy running-config startup-config</b>  | Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.                                                       |

### Example

The following example configures the HSRP virtual MAC address manually:

```
switch # configure terminal
switch(config)# interface ethernet 1/2
switch(config-if)# hsrp 2
switch(config-if-hsrp)# mac-address 5000.1000.1060
switch(config-if-hsrp)# copy running-config startup-config
```

## Configuring the HSRP Virtual MAC Address Using Burned-in MAC Address

You can override the default virtual MAC address that HSRP derives from the configured group number. You must configure the same virtual MAC address on both vPC peers of a vPC link.

### Procedure

|               | Command or Action                                                 | Purpose                                                                                |
|---------------|-------------------------------------------------------------------|----------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configure terminal</b>                                 | Enters global configuration mode.                                                      |
| <b>Step 2</b> | switch(config)# <b>interface</b> <i>type number</i>               | Enters interface configuration mode.                                                   |
| <b>Step 3</b> | switch(config-if)# <b>hsrp use-bia</b> [ <b>scope interface</b> ] | Configures HSRP to use the burned-in MAC address of the interface for the HSRP virtual |

|               | Command or Action                                            | Purpose                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|---------------|--------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|               |                                                              | <p>MAC address. Optionally, you can configure HSRP to use the burned-in MAC address for all groups on this interface by using the scope interface keyword.</p> <p><b>Note</b> Proxy ARP breaks when HSRP is configured with use-bia command. A standby router cannot cover for the lost proxy ARP database of the failed router.</p> <p>When the use-bia option is configured, the ARP process on the HSRP active device mistakenly sees the HSRP group as the standby device because of the lack of virtual address that it looks for. As a result, both the HSRP active and the standby devices suppress ARP replies to proxy ARP requests.</p> |
| <b>Step 4</b> | switch(config-if)# <b>copy running-config startup-config</b> | (Optional) Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |

### Example

The following example configures the HSRP virtual MAC address manually:

```
switch # configure terminal
switch(config)# interface ethernet 1/2
switch(config-if)# hsrp use-bia
switch(config-if)# copy running-config startup-config
```

## Configuring MAC Address Table Reservation for HSRP

Before Release 8.4(6), the behaviour was to enable HSRP Virtual MAC reservation. Beginning with Cisco NX-OS Release 8.4(6), the default behaviour is changed. If required you can enable default Virtual MAC reservation.

You can add MAC reservation for the HSRP Virtual MAC for all VLANs this overrides the default behaviour. By default, If SVI interface is shut on both vPC+peers, the mac reservation is removed. This configuration overrides and MAC reservation is assigned. Enable Virtual MAC reservation with below configuration commands:







```

hsrp version 2
hsrp 1001
mac-address 0000.0044.0004
ip 10.10.1.1
hsrp 1001 ipv6
ip 10:10:1::1
switch(config)#

```

## Authenticating HSRP

You can configure HSRP to authenticate the protocol using cleartext or MD5 digest authentication. MD5 authentication uses a key chain. For more details, see the *Cisco Nexus 7000 Series NX-OS Security Configuration Guide*.

### Before you begin

- You must enable HSRP.
- You must configure the same authentication and keys on all members of the HSRP group.
- Ensure that you have created the key chain if you are using MD5 authentication.
- Confirm that you are in the correct VDC. To change the VDC, use the **switchto vdc** command.

### Procedure

|               | Command or Action                                                                                                                                                                                                      | Purpose                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|---------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configure terminal</b>                                                                                                                                                                                      | Enters global configuration mode.                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Step 2</b> | switch(config)# <b>interface</b> <i>interface-type slot/port</i>                                                                                                                                                       | Enters interface configuration mode.                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Step 3</b> | switch(config-if)# <b>hsrp</b> <i>group-number</i> [ <b>ipv4</b>   <b>ipv6</b> ]                                                                                                                                       | Creates an HSRP group and enters HSRP configuration mode.                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Step 4</b> | switch(config-if-hsrp)# <b>authentication</b> { <b>text</b> <i>string</i>   <b>md5</b> { <b>key-chain</b> <i>key-chain</i>   <b>key-string</b> { <b>0</b>   <b>7</b> } <i>text</i> [ <b>timeout</b> <i>seconds</i> ]}} | Configures cleartext authentication for HSRP on this interface by using the <b>authentication text</b> command, or you can configure MD5 authentication for HSRP on this interface using the <b>authentication md5</b> command. If you configure MD5 authentication, you can use a key chain or key string. If you use a key string, you can optionally set the timeout for when HSRP only accepts a new key. The range is from 0 to 32767 seconds. |
| <b>Step 5</b> | switch(config-if-hsrp)# <b>copy running-config startup-config</b>                                                                                                                                                      | Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.                                                                                                                                                                                                                                                                                                                       |
| <b>Step 6</b> | (Optional) switch(config-if-hsrp)# <b>show hsrp</b> [ <b>group</b> <i>group-number</i> ]                                                                                                                               | Displays HSRP information.                                                                                                                                                                                                                                                                                                                                                                                                                          |

### Example

The following example shows how to configure MD5 authentication for HSRP on Ethernet 1/2 after creating the key chain:

```
switch# configure terminal
switch (config)# interface ethernet 1/2
switch(config)# key chain hsrp-keys
switch(config-keychain)# key 0
switch(config-keychain-key)# key-string 7 zqdest
switch(config-keychain-key) accept-lifetime 00:00:00 Jun 01 2010 23:59:59 Sep 12 2010
switch(config-keychain-key) send-lifetime 00:00:00 Jun 01 2010 23:59:59 Aug 12 2010
switch(config-keychain-key) key 1
switch(config-keychain-key) key-string 7 uaeqdyito
switch(config-keychain-key) accept-lifetime 00:00:00 Aug 12 2010 23:59:59 Dec 12 2010
switch(config-keychain-key) send-lifetime 00:00:00 Sep 12 2010 23:59:59 Nov 12 2010
switch(config-keychain-key)# interface ethernet 1/2

switch(config-if)# hsrp 2
switch(config-if-hsrp)# authenticate md5 key-chain hsrp-keys
switch(config-if-hsrp)# copy running-config startup-config
```

## Configuring HSRP Object Tracking

You can configure an HSRP group to adjust its priority based on the availability of other interfaces or routes. The priority of a device can change dynamically if it has been configured for object tracking and the object that is being tracked goes down.

The tracking process periodically polls the tracked objects and notes any value change. The value change triggers HSRP to recalculate the priority. The HSRP interface with the higher priority becomes the active router if you configure the HSRP interface for preemption.

### Procedure

|               | Command or Action                                                                              | Purpose                                                                                                                                                                                |
|---------------|------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configure terminal</b>                                                              | Enters global configuration mode.                                                                                                                                                      |
| <b>Step 2</b> | Perform one of the following tasks:                                                            |                                                                                                                                                                                        |
| <b>Step 3</b> | switch(config)# <b>interface</b> <i>interface-type</i><br><i>slot/port</i>                     | Enters interface configuration mode.                                                                                                                                                   |
| <b>Step 4</b> | switch(config-if)# <b>hsrp</b> <i>group-number</i> [ <b>ipv4</b>   <b>ipv6</b> ]               | Creates an HSRP group and enters hsrp configuration mode.                                                                                                                              |
| <b>Step 5</b> | switch(config-if-hsrp)# <b>priority</b> [ <i>value</i> ]                                       | Sets the priority level used to select the active router in an HSRP group. The range is from 0 to 255. The default is 100.                                                             |
| <b>Step 6</b> | switch(config-if-hsrp)# <b>track</b> <i>object-number</i><br>[ <b>decrement</b> <i>value</i> ] | Specifies an object to be tracked that affects the weighting of an HSRP interface.<br><br>The value argument specifies a reduction in the priority of an HSRP interface when a tracked |

|               | Command or Action                                                                                                                   | Purpose                                                                                                                                                                                                         |
|---------------|-------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|               |                                                                                                                                     | object fails. The range is from 1 to 255. The default is 10.                                                                                                                                                    |
| <b>Step 7</b> | switch(config-if-hsrp)# <b>preempt</b> [ <b>delay</b> [ <b>minimum seconds</b> ] [ <b>reload seconds</b> ] [ <b>sync seconds</b> ]] | Configures the router to take over as the active router for an HSRP group if it has a higher priority than the current active router. This command is disabled by default. The range is from 0 to 3600 seconds. |
| <b>Step 8</b> | switch(config-if-hsrp)# <b>copy running-config startup-config</b>                                                                   | Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.                                                                                   |
| <b>Step 9</b> | (Optional) switch(config-if-hsrp)# <b>show hsrp interface interface-type number</b>                                                 | Displays HSRP information for an interface.                                                                                                                                                                     |

### Example

The following example shows how to configure HSRP object tracking on Ethernet 1/2:

```
switch# configure terminal
switch(config)# track 1 interface ethernet 2/2 line-protocol
switch(config)# interface ethernet 1/2

switch(config-if)# hsrp 2
switch(config-if-hsrp)# track 1 decrement 20
switch(config-if-hsrp)# copy running-config startup-config
```

## Configuring the HSRP Priority

You can configure the HSRP priority on an interface. HSRP uses the priority to determine which HSRP group member acts as the active router. If you configure HSRP on a vPC-enabled interface, you can optionally configure the upper and lower threshold values to control when to fail over to the vPC trunk. If the standby router priority falls below the lower threshold, HSRP sends all standby router traffic across the vPC trunk to forward through the active HSRP router. HSRP maintains this scenario until the standby HSRP router priority increases above the upper threshold.

For IPv6 HSRP groups, if all group members have the same priority, HSRP selects the active router based on the IPv6 link-local address.

For IPv4 HSRP groups, HSRP selects the active router based on the interface IP address when the priority is same.



**Note** Prior to Cisco NX-OS Release 7.2(0)D1(1), if the HSRP peer has a higher source interface IP address than the existing HSRP active peer and if preemption is enabled, the HSRP peer that has the same priority as the existing HSRP active peer preempts the existing HSRP active peer in the network.

After Cisco NX-OS Release 7.2(0)D1(1), even if the HSRP peer has a higher source interface IP address than the existing HSRP active peer and if preemption is enabled, the HSRP peer that has the same priority as the existing HSRP active peer does not preempt the existing HSRP active peer in the network.

### Procedure

|               | Command or Action                                                                                                                                   | Purpose                                                                                                                                                                                                                                                                                                                                                                                                   |
|---------------|-----------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configure terminal</b>                                                                                                                   | Enters global configuration mode.                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Step 2</b> | switch(config)# <b>interface</b> <i>type number</i>                                                                                                 | Enters interface configuration mode.                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Step 3</b> | switch(config-if)# <b>hsrp</b> <i>group-number</i> [ <b>ipv4</b> ]                                                                                  | Creates an HSRP group and enters hsrp configuration mode. The range for HSRP version 1 is from 0 to 255. The range is for HSRP version 2 is from 0 to 4095. The default value is 0.                                                                                                                                                                                                                       |
| <b>Step 4</b> | switch(config-if-hsrp)# <b>priority</b> <i>level</i><br>[ <b>forwarding-threshold lower</b> <i>lower-value</i><br><b>upper</b> <i>upper-value</i> ] | Sets the priority level used to select the active router in an HSRP group in interface configuration mode. The level range is from 0 to 255. The default is 100. Optionally, sets the upper and lower threshold values used by vPC to determine when to fail over to the vPC trunk. The lower-value range is from 1 to 255. The default is 1. The upper-value range is from 1 to 255. The default is 255. |
| <b>Step 5</b> | switch(config-if-hsrp)# <b>copy running-config startup-config</b>                                                                                   | Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.                                                                                                                                                                                                                                                                             |

### Example

The following example configures the HSRP priority level to 60 and the upper and lower threshold values used by vPC:

```
switch# configure terminal
switch(config)# interface ethernet 1/2
switch(config-if)# hsrp 2
switch(config-if-hsrp)# priority 60 forwarding-threshold lower 40 upper 50
switch(config-if-hsrp)# copy running-config startup-config
```

## Customizing HSRP in HSRP Configuration Mode

You can optionally customize the behavior of HSRP. Be aware that as soon as you enable an HSRP group by configuring a virtual IP address, that group is now operational. If you first enable an HSRP group before customizing HSRP, the router could take control over the group and become the active router before you finish customizing the feature. If you plan to customize HSRP, you should do so before you enable the HSRP group. To customize HSRP, use the following commands in HSRP configuration mode.

## Procedure

|               | Command or Action                                                                                                                                        | Purpose                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|---------------|----------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch(config-if-hsrp)# <b>name</b> <i>string</i>                                                                                                        | Specifies the IP redundancy name for an HSRP group. The string is from 1 to 255 characters. The default string has the following format:<br><br>hsrp-interface short-name group-id. For example, hsrp-Eth2/1-1.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Step 2</b> | switch(config-if-hsrp)# <b>preempt</b> [ <b>delay</b> [ <b>minimum</b> <i>seconds</i> ] [ <b>reload</b> <i>seconds</i> ] [ <b>sync</b> <i>seconds</i> ]] | Configures the router to take over as an active router for an HSRP group if it has a higher priority than the current active router. This command is disabled by default. The range is from 0 to 3600 seconds.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Step 3</b> | switch(config-if-hsrp)# <b>timers</b> [ <b>msec</b> ] <i>hellotime</i> [ <b>msec</b> ] <i>holdtime</i>                                                   | Configures the hello and hold time for this HSRP member as follows: <ul style="list-style-type: none"> <li>• <b>hellotime</b>—The interval between successive hello packets sent. The range is from 1 to 254 seconds.</li> <li>• <b>holdtime</b>—The interval before the information in the hello packet is considered invalid. The range is from 3 to 255.</li> </ul> <p>The optional msec keyword specifies that the argument is expressed in milliseconds instead of the default seconds. The timer ranges for milliseconds are as follows:</p> <ul style="list-style-type: none"> <li>• <b>hellotime</b>—The interval between successive hello packets sent. The range is from 255 to 999 milliseconds.</li> <li>• <b>holdtime</b>—The interval before the information in the hello packet is considered invalid. The range is from 750 to 3000 milliseconds.</li> </ul> |

## Example

The following example shows how to customize HSRP in HSRP configuration mode:

```
switch(config-if-hsrp) # name HSRP-1
switch(config-if-hsrp) # preempt delay minimum 60
switch(config-if-hsrp) # timers 5 18
```

## Customizing HSRP in Interface Configuration Mode

You can optionally customize the behavior of HSRP. Be aware that as soon as you enable an HSRP group by configuring a virtual IP address, that group is now operational. If you first enable an HSRP group before customizing HSRP, the router could take control over the group and become the active router before you finish customizing the feature. If you plan to customize HSRP, you should do so before you enable the HSRP group. To customize HSRP, use the following commands in interface configuration mode.

### Procedure

|               | Command or Action                                                | Purpose                                                                                                                                                                  |
|---------------|------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configure terminal</b>                                | Enters global configuration mode.                                                                                                                                        |
| <b>Step 2</b> | switch(config)# <b>interface</b> <i>interface-type slot/port</i> | Enters interface configuration mode.                                                                                                                                     |
| <b>Step 3</b> | switch(config-if)# <b>hsrp delay minimum</b> <i>seconds</i>      | Specifies the minimum amount of time that HSRP waits after a group is enabled before participating in the group. The range is from 0 to 10000 seconds. The default is 0. |
| <b>Step 4</b> | switch(config-if)# <b>hsrp delay reload</b> <i>seconds</i>       | Specifies the minimum amount of time that HSRP waits after reload before participating in the group. The range is from 0 to 10000 seconds. The default is 0.             |

### Example

The following example shows how to customize HSRP in interface configuration mode:

```
switch # configure terminal
switch(config)# interface ethernet 1/2
switch(config-if)# hsrp delay minimum 30
switch(config-if)# hsrp delay reload 30
```

## Configuring Extended Hold Timers for HSRP

You can configure HSRP to use extended hold timers to support extended NSF during a controlled (graceful) switchover or ISSU, including software upgrades and supervisor switchovers. You should configure extended hold timers on all HSRP routers.

You must configure extended hold timers on all HSRP routers if you configure extended hold timers. If you configure a nondefault hold timer, you should configure the same value on all HSRP routers when you configure HSRP extended hold timers.

HSRP extended hold timers are not applied if you configure millisecond hello and hold timers for HSRPv1. This statement does not apply to HSRPv2.



**Procedure**

|               | <b>Command or Action</b>                                             | <b>Purpose</b>                                                                                                                                                                                                                                                            |
|---------------|----------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configure terminal</b>                                    | Enters global configuration mode.                                                                                                                                                                                                                                         |
| <b>Step 2</b> | switch(config)# <b>hsrp timers extended-hold</b><br>[ <i>timer</i> ] | Sets the HSRP extended hold timer, in seconds, for both IPv4 and IPv6 groups. The timer range is from 10 to 255. The default is 10.<br><br><b>Note</b> Use the <b>show hsrp</b> command or the <b>show running-config hsrp</b> command to display the extended hold time. |
| <b>Step 3</b> | (Optional) switch(config)# <b>copy running-config startup-config</b> | Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.                                                                                                                                             |

**Example**

The following example shows how to configure extended hold timers for HSRP:

```
switch # configure terminal
switch(config)# hsrp timers extended-hold
switch(config)# copy running-config startup-config
```

## Verifying the HSRP Configuration

Use one of the following commands to verify the configuration:

| <b>Command</b>                                                                                                                                                                                                                                  | <b>Purpose</b>                                                                                                                                                                                               |
|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>show hsrp</b> [ <b>group</b> <i>group-number</i> ]                                                                                                                                                                                           | Displays the HSRP status for all groups or one group.                                                                                                                                                        |
| <b>show hsrp delay</b> [ <b>interface</b> <i>interface-type slot/port</i> ]                                                                                                                                                                     | Displays the HSRP delay value for all interfaces or one interface.                                                                                                                                           |
| <b>show hsrp</b> [ <b>interface</b> <i>interface-type slot/port</i> ]                                                                                                                                                                           | Displays the HSRP status for an interface.                                                                                                                                                                   |
| <b>show hsrp</b> [ <b>group</b> <i>group-number</i> ] [ <b>interface</b> <i>interface-type slot/port</i> ] [ <b>active</b> ] [ <b>all</b> ] [ <b>init</b> ] [ <b>learn</b> ] [ <b>listen</b> ] [ <b>speak</b> ] [ <b>standby</b> ]              | Displays the HSRP status for a group or interface for virtual forwarders in the active, init, learn, listen, or standby state. Use the all keyword to see all states, including disabled.                    |
| <b>show hsrp</b> [ <b>group</b> <i>group-number</i> ] [ <b>interface</b> <i>interface-type slot/port</i> ] [ <b>active</b> ] [ <b>all</b> ] [ <b>init</b> ] [ <b>learn</b> ] [ <b>listen</b> ] [ <b>speak</b> ] [ <b>standby</b> ] <b>brief</b> | Displays a brief summary of the HSRP status for a group or interface for virtual forwarders in the active, init, learn, listen, or standby state. Use the all keyword to see all states, including disabled. |
| <b>show hsrp mgo</b> [ <b>namename</b> ] [ <b>brief</b> ]                                                                                                                                                                                       | Displays the relationships between HSRP groups that are in use for MGO and their slave sessions.                                                                                                             |

## Configuration Examples for HSRP

This example shows how to enable HSRP on an interface with MD5 authentication and interface tracking:

```
key chain hsrp-keys
 key 0
 key-string 7 zqdest
 accept-lifetime 00:00:00 Jun 01 2008 23:59:59 Sep 12 2008
 send-lifetime 00:00:00 Jun 01 2008 23:59:59 Aug 12 2008
 key 1
 key-string 7 uaeqdyito
 accept-lifetime 00:00:00 Aug 12 2008 23:59:59 Dec 12 2008
 send-lifetime 00:00:00 Sep 12 2008 23:59:59 Nov 12 2008
feature hsrp
track 2 interface ethernet 2/2 ip
interface ethernet 1/2
 ip address 192.0.2.2/8
 hsrp 1
 authenticate md5 key-chain hsrp-keys
 priority 90
 track 2 decrement 20
 ip 192.0.2.10
no shutdown
```

This example shows how to configure the HSRP priority on an interface:

```
interface vlan 1
 hsrp 0
 preempt
 priority 100 forwarding-threshold lower 80 upper 90
 ip 192.0.2.2
 track 1 decrement 30
```

## Related Documents for HSRP

| Related Topic                                      | Document Title                                                              |
|----------------------------------------------------|-----------------------------------------------------------------------------|
| Configuring the Gateway Load Balancing protocol    | <a href="#">Configuring GLBP</a>                                            |
| Configuring the Virtual Router Redundancy protocol | <a href="#">Configuring VRRP</a>                                            |
| HSRP CLI commands                                  | <i>Cisco Nexus 7000 Series NX-OS Unicast Routing Command Reference</i>      |
| Configuring high availability                      | <i>Cisco Nexus 7000 Series NX-OS High Availability and Redundancy Guide</i> |

## MIBs

| Related Topic  | Document Title                                                                                                                                                                                     |
|----------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| CISCO-HSRP-MIB | To locate and download MIBs, go to the following URL:<br><a href="http://www.cisco.com/public/sw-center/netmgmt/cmtk/mibs.shtml">http://www.cisco.com/public/sw-center/netmgmt/cmtk/mibs.shtml</a> |

## Feature History for HSRP

This table includes only the updates for those releases that have resulted in additions or changes to the feature.

| Feature Name            | Releases | Feature Information                                                                                                      |
|-------------------------|----------|--------------------------------------------------------------------------------------------------------------------------|
| MGO                     | 6.2(2)   | This feature was introduced.                                                                                             |
| FabricPath anycast HSRP | 6.2(2)   | This feature was introduced.                                                                                             |
| BFD                     | 5.0(2)   | Added support for BFD. See the <i>Cisco Nexus 7000 Series NX-OS Interfaces Configuration Guide</i> for more information. |
| IPv6                    | 5.0(2)   | Added support to IPv6.                                                                                                   |
| Object track lists      | 4.2(1)   | Added support for object track lists.                                                                                    |
| Extended hold timers    | 4.2(1)   | Added support for extended hold timers for extended NSF support.                                                         |
| CISCO-HSRP-MIB          | 4.2(1)   | Added support for CISCO-HSRP-MIB                                                                                         |
| Priority thresholds     | 4.1(3)   | Added support for vPC threshold values on HSRP priority.                                                                 |
| HSRP                    | 4.0(1)   | This feature was introduced.                                                                                             |





## CHAPTER 22

# Configuring VRRP

---

This chapter contains the following sections:

- [Finding Feature Information, on page 483](#)
- [Information About VRRP, on page 483](#)
- [Guidelines and Limitations for VRRP, on page 489](#)
- [Default Settings for VRRP Parameters, on page 490](#)
- [Configuring VRRP, on page 490](#)
- [Verifying the VRRP Configuration, on page 503](#)
- [Monitoring VRRP Statistics, on page 504](#)
- [Configuration Example for VRRP, on page 504](#)
- [Related Documents for VRRP, on page 506](#)
- [Feature History for VRRP, on page 506](#)

## Finding Feature Information

Your software release might not support all the features documented in this module. For the latest caveats and feature information, see the Bug Search Tool at <https://tools.cisco.com/bugsearch/> and the release notes for your software release. To find information about the features documented in this module, and to see a list of the releases in which each feature is supported, see the "New and Changed Information" chapter or the Feature History table in this chapter.

## Information About VRRP

VRRP allows for transparent failover at the first-hop IP router by configuring a group of routers to share a virtual IP address. VRRP selects a master router in that group to handle all packets for the virtual IP address. The remaining routers are in standby and take over if the master router fails.

## VRRP Operation

A LAN client can determine which router should be the first hop to a particular remote destination by using a dynamic process or static configuration. Examples of dynamic router discovery are as follows:

Proxy ARP—The client uses Address Resolution Protocol (ARP) to get the destination it wants to reach, and a router responds to the ARP request with its own MAC address.

Routing protocol—The client listens to dynamic routing protocol updates (for example, from Routing Information Protocol [RIP]) and forms its own routing table.

ICMP Router Discovery Protocol (IRDP) client—The client runs an Internet Control Message Protocol (ICMP) router discovery client.

The disadvantage to dynamic discovery protocols is that they incur some configuration and processing overhead on the LAN client. Also, if a router fails, the process of switching to another router can be slow.

An alternative to dynamic discovery protocols is to statically configure a default router on the client. Although, this approach simplifies client configuration and processing, it creates a single point of failure. If the default gateway fails, the LAN client is limited to communicating only on the local IP network segment and is cut off from the rest of the network.

VRRP can solve the static configuration problem by enabling a group of routers (a VRRP group) to share a single virtual IP address.




**Note** The VRRP IP address must be different to any physical IP address of the devices participating in the VRRP, otherwise the ARP or MAC entries will be corrupted and may cause forwarding problems.

You can then configure the LAN clients with the virtual IP address as their default gateway.

**Figure 49: Basic VRRP Topology**

This image shows a basic VLAN topology where routers A, B, and C form a VRRP group. The IP address of the VRRP group must be different to the address that was configured for the Ethernet interface of Router A (10.0.0.1).



Because the virtual IP address uses the IP address of the physical Ethernet interface of Router A, Router A is the master (also known as the IP address owner). As the master, Router A owns the virtual IP address of the VRRP group and forwards packets sent to this IP address. Clients 1 through 3 are configured with the default gateway IP address of 10.0.0.1.

Routers B and C function as backups. If the master fails, the backup router with the highest priority becomes the master and takes over the virtual IP address to provide uninterrupted service for the LAN hosts. When router A recovers, it becomes the master again.



---

**Note** In Cisco NX-OS Release 4.1(2) and later, packets received on a routed port destined for the VRRP virtual IP address terminates on the local router, regardless of whether that router is the master VRRP router or a backup VRRP router. This includes ping and Telnet traffic. Packets received on a Layer 2 (VLAN) interface destined for the VRRP virtual IP address terminates on the master router.

---

## VRRP Benefits

The benefits of VRRP are as follows:

- **Redundancy**—Enables you to configure multiple routers as the default gateway router, which reduces the possibility of a single point of failure in a network.
- **Load sharing**—Allows traffic to and from LAN clients to be shared by multiple routers. The traffic load is shared more equitably among available routers.
- **Multiple VRRP groups**—Supports up to 255 VRRP groups on a router physical interface if the platform supports multiple MAC addresses. Multiple VRRP groups enable you to implement redundancy and load sharing in your LAN topology.
- **Multiple IP addresses**—Allows you to manage multiple IP addresses, including secondary IP addresses. If you have multiple subnets configured on an Ethernet interface, you can configure VRRP on each subnet.
- **Preemption**—Enables you to preempt a backup router that has taken over for a failing master with a higher priority backup router that has become available.
- **Advertisement protocol**—Uses a dedicated Internet Assigned Numbers Authority (IANA) standard multicast address (224.0.0.18) for VRRP advertisements. This addressing scheme minimizes the number of routers that must service the multicasts and allows test equipment to accurately identify VRRP packets on a segment. IANA has assigned the IP protocol number 112 to VRRP.
- **VRRP tracking**—Ensures that the best VRRP router is the master for the group by altering VRRP priorities based on interface states.
- The benefits of VRRPv3 are as follows:
  - Interoperability in multi-vendor environments.
  - Support for the IPv4 and IPv6 address families.
  - Improved scalability through the use of VRRS pathways.

## Multiple VRRP Groups


You can configure up to 255 VRRP groups on a physical interface. The number of VRRP groups that a router interface can support depends on the following factors:

- Router processing capability
- Router memory capability

In a topology where multiple VRRP groups are configured on a router interface, the interface can act as a master for one VRRP group and as a backup for one or more other VRRP groups.

**Figure 50: Load Sharing and Redundancy VRRP Topology**

This image shows a LAN topology in which VRRP is configured so that Routers A and B share the traffic to and from clients 1 through 4. Routers A and B act as backups to each other if either router fails.



This topology contains two virtual IP addresses for two VRRP groups that overlap. For VRRP group 1, Router A is the owner of IP address 10.0.0.1 and is the master. Router B is the backup to Router A. Clients 1 and 2 are configured with the default gateway IP address of 10.0.0.1.

For VRRP group 2, Router B is the owner of IP address 10.0.0.2 and is the master. Router A is the backup to router B. Clients 3 and 4 are configured with the default gateway IP address of 10.0.0.2.

## VRRP Router Priority and Preemption

An important aspect of the VRRP redundancy scheme is the VRRP router priority because the priority determines the role that each VRRP router plays and what happens if the master router fails.

If a VRRP router owns the virtual IP address and the IP address of the physical interface, this router functions as the master. The priority of the master is 255.

Priority also determines if a VRRP router functions as a backup router and the order of ascendancy to becoming a master if the master fails.

For example, if Router A, the master in a LAN topology, fails, VRRP must determine if backups B or C should take over. If you configure Router B with priority 101 and Router C with the default priority of 100, VRRP selects Router B to become the master because it has the higher priority. If you configure routers B and C with the default priority of 100, VRRP selects the backup with the higher IP address to become the master.

VRRP uses preemption to determine what happens after a VRRP backup router becomes the master. With preemption enabled by default, VRRP switches to a backup if that backup comes online with a priority higher than the new master. For example, if Router A is the master and fails, VRRP selects Router B (next in order of priority). If Router C comes online with a higher priority than Router B, VRRP selects Router C as the new master, even though Router B has not failed.

If you disable preemption, VRRP switches only if the original master recovers or the new master fails.



## vPC and VRRP

VRRP interoperates with virtual port channels (vPCs). vPCs allow links that are physically connected to two different Cisco Nexus 7000 series devices to appear as a single port channel by a third device. See the *Cisco Nexus 7000 Series NX-OS Layer 2 Switching Configuration Guide*, for more information on vPCs.

vPC forwards traffic through both the master VRRP router as well as the backup VRRP router.



---

**Note** You should configure VRRP on the primary vPC peer device as active and VRRP on the vPC secondary device as standby.

---

## VRRP Advertisements

The VRRP master sends VRRP advertisements to other VRRP routers in the same group. The advertisements communicate the priority and state of the master. Cisco NX-OS encapsulates the VRRP advertisements in IP packets and sends them to the IP multicast address assigned to the VRRP group. Cisco NX-OS sends the advertisements once every second by default, but you can configure a different advertisement interval.

## VRRP Authentication

VRRP supports the following authentication functions:

- No authentication
- Plain text authentication

VRRP rejects packets in any of the following cases:

- The authentication schemes differ on the router and in the incoming packet.
- Text authentication strings differ on the router and in the incoming packet.

## VRRP Tracking

VRRP supports the following two options for tracking:

- Native interface tracking—Tracks the state of an interface and uses that state to determine the priority of the VRRP router in a VRRP group. The tracked state is down if the interface is down or if the interface does not have a primary IP address.
- Object tracking—Tracks the state of a configured object and uses that state to determine the priority of the VRRP router in a VRRP group.

If the tracked state (interface or object) goes down, VRRP updates the priority based on what you configure the new priority to be for the tracked state. When the tracked state comes up, VRRP restores the original priority for the virtual router group.

For example, you may want to lower the priority of a VRRP group member if its uplink to the network goes down so another group member can take over as master for the VRRP group.



---

**Note** VRRP does not support Layer 2 interface tracking.

---

## VRRPv3 and VRRS

VRRP version 3 (VRRPv3) enables a group of switches to form a single virtual switch in order to provide redundancy and reduce the possibility of a single point of failure in a network. The LAN clients can then be configured with the virtual switch as their default gateway. The virtual switch, representing a group of switches, is also known as a VRRPv3 group.

Virtual router redundancy service (VRRS) improves the scalability of VRRPv3 by providing a stateless redundancy service to VRRS pathways and VRRS clients by monitoring VRRPv3. VRRPv3 acts as a VRRS server that pushes VRRPv3 status information (such as current and previous redundancy states, active and inactive Layer 2 and Layer 3 addresses, and so on) to VRRS pathways and all registered VRRS clients.

VRRS clients are other Cisco processes or applications that use VRRPv3 to provide or withhold a service or resource dependent upon the state of the group. VRRS pathways are special VRRS clients that use the VRRS database information to provide scaled first-hop gateway redundancy across scaled interface environments.

VRRS by itself is limited to maintaining its own state. Linking a VRRS client to a VRRPv3 group provides a mechanism that allows VRRS to provide a service to client applications so that they can implement stateless or stateful failovers. A stateful failover requires communication with a nominated backup before the failure so that operational data is not lost when the failover occurs.

VRRS pathways operate in a similar way to clients but are integrated with the VRRS architecture. They provide a means to scale first-hop gateway redundancy by allowing you to configure a virtual address across hundreds of interfaces. The virtual gateway state of a VRRS pathway follows the state of a First-Hop Redundancy Protocol (FHRP) VRRS server.

VRRPv3 notifies VRRS of its current state (master, backup, or nonoperational initial state [INIT]) and passes that information to pathways or clients. The VRRPv3 group name activates VRRS and associates the VRRPv3 group with any clients or pathways that are configured as part of VRRS with the same name.

Pathways and clients act on the VRRPv3 server state. When a VRRPv3 group changes states, VRRS pathways and clients alter their behavior (performing tasks such as shutting down interfaces or appending accounting logs) depending on the state received from VRRS.

## BFD for VRRP

This feature supports bidirectional forwarding detection (BFD). BFD is a detection protocol that provides fast-forwarding and path-failure detection times. BFD provides subsecond failure detection between two adjacent devices and can be less CPU-intensive than protocol hello messages because some of the BFD load can be distributed onto the data plane on supported modules. See the *Cisco Nexus 7000 Series NX-OS Interfaces Configuration Guide*, for more information.

## High Availability

VRRP supports high availability through stateful restarts and stateful switchovers. A stateful restart occurs when the VRRP process fails and is restarted. Stateful switchover occurs when the active supervisor switches to the standby supervisor. Cisco NX-OS applies the run-time configuration after the switchover.

VRRPv3 does not support stateful switchovers.

## Virtualization Support

VRRP supports virtual routing and forwarding (VRF) instances. VRF exists within virtual device contexts (VDCs). By default, Cisco NX-OS places you in the default VDC and default VRF unless you specifically configure another VDC and VRF. If you change the VRF membership of an interface, Cisco NX-OS removes all Layer 3 configurations, including VRRP.

For more information, see the *Cisco Nexus 7000 Series NX-OS Virtual Device Context Configuration Guide* and [Configuring Layer 3 Virtualization, on page 371](#).

## Guidelines and Limitations for VRRP

- You cannot configure VRRP on the management interface.
- When VRRP is enabled, you should replicate the VRRP configuration across devices in your network.
- We recommend that you do not configure more than one first-hop redundancy protocol on the same interface.
- You must configure an IP address for the interface where you configure VRRP and enable that interface before VRRP becomes active.
- Cisco NX-OS removes all Layer 3 configurations on an interface when you change the interface VRF membership, port channel membership, or when you change the port mode to Layer 2.
- When you configure VRRP to track a Layer 2 interface, you must shut down the Layer 2 interface and reenabling the interface to update the VRRP priority to reflect the state of the Layer 2 interface.
- BFD for VRRP can only be configured between two routers.
- The VRRP IP address must be different than any physical IP address of the devices participating in the VRRP, otherwise the ARP or MAC entries will be corrupted and may cause forwarding problems.
- VRRPv3 has the following configuration guidelines and limitations:
  - VRRPv3 is not intended as a replacement for existing dynamic protocols. VRRPv3 is designed for use over multi-access, multicast, or broadcast-capable Ethernet LANs.
  - VRRPv3 is supported only on Ethernet and Fast Ethernet interfaces, bridge group virtual interfaces (BVI), and Gigabit Ethernet interfaces as well as on Multiprotocol Label Switching (MPLS) virtual private networks (VPNs), VRF-aware MPLS VPNs, and VLANs.
  - When VRRPv3 is in use, VRRPv2 is unavailable. To configure VRRPv3, you must disable any VRRPv2 configuration.
  - VRRS is currently available only for use with VRRPv3.
  - Use VRRPv3 millisecond timers only where absolutely necessary and with careful consideration and testing. Millisecond values work only under favorable circumstances. The millisecond timer values are compatible with third-party vendors, as long as they also support VRRPv3.
  - Full network redundancy can be achieved only if VRRPv3 operates over the same network path as the VRRS pathway redundant interfaces. For full redundancy, the following restrictions apply:

- VRRS pathways should use the same physical interface as the parent VRRPv3 group or be configured on a subinterface with the same physical interface as the parent VRRPv3 group.
- VRRS pathways can be configured on switch virtual interfaces (SVIs) only if the associated VLAN shares the same trunk as the VLAN on which the parent VRRPv3 group is configured.

## Default Settings for VRRP Parameters

### Default RIP Parameters

| Parameters                        | Default           |
|-----------------------------------|-------------------|
| Advertisement interval            | 1 second          |
| Authentication                    | No authentication |
| Preemption                        | Enabled           |
| Priority                          | 100               |
| VRRP feature                      | Disabled          |
| VRRPv3                            | Disabled          |
| VRRS                              | Disabled          |
| VRRPv3 secondary address matching | Enables           |
| Priority of a VRRPv3 group        | 100               |
| VRRPv3 advertisement timer        | 1000 milliseconds |

## Configuring VRRP

### Enabling VRRP

You must globally enable the VRRP feature before you configure and enable any VRRP groups.

#### Procedure

|               | Command or Action                        | Purpose                                                                                          |
|---------------|------------------------------------------|--------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configure terminal</b>        | Enters global configuration mode.                                                                |
| <b>Step 2</b> | switch(config)# <b>[no] feature vrrp</b> | Enables the VRRP feature.<br><br>Use the <b>no</b> form of this command to disable this feature. |

|               | Command or Action                                                    | Purpose                                                                                                                       |
|---------------|----------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------|
|               |                                                                      | Using the no form of this command will disable the feature in a VDC and remove all associated configurations.                 |
| <b>Step 3</b> | (Optional) switch(config)# <b>copy running-config startup-config</b> | Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration. |

### Example

The following example enables VRRP:

```
switch # configure terminal
switch(config)# feature vrrp
switch(config)# copy running-config startup-config
```

## Configuring VRRP Groups

You can create a VRRP group, assign the virtual IP address, and enable the group.

You can configure one virtual IPv4 address for a VRRP group. By default, the master VRRP router drops the packets addressed directly to the virtual IP address because the VRRP master is only intended as a next-hop router to forward packets. Some applications require that Cisco NX-OS accept packets addressed to the virtual router IP. Use the secondary option to the virtual IP address to accept these packets when the local router is the VRRP master.

Once you have configured the VRRP group, you must explicitly enable the group before it becomes active.

### Before you begin

Confirm that you are in the correct VDC. To change the VDC, use the **switchto vdc** command.

Ensure that you have configured an IP address on the interface.

### Procedure

|               | Command or Action                                                    | Purpose                                                                                                                                           |
|---------------|----------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configure terminal</b>                                    | Enters global configuration mode.                                                                                                                 |
| <b>Step 2</b> | switch(config)# <b>interface</b> <i>interface-type slot/port</i>     | Enters interface configuration mode.                                                                                                              |
| <b>Step 3</b> | switch(config-if)# <b>vrrp</b> <i>number</i>                         | Creates a virtual router group. the range is from 1 to 255.                                                                                       |
| <b>Step 4</b> | switch(config-if-vrrp)# <b>address</b> <i>ip-address</i> [secondary] | Configures the virtual IPv4 address for the specified VRRP group. This address should be in the same subnet as the IPv4 address of the interface. |

|               | Command or Action                                                 | Purpose                                                                                                                                                         |
|---------------|-------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------|
|               |                                                                   | Use the secondary option only if applications require that VRRP routers accept the packets sent to the virtual router's IP address and deliver to applications. |
| <b>Step 5</b> | switch(config-if-vrrp)# <b>no shutdown</b>                        | Enables the VRRP group. Disabled by default.                                                                                                                    |
| <b>Step 6</b> | switch(config-if-vrrp)# <b>copy running-config startup-config</b> | Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.                                   |
| <b>Step 7</b> | (Optional) switch(config-if-vrrp)# <b>show vrrp</b>               | (Optional) Displays VRRP information.                                                                                                                           |

### Example

The following example enables VRRP:

```
switch # configure terminal
switch(config)# interface ethernet 2/1

switch(config-if)# vrrp 250
switch(config-if-vrrp)# address 192.0.2.8
switch(config-if-vrrp)# no shutdown
switch(config-if-vrrp)# copy running-config startup-config
switch(config-if-vrrp)# show vrrp
```

## Configuring VRRP Priority

The valid priority range for a virtual router is from 1 to 254 (1 is the lowest priority and 254 is the highest). The default priority value for backups is 100. For devices whose interface IP address is the same as the primary virtual IP address (the master), the default value is 255.

If you configure VRRP on a vPC-enabled interface, you can optionally configure the upper and lower threshold values to control when to fail over to the vPC trunk. If the backup router priority falls below the lower threshold, VRRP sends all backup router traffic across the vPC trunk to forward through the master VRRP router. VRRP maintains this scenario until the backup VRRP router priority increases above the upper threshold.

### Before you begin

- Confirm that you are in the correct VDC. To change the VDC, use the **switchto vdc** command.
- Ensure that you have configured an IP address on the interface.
- You must enable VRRP.

### Procedure

|               | Command or Action                 | Purpose                           |
|---------------|-----------------------------------|-----------------------------------|
| <b>Step 1</b> | switch# <b>configure terminal</b> | Enters global configuration mode. |

|               | Command or Action                                                                                                                                          | Purpose                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|---------------|------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 2</b> | switch(config)# <b>interface</b> <i>interface-type slot/port</i>                                                                                           | Enters interface configuration mode.                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Step 3</b> | switch(config-if)# <b>vrrp</b> <i>number</i>                                                                                                               | Creates a virtual router group. the range is from 1 to 255.                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Step 4</b> | switch(config-if-vrrp)# <b>shutdown</b>                                                                                                                    | Disables the VRRP group. Disabled by default.                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Step 5</b> | switch(config-if-vrrp)# <b>priority</b> <i>level</i><br>[ <b>forwarding-threshold</b> <b>lower</b> <i>lower-value</i><br><b>upper</b> <i>upper-value</i> ] | Sets the priority level used to select the active router in a VRRP group. The level range is from 1 to 254. The default is 100 for backups and 255 for a master that has an interface IP address equal to the virtual IP address.<br><br>Optionally, sets the upper and lower threshold values used by vPC to determine when to fail over to the vPC trunk. The lower-value range is from 1 to 255. The default is 1. The upper-value range is from 1 to 255. The default is 255. |
| <b>Step 6</b> | switch(config-if-vrrp)# <b>no shutdown</b>                                                                                                                 | Enables the VRRP group. Disabled by default.                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Step 7</b> | switch(config-if-vrrp)# <b>copy running-config startup-config</b>                                                                                          | Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.                                                                                                                                                                                                                                                                                                                                                     |
| <b>Step 8</b> | (Optional) switch(config-if-vrrp)# <b>show vrrp</b>                                                                                                        | Displays VRRP information.                                                                                                                                                                                                                                                                                                                                                                                                                                                        |

### Example

The following example enables VRRP:

```
switch # configure terminal
switch(config)# interface ethernet 2/1

switch(config-if)# vrrp 250
switch(config-if)# shutdown
switch(config-if-vrrp)# priority 60 forwarding-threshold lower 40 upper 50
switch(config-if-vrrp)# no shutdown
switch(config)# copy running-config startup-config
switch(config-if-vrrp)# show vrrp
```

## Configuring VRRP Authentication

You can configure simple text authentication for a VRRP group.

### Before you begin

- Ensure that the authentication configuration is identical for all VRRP devices in the network.
- Confirm that you are in the correct VDC. To change the VDC, use the **switchto vdc** command.

- Ensure that you have configured an IP address on the interface. See [#unique\\_773](#).
- You must enable VRRP.

### Procedure

|               | Command or Action                                                  | Purpose                                                                                                                                                                                                                                      |
|---------------|--------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configure terminal</b>                                  | Enters global configuration mode.                                                                                                                                                                                                            |
| <b>Step 2</b> | switch(config)# <b>interface</b> <i>interface-type slot/port</i>   | Enters interface configuration mode.                                                                                                                                                                                                         |
| <b>Step 3</b> | switch(config-if)# <b>vrrp</b> <i>number</i>                       | Creates a virtual router group. The range is from 1 to 255.                                                                                                                                                                                  |
| <b>Step 4</b> | switch(config-if-vrrp)# <b>shutdown</b>                            | Disables the VRRP group. Disabled by default.                                                                                                                                                                                                |
| <b>Step 5</b> | switch(config-if-vrrp)# <b>authentication text</b> <i>password</i> | Assigns the simple text authentication option and specifies the keyname password. The keyname range is from 1 to 255 characters. We recommend that you use at least 16 characters. The text password is up to eight alphanumeric characters. |
| <b>Step 6</b> | switch(config-if-vrrp)# <b>no shutdown</b>                         | Enables the VRRP group. Disabled by default.                                                                                                                                                                                                 |
| <b>Step 7</b> | switch(config-if-vrrp)# <b>copy running-config startup-config</b>  | Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.                                                                                                                |
| <b>Step 8</b> | (Optional) switch(config-if-vrrp)# <b>show vrrp</b>                | Displays VRRP information.                                                                                                                                                                                                                   |

### Example

The following example enables VRRP:

```
switch # configure terminal
switch(config)# interface ethernet 2/1

switch(config-if)# vrrp 250
switch(config-if)# shutdown
switch(config-if-vrrp)# authentication text aPassword
switch(config-if-vrrp)# no shutdown
switch(config)# copy running-config startup-config
switch(config-if-vrrp)# show vrrp
```

## Configuring Time Intervals for Advertisement Packets

You can configure the time intervals for advertisement packets.

### Before you begin

- Confirm that you are in the correct VDC. To change the VDC, use the **switchto vdc** command.



- Ensure that you have configured an IP address on the interface.
- You must enable VRRP.

### Procedure

|               | Command or Action                                                    | Purpose                                                                                                                       |
|---------------|----------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configure terminal</b>                                    | Enters global configuration mode.                                                                                             |
| <b>Step 2</b> | switch(config)# <b>interface</b> <i>interface-type slot/port</i>     | Enters interface configuration mode.                                                                                          |
| <b>Step 3</b> | switch(config-if)# <b>vrrp</b> <i>number</i>                         | Creates a virtual router group. The range is from 1 to 255.                                                                   |
| <b>Step 4</b> | switch(config-if-vrrp)# <b>shutdown</b>                              | Disables the VRRP group. Disabled by default.                                                                                 |
| <b>Step 5</b> | switch(config-if-vrrp)# <b>advertisement interval</b> <i>seconds</i> | Sets the interval time in seconds between sending advertisement frames. The range is from 1 to 255. The default is 1 second.  |
| <b>Step 6</b> | switch(config-if-vrrp)# <b>no shutdown</b>                           | Enables the VRRP group. Disabled by default.                                                                                  |
| <b>Step 7</b> | switch(config-if-vrrp)# <b>copy running-config startup-config</b>    | Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration. |
| <b>Step 8</b> | (Optional) switch(config-if-vrrp)# <b>show vrrp</b>                  | Displays VRRP information.                                                                                                    |

### Example

The following example enables VRRP:

```
switch # configure terminal
switch(config)# interface ethernet 2/1

switch(config-if)# vrrp 250
switch(config-if)# shutdown
switch(config-if-vrrp)# advertisement-interval 15
switch(config-if-vrrp)# no shutdown
switch(config)# copy running-config startup-config
switch(config-if-vrrp)# show vrrp
```

## Disabling Preemption

You can disable preemption for a VRRP group member. If you disable preemption, a higher-priority backup router does not take over for a lower-priority master router. Preemption is enabled by default.

### Before you begin

- You must enable VRRP.
- Ensure that you have configured an IP address on the interface.

- Confirm that you are in the correct VDC. To change the VDC, use the **switchto vdc** command.

### Procedure

|               | Command or Action                                                 | Purpose                                                                                                                       |
|---------------|-------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configure terminal</b>                                 | Enters global configuration mode.                                                                                             |
| <b>Step 2</b> | switch(config)# <b>interface</b> <i>interface-type slot/port</i>  | Enters interface configuration mode.                                                                                          |
| <b>Step 3</b> | switch(config-if)# <b>vrrp</b> <i>number</i>                      | Creates a virtual router group. The range is from 1 to 255.                                                                   |
| <b>Step 4</b> | switch(config-if-vrrp)# <b>shutdown</b>                           | Disables the VRRP group. Disabled by default.                                                                                 |
| <b>Step 5</b> | switch(config-if-vrrp)# <b>no preempt</b>                         | Disables the preempt option and allows the master to remain when a higher-priority backup appears.                            |
| <b>Step 6</b> | switch(config-if-vrrp)# <b>no shutdown</b>                        | Enables the VRRP group. Disabled by default.                                                                                  |
| <b>Step 7</b> | switch(config-if-vrrp)# <b>copy running-config startup-config</b> | Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration. |
| <b>Step 8</b> | (Optional) switch(config-if-vrrp)# <b>show vrrp</b>               | Displays VRRP information.                                                                                                    |

### Example

The following example enables VRRP:

```
switch # configure terminal
switch(config)# interface ethernet 2/1

switch(config-if)# vrrp 250
switch(config-if)# shutdown
switch(config-if-vrrp)# no preempt
switch(config-if-vrrp)# no shutdown
switch(config)# copy running-config startup-config
switch(config-if-vrrp)# show vrrp
```

## Configuring VRRP Interface State Tracking

Interface state tracking changes the priority of the virtual router based on the state of another interface in the device. When the tracked interface goes down or the IP address is removed, Cisco NX-OS assigns the tracking priority value to the virtual router. When the tracked interface comes up and an IP address is configured on this interface, Cisco NX-OS restores the configured priority to the virtual router.



**Note** For interface state tracking to function, you must enable preemption on the interface.



**Note** VRRP does not support Layer 2 interface tracking.

### Before you begin

- Confirm that you are in the correct VDC. To change the VDC, use the **switchto vdc** command.
- Ensure that you have configured an IP address on the interface. See [#unique\\_773](#).
- You must enable VRRP.
- Ensure that you have enabled the virtual router.

### Procedure

|               | Command or Action                                                                | Purpose                                                                                                                       |
|---------------|----------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configure terminal</b>                                                | Enters global configuration mode.                                                                                             |
| <b>Step 2</b> | switch(config)# <b>interface</b> <i>interface-type slot/port</i>                 | Enters interface configuration mode.                                                                                          |
| <b>Step 3</b> | switch(config-if)# <b>vrrp</b> <i>number</i>                                     | Creates a virtual router group. The range is from 1 to 255.                                                                   |
| <b>Step 4</b> | switch(config-if-vrrp)# <b>shutdown</b>                                          | Disables the VRRP group. Disabled by default.                                                                                 |
| <b>Step 5</b> | switch(config-if-vrrp)# <b>track interface</b> <i>type number priority value</i> | Enables interface priority tracking for a VRRP group. The priority range is from 1 to 254.                                    |
| <b>Step 6</b> | switch(config-if-vrrp)# <b>no shutdown</b>                                       | Enables the VRRP group. Disabled by default.                                                                                  |
| <b>Step 7</b> | switch(config-if-vrrp)# <b>copy running-config startup-config</b>                | Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration. |
| <b>Step 8</b> | (Optional) switch(config-if-vrrp)# <b>show vrrp</b>                              | Displays VRRP information.                                                                                                    |

### Example

The following example enables VRRP:

```
switch # configure terminal
switch(config)# interface ethernet 2/1

switch(config-if)# vrrp 250
switch(config-if)# shutdown
switch(config-if-vrrp)# track interface ethernet 2/10 priority 254
switch(config-if-vrrp)# no shutdown
switch(config)# copy running-config startup-config
switch(config-if-vrrp)# show vrrp
```

## Enabling the VRRPv3 Feature

You must globally enable the VRRPv3 feature before you can configure and enable any VRRPv3 groups.

### Procedure

|               | Command or Action                     | Purpose                                                                                                                                                                                                                                                                                                                                                      |
|---------------|---------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configure terminal</b>     | Enters global configuration mode.                                                                                                                                                                                                                                                                                                                            |
| <b>Step 2</b> | switch(config)# <b>feature vrrpv3</b> | Enables VRRP version 3 and Virtual Router Redundancy Service (VRRS). The <b>no</b> form of this command disables VRRPv3 and VRRS in a VDC.<br><br>If VRRPv2 is currently configured, use the <b>no feature vrrp</b> command in global configuration mode to remove the VRRPv2 configuration and then use the <b>feature vrrpv3</b> command to enable VRRPv3. |

### Example

The following example shows how to enable VRRPv3:

```
switch# configure terminal
switch(config)# enable vrrpv3
```

## Creating VRRPv3 Groups

You can create a VRRPv3 group, assign the virtual IP address, and enable the group.

### Before you begin

- Ensure that the VRRPv3 feature is enabled.
- Ensure that you are in the correct VDC (or use the **switchto vdc** command).
- Ensure that you configure an IP address on the interface.

### Procedure

|               | Command or Action                                                                                     | Purpose                                                                                        |
|---------------|-------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configure terminal</b>                                                                     | Enters global configuration mode.                                                              |
| <b>Step 2</b> | switch(config)# <b>interface</b> <i>type/number</i>                                                   | Enters interface configuration mode.                                                           |
| <b>Step 3</b> | switch(config-if)# <b>vrrpv3</b> <i>number</i><br><b>address-family</b> [ <b>ipv4</b>   <b>ipv6</b> ] | Creates a VRRPv3 group and enters VRRPv3 group configuration mode. The range is from 1 to 255. |

|                | Command or Action                                                                                         | Purpose                                                                                                                                                                                                                                                                                                    |
|----------------|-----------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 4</b>  | switch(config-if-vrrpv3-group)# <b>address</b><br><i>ip-address</i> [ <b>primary</b>   <b>secondary</b> ] | (Optional)<br>Specifies a primary or secondary IPv4 or IPv6 address for the VRRPv3 group.<br><br>To utilize secondary IP addresses in a VRRPv3 group, you must first configure a primary IP address on the same group.                                                                                     |
| <b>Step 5</b>  | switch(config-if-vrrpv3-group)# <b>description</b><br><i>description</i>                                  | (Optional)<br>Specifies a description for the VRRPv3 group. You can enter up to 80 alphanumeric characters.                                                                                                                                                                                                |
| <b>Step 6</b>  | switch(config-if-vrrpv3-group)# <b>match-address</b>                                                      | (Optional)<br>Matches the secondary address in the advertisement packet against the configured address.                                                                                                                                                                                                    |
| <b>Step 7</b>  | switch(config-if-vrrpv3-group)# <b>preempt</b><br>[ <b>delay minimum seconds</b> ]                        | (Optional)<br>Enables preemption of a lower priority master switch with an optional delay. The range is from 0 to 3600.                                                                                                                                                                                    |
| <b>Step 8</b>  | switch(config-if-vrrpv3-group)# <b>priority level</b>                                                     | (Optional)<br>Specifies the priority of the VRRPv3 group. The range is from 1 to 254.                                                                                                                                                                                                                      |
| <b>Step 9</b>  | switch(config-if-vrrpv3-group)# <b>timers</b><br><b>advertise interval</b>                                | (Optional)<br>Sets the advertisement timer in milliseconds. The range is from 100 to 40950.<br><br>Cisco recommends that you set this timer to a value greater than or equal to 1 second.                                                                                                                  |
| <b>Step 10</b> | switch(config-if-vrrpv3-group)# <b>vrrp2</b>                                                              | (Optional)<br>Enables support for VRRPv2 simultaneously, to ensure interoperability with devices that support only VRRPv2.<br><br>VRRPv2 compatibility mode is provided to allow an upgrade from VRRPv2 to VRRPv3. This is not a full VRRPv2 implementation and should be used only to perform an upgrade. |
| <b>Step 11</b> | switch(config-if-vrrpv3-group)# <b>vrrs leader</b><br><i>vrrs-leader-name</i>                             | (Optional)<br>Specifies a leader's name to be registered with VRRS.                                                                                                                                                                                                                                        |
| <b>Step 12</b> | switch(config-if-vrrpv3-group)# <b>shutdown</b>                                                           | (Optional)                                                                                                                                                                                                                                                                                                 |

|                | Command or Action                                                                                                 | Purpose                                                                                                                               |
|----------------|-------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------|
|                |                                                                                                                   | Disables VRRP configuration for the VRRPv3 group.                                                                                     |
| <b>Step 13</b> | switch(config-if-vrrpv3-group)# <b>show fhrp</b><br>[ <i>interface-type interface-number</i> ] [ <b>verbose</b> ] | (Optional)<br>Displays First Hop Redundancy Protocol (FHRP) information. Use the <b>verbose</b> keyword to view detailed information. |
| <b>Step 14</b> | switch(config-if-vrrpv3-group)# <b>copy running-config startup-config</b>                                         | (Optional)<br>Saves this configuration change.                                                                                        |

### Example

The following example shows how to create a VRRPv3 group:

```
switch# configure terminal
switch(config)# interface ethernet 1/2
switch(config-if)# vrrpv3 5 address-family ipv4
switch(config-if)# hsrp version 2
switch(config-if-vrrpv3-group)# address 100.0.1.10 primary
switch(config-if-vrrpv3-group)# description group3
switch(config-if-vrrpv3-group)# match-address
switch(config-if-vrrpv3-group)# preempt delay minimum 30
switch(config-if-vrrpv3-group)# priority 3
switch(config-if-vrrpv3-group)# timers advertise 1000
switch(config-if-vrrpv3-group)# vrrp2
switch(config-if-vrrpv3-group)# vrrs leader leader1
switch(config-if-vrrpv3-group)# shutdown
switch(config-if-vrrpv3-group)# show fhrp ethernet 1/2 verbose
switch(config-if-vrrpv3-group)# show running-config startup-config
```

## Configuring the Delay Period for FHRP Client Initialization

### Procedure

|               | Command or Action                                                                           | Purpose                                                                                                                                                                                                                                                                                            |
|---------------|---------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configure terminal</b>                                                           | Enters global configuration mode.                                                                                                                                                                                                                                                                  |
| <b>Step 2</b> | switch(config)# <b>fhrp delay</b> {[ <b>minimum</b> ]<br>[ <b>reload</b> ] <i>seconds</i> } | Specifies the delay period for the initialization of FHRP clients. The range is from 0 to 3600 seconds.<br><br>The <b>minimum</b> keyword configures the delay period after an interface becomes available.<br><br>The <b>reload</b> command configures the delay period after the device reloads. |
| <b>Step 3</b> | switch(config)# <b>copy running-config startup-config</b>                                   | (Optional)<br>Saves this configuration change.                                                                                                                                                                                                                                                     |

**Example**

The following example shows how to configure the delay period for initializing FHRP clients:

```
switch# configure terminal
switch(config)# fhrp delay minimum 14
```

## Configuring VRRPv3 Control Groups

**Before you begin**

- Ensure that the VRRPv3 feature is enabled.
- Ensure that you are in the correct VDC (or use the **switchto vdc** command).
- Ensure that you configure an IP address on the interface.

**Procedure**

|               | Command or Action                                                                                                 | Purpose                                                                                                                                            |
|---------------|-------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configure terminal</b>                                                                                 | Enters global configuration mode.                                                                                                                  |
| <b>Step 2</b> | switch(config)# <b>interface</b> <i>type/number</i>                                                               | Enters interface configuration mode.                                                                                                               |
| <b>Step 3</b> | switch(config-if)# <b>ip address</b> <i>ip address mask</i><br>[ <b>secondary</b> ]                               | Configures the IP address on the interface.<br><br>You can use the <b>secondary</b> keyword to configure additional IP addresses on the interface. |
| <b>Step 4</b> | switch(config-if)# <b>vrrpv3</b> <i>number</i><br><b>address-family</b> [ <b>ipv4</b>   <b>ipv6</b> ]             | Creates a VRRPv3 group and enters VRRPv3 group configuration mode. The range is from 1 to 255.                                                     |
| <b>Step 5</b> | switch(config-if-vrrpv3-group)# <b>address</b><br><i>ip-address</i> [ <b>primary</b>   <b>secondary</b> ]         | (Optional)<br><br>Specifies a primary or secondary IPv4 or IPv6 address for the VRRPv3 group.                                                      |
| <b>Step 6</b> | switch(config-if-vrrpv3-group)# <b>vrrs leader</b><br><i>vrrs-leader-name</i>                                     | (Optional)<br><br>Specifies a leader's name to be registered with VRRS.                                                                            |
| <b>Step 7</b> | switch(config-if-vrrpv3-group)# <b>shutdown</b>                                                                   | (Optional)<br><br>Disables VRRP configuration for the VRRPv3 group.                                                                                |
| <b>Step 8</b> | switch(config-if-vrrpv3-group)# <b>show fhrp</b><br>[ <i>interface-type interface-number</i> ] [ <b>verbose</b> ] | (Optional)<br><br>Displays First Hop Redundancy Protocol (FHRP) information.                                                                       |

|               | Command or Action                                                               | Purpose                                                     |
|---------------|---------------------------------------------------------------------------------|-------------------------------------------------------------|
|               |                                                                                 | Use the <b>verbose</b> keyword to view detailed information |
| <b>Step 9</b> | <code>switch(config-if-vrrpv3-group)# copy running-config startup-config</code> | (Optional)<br>Saves this configuration change.              |

### Example

The following example shows how to configure a VRRPv3 control group:

```
switch# configure terminal
switch(config)# interface ethernet 1/2
switch(config-if)# ip address 209.165.200.230 255.255.255.224
switch(config-if)# vrrpv3 5 address-family ipv4
switch(config-if-vrrpv3-group)# address 209.165.200.227 primary
switch(config-if-vrrpv3-group)# vrrs leader leader1
switch(config-if-vrrpv3-group)# shutdown
switch(config-if-vrrpv3-group)# show fhrp ethernet 1/2 verbose
switch(config-if-vrrpv3-group)# show running-config startup-config
```

## Configuring VRRS Pathways

You can configure a Virtual Router Redundancy Service (VRRS) pathway. In scaled environments, VRRS pathways should be used in combination with VRRPv3 control groups.

### Before you begin

- Ensure that the VRRPv3 feature is enabled.
- Ensure that you are in the correct VDC (or use the **switchto vdc** command).
- Ensure that you configure an IP address on the interface.

### Procedure

|               | Command or Action                                                      | Purpose                                                                                                                                            |
|---------------|------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | <code>switch# configure terminal</code>                                | Enters global configuration mode.                                                                                                                  |
| <b>Step 2</b> | <code>switch(config)# interface type/number</code>                     | Enters interface configuration mode.                                                                                                               |
| <b>Step 3</b> | <code>switch(config-if)# ip address ip-address mask [secondary]</code> | Configures the IP address on the interface.<br><br>You can use the <b>secondary</b> keyword to configure additional IP addresses on the interface. |
| <b>Step 4</b> | <code>switch(config-if)# vrrs pathway vrrs-tag</code>                  | Defines the VRRS pathway for a VRRS group and enters VRRS pathway configuration mode.                                                              |



|               | Command or Action                                                                         | Purpose                                                                                                                                                                                    |
|---------------|-------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|               |                                                                                           | The <i>vrrs-tag</i> argument specifies the name of the VRRS tag that is being associated with the pathway.                                                                                 |
| <b>Step 5</b> | <code>switch(config-if-vrrs-pw)# mac address {mac-address   inherit}</code>               | Specifies a MAC address for the pathway.<br><br>The <b>inherit</b> keyword causes the pathway to inherit the virtual MAC address of the VRRPv3 group with which the pathway is associated. |
| <b>Step 6</b> | <code>switch(config-if-vrrs-pw)# address ip-address</code>                                | Defines the virtual IPv4 or IPv6 address for a pathway.<br><br>A VRRPv3 group is capable of controlling more than one pathway.                                                             |
| <b>Step 7</b> | <code>switch(config-if-vrrs-pw)# show vrrs pathway interface-type interface-number</code> | (Optional)<br><br>Displays the VRRS pathway information for different pathway states, such as active, inactive, and not ready.                                                             |
| <b>Step 8</b> | <code>switch(config-if-vrrs-pw)# copy running-config startup-config</code>                | (Optional)<br><br>Saves this configuration change.                                                                                                                                         |

### Example

The following example shows how to configure VRRS pathways:

```
switch# configure terminal
switch(config)# interface ethernet 1/2
switch(config-if)# ip address 209.165.200.230 255.255.255.224
switch(config-if)# vrrs pathway path1
switch(config-if-vrrs-pw)# mac address fe24.fe24.fe24
switch(config-if-vrrs-pw)# address 209.165.201.10
switch(config-if-vrrs-pw)# show vrrs pathway ethernet 1/2
switch(config-if-vrrs-pw)# show running-config startup-config
```

## Verifying the VRRP Configuration

Use one of the following commands to verify the configuration:

| Command                                     | Purpose                                    |
|---------------------------------------------|--------------------------------------------|
| <code>show vrrp</code>                      | Displays the VRRP status for all groups.   |
| <code>show vrrp vr group-number</code>      | Displays the VRRP status for a VRRP group. |
| <code>show vrrs client [client-name]</code> | Displays the VRRS client information.      |

| Command                                                                      | Purpose                                                                                                      |
|------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------|
| <b>show vrrs pathway</b> [ <i>interface type/number</i> ]                    | Displays the VRRS pathway information for different pathway states, such as active, inactive, and not ready. |
| <b>show vrrs server</b>                                                      | Displays the VRRS server information.                                                                        |
| <b>show vrrs tag</b> [ <i>tag-name</i> ]                                     | Displays the VRRS tag information.                                                                           |
| <b>show fhrp</b> [ <i>interface-type interface-name</i> ] [ <b>verbose</b> ] | Displays First Hop Redundancy Protocol (FHRP) information.                                                   |
| <b>show interface</b> interface-type                                         | Displays the virtual router configuration for an interface.                                                  |

## Monitoring VRRP Statistics

Use one of the following commands to display statistics about the feature:

| Command                     | Purpose                       |
|-----------------------------|-------------------------------|
| <b>show vrrp statistics</b> | Displays the VRRP statistics. |

- Use the **clear vrrp vr** command to clear the IPv4 VRRP statistics for the specified interface.
- Use the **clear vrrp statistics** command to clear all the VRRP statistics for all interfaces in the device.
- Use the **clear vrrp ipv4** command to clear all the statistics for the specified IPv4 virtual router.

## Configuration Example for VRRP

In this example, Router A and Router B each belong to three VRRP groups. In the configuration, each group has the following properties:

- Group 1:
  - Virtual IP address is 10.1.0.10.
  - Router A will become the master for this group with priority 120.
  - Advertising interval is 3 seconds.
  - Preemption is enabled.
- Group 5:
  - Router B will become the master for this group with priority 200.
  - Advertising interval is 30 seconds.
  - Preemption is enabled.

- Group 100:
  - Router A will become the master for this group first because it has a higher IP address (10.1.0.2).
  - Advertising interval is the default 1 second.
  - Preemption is disabled.

#### Router A

```
switch (config)# interface ethernet 1/0

switch (config-if)# ip address 10.1.0.2/16
switch (config-if)# no shutdown
switch (config-if)# vrrp 1
switch (config-if-vrrp)# priority 120
switch (config-if-vrrp)# authentication text cisco
switch (config-if-vrrp)# advertisement-interval 3
switch (config-if-vrrp)# address 10.1.0.10
switch (config-if-vrrp)# no shutdown
switch (config-if-vrrp)# exit
switch (config-if)# vrrp 5
switch (config-if-vrrp)# priority 100
switch (config-if-vrrp)# advertisement-interval 30
switch (config-if-vrrp)# address 10.1.0.50
switch (config-if-vrrp)# no shutdown
switch (config-if-vrrp)# exit
switch (config-if)# vrrp 100
switch (config-if-vrrp)# no preempt
switch (config-if-vrrp)# address 10.1.0.100
switch (config-if-vrrp)# no shutdown
```

#### Router B

```
switch (config)# interface ethernet 1/0

switch (config-if)# ip address 10.2.0.1/2
switch (config-if)# no shutdown
switch (config-if)# vrrp 1
switch (config-if-vrrp)# priority 100
switch (config-if-vrrp)# authentication text cisco
switch (config-if-vrrp)# advertisement-interval 3
switch (config-if-vrrp)# address 10.2.0.10
switch (config-if-vrrp)# no shutdown
switch (config-if-vrrp)# exit
switch (config-if)# vrrp 5
switch (config-if-vrrp)# priority 200
switch (config-if-vrrp)# advertisement-interval 30
switch (config-if-vrrp)# address 10.2.0.50
switch (config-if-vrrp)# no shutdown
switch (config-if-vrrp)# exit
switch (config-if)# vrrp 100
switch (config-if-vrrp)# no preempt
switch (config-if-vrrp)# address 10.2.0.100
switch (config-if-vrrp)# no shutdown
```

This example shows how to enable VRRPv3 and create and customize a VRRPv3 group:

```
switch# configure terminal
switch(config)# feature vrrp
switch(config)# interface ethernet 4/6
switch (config-if)# vrrpv3 5 address-family ipv4
switch (config-if-vrrp3-group)# address 209.165.200.225 primary
```

```
switch (config-if-vrrp3-group)# description group3
switch (config-if-vrrp3-group)# match-address
switch (config-if-vrrp3-group)# preempt delay minimum 30
```

## Related Documents for VRRP

| Related Topic                                   | Document Title                                                              |
|-------------------------------------------------|-----------------------------------------------------------------------------|
| Configuring the gateway load balancing protocol | <a href="#">#unique_778</a>                                                 |
| Configuring the hot standby routing protocol    | <a href="#">#unique_779</a>                                                 |
| VRRP CLI commands                               | <i>Cisco Nexus 7000 Series Unicast Routing Command Reference</i>            |
| Configuring high availability                   | <i>Cisco Nexus 7000 Series NX-OS High Availability and Redundancy Guide</i> |

## Feature History for VRRP

This table includes only the updates for those releases that have resulted in additions or changes to the feature.

| Feature Name             | Releases | Feature Information                                       |
|--------------------------|----------|-----------------------------------------------------------|
| VRRPv3 and VRRS          | 6.2(2)   | These features were introduced.                           |
| BFD for VRRP             | 5.2(1)   | Added support for BFD.                                    |
| VRRP priority thresholds | 4.2(1)   | Added support for priority thresholds and vPC.            |
| VRRP object tracking     | 4.2(1)   | Added support for tracking multiple object types in VRRP. |
| VRRP                     | 4.0(1)   | This feature was introduced.                              |



## CHAPTER 23

# Configuring Object Tracking

This chapter contains the following sections:

- [Finding Feature Information, on page 507](#)
- [Information About Object Tracking, on page 507](#)
- [Prerequisites for Object Tracking, on page 509](#)
- [Guidelines and Limitations for Object Tracking, on page 509](#)
- [Default Settings for Object Tracking Parameters, on page 509](#)
- [Configuring Object Tracking, on page 510](#)
- [Verifying the Object Tracking Configuration, on page 518](#)
- [Configuration Example for Object Tracking, on page 519](#)
- [Related Documents for Object Tracking, on page 519](#)
- [Standards for Object Tracking, on page 519](#)
- [Feature History for Object Tracking, on page 519](#)

## Finding Feature Information

Your software release might not support all the features documented in this module. For the latest caveats and feature information, see the Bug Search Tool at <https://tools.cisco.com/bugsearch/> and the release notes for your software release. To find information about the features documented in this module, and to see a list of the releases in which each feature is supported, see the "New and Changed Information" chapter or the Feature History table in this chapter.

## Information About Object Tracking

Object tracking allows you to track specific objects on the device, such as the interface line protocol state, IP routing, and route reachability, and to take action when the state of the tracked object changes. This feature allows you to increase the availability of the network and shorten recovery time if an object state goes down.

The object tracking feature allows you to create a tracked object that multiple clients can use to modify the client behavior when a tracked object changes. Several clients register their interest with the tracking process, track the same object, and take different actions when the object state changes.

Clients include the following features:

- Embedded Event Manager (EEM)

- Gateway Load Balancing Protocol (GLBP)
- Hot Standby Redundancy Protocol (HSRP)
- Virtual port channel (vPC)
- Virtual Router Redundancy Protocol (VRRP)

The object tracking monitors the status of the tracked objects and communicates any changes made to interested clients. Each tracked object is identified by a unique number that clients can use to configure the action to take when a tracked object changes state.

Cisco NX-OS tracks the following object types:

- Interface line protocol state—Tracks whether the line protocol state is up or down.
- Interface IP routing state—Tracks whether the interface has an IPv4 or IPv6 address and if IPv4 or IPv6 routing is enabled and active.
- IP route reachability—Tracks whether an IPv4 or IPv6 route exists and is reachable from the local device.

For example, you can configure HSRP to track the line protocol of the interface that connects one of the redundant routers to the rest of the network. If that link protocol goes down, you can modify the priority of the affected HSRP router and cause a switchover to a backup router that has better network connectivity.

For more information related to object tracking, see the following chapters:

- [#unique\\_778](#)
- [#unique\\_779](#)
- [Configuring Layer 3 Virtualization, on page 371](#)

## Object Track List

An object track list allows you to track the combined states of multiple objects. Object track lists support the following capabilities:

- Boolean "and" function—Each object defined within the track list must be in an up state so that the track list object can become up.
- Boolean "or" function—At least one object defined within the track list must be in an up state so that the tracked object can become up.
- Threshold percentage—The percentage of up objects in the tracked list must be greater than the configured up threshold for the tracked list to be in the up state. If the percentage of down objects in the tracked list is above the configured track list down threshold, the tracked list is marked as down.
- Threshold weight—Assign a weight value to each object in the tracked list, and a weight threshold for the track list. If the combined weights of all up objects exceeds the track list weight up threshold, the track list is in an up state. If the combined weights of all the down objects exceeds the track list weight down threshold, the track list is in the down state.

Other entities, such as virtual Port Channels (vPCs) can use an object track list to modify the state of a vPC based on the state of the multiple peer links that create the vPC.

See the *Cisco Nexus 7000 Series NX-OS Interfaces Configuration Guide*, for more information on vPCs.

## High Availability

Object tracking supports high availability through stateful restarts. A stateful restart occurs when the object tracking process crashes. Object tracking also supports a stateful switchover on a dual supervisor system. Cisco NX-OS applies the runtime configuration after the switchover.

You can also use object tracking to modify the behavior of a client to improve overall network availability.

## Virtualization Support

Object tracking supports Virtual Routing and Forwarding (VRF) instances. VRFs exist within virtual device contexts (VDCs). By default, Cisco NX-OS places you in the default VDC and default VRF unless you specifically configure another VDC and VRF. By default, Cisco NX-OS tracks the route reachability state of objects in the default VRF. If you want to track objects in another VRF, you must configure the object to be a member of that VRF.

For more information, see the *Cisco Nexus 7000 Series NX-OS Virtual Device Context Configuration Guide*, and [Configuring Layer 3 Virtualization, on page 371](#).

## Prerequisites for Object Tracking



**Note** For a full list of feature-specific prerequisites, see the platform-specific documentation.

## Guidelines and Limitations for Object Tracking

Object Tracking has the following configuration guidelines and limitations:

- Supports up to 500 tracked objects per VDC.
- Supports Ethernet, subinterfaces, tunnels, port channels, loopback interfaces, and VLAN interfaces.
- Supports one tracked object per HSRP group or GLBP group.
- If you are familiar with the Cisco IOS CLI, be aware that the Cisco NX-OS commands for this feature might differ from the Cisco IOS commands that you would use.

## Default Settings for Object Tracking Parameters

### Default Object Tracking Parameters

| Parameters         | Default               |
|--------------------|-----------------------|
| Tracked Object VRF | Member of default VRF |

# Configuring Object Tracking

## Configuring Object Tracking for an Interface

You can configure Cisco NX-OS to track the line protocol or IPv4 or IPv6 routing state of an interface.

### Before you begin

Confirm that you are in the correct VDC. To change the VDC, use the **switchto vdc** command.

### Procedure

|               | Command or Action                                                                                                                      | Purpose                                                                                                                       |
|---------------|----------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configure terminal</b>                                                                                                      | Enters global configuration mode.                                                                                             |
| <b>Step 2</b> | switch(config)# <b>track</b> <i>object-id</i> <b>interface</b> <i>interface-type number</i> <b>{ip   ipv6} routing   line-protocol</b> | Creates a tracked object for an interface and enters tracking configuration mode. The object-id range is from 1 to 500.       |
| <b>Step 3</b> | (Optional) switch(config-track)# <b>show track</b> [ <i>object-id</i> ]                                                                | Displays object tracking information.                                                                                         |
| <b>Step 4</b> | (Optional) switch(config)# <b>copy running-config startup-config</b>                                                                   | Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration. |

### Example

The following example shows how to configure object tracking for the line protocol state on Ethernet 1/2:

```
switch # configure terminal
switch(config)# track 1 interface ethernet 1/2 line-protocol
switch(config)# copy running-config startup-config
```

The following example shows how to configure object tracking for the IPv4 routing state on Ethernet 1/2:

```
switch # configure terminal
switch(config)# track 2 interface ethernet 1/2 ip routing
switch(config)# copy running-config startup-config
```

The following example shows how to configure object tracking for the IPv6 routing state on Ethernet 1/2:

```
switch # configure terminal
switch(config)# track 3 interface ethernet 1/2 ipv6 routing
switch(config)# copy running-config startup-config
```



## Deleting a Tracking Object

### Before you begin

Confirm that you are in the correct VDC. To change the VDC, use the **switchto vdc** command.

### Procedure

|               | Command or Action                                                    | Purpose                                                                                                                       |
|---------------|----------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configure terminal</b>                                    | Enters global configuration mode.                                                                                             |
| <b>Step 2</b> | switch(config)# <b>no track /</b>                                    | Deletes a tracked object for an interface. The object-id range is from 1 to 500.                                              |
| <b>Step 3</b> | (Optional) switch(config)# <b>copy running-config startup-config</b> | Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration. |

### Example

The following example shows how to delete an object tracking:

```
switch # configure terminal
switch(config)# no track 1
switch(config)# copy running-config startup-config
```

## Configuring Object Tracking for Route Reachability

You can configure Cisco NX-OS to track the existence and reachability of an IP route or IPv6 route.

### Before you begin

Confirm that you are in the correct VDC. To change the VDC, use the **switchto vdc** command.

### Procedure

|               | Command or Action                                                                   | Purpose                                                                                                                                                                                                                                                                                                                                                                                  |
|---------------|-------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configure terminal</b>                                                   | Enters global configuration mode.                                                                                                                                                                                                                                                                                                                                                        |
| <b>Step 2</b> | switch(config)# <b>track object-id {ip   ipv6} route prefix/length reachability</b> | Creates a tracked object for a route and enters tracking configuration mode. The object-id range is from 1 to 500. The prefix format for IP is A.B.C.D/length, where the length range is from 1 to 32. The prefix format for IPv6 is A::B::C:D/length, where the length range is from 1 to 128. The prefix format for IPv6 is A::B::C:D/length, where the length range is from 1 to 128. |

|               | Command or Action                                               | Purpose                                                                                                                       |
|---------------|-----------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 3</b> | (Optional) switch(config-track)# <b>show track</b> [object-id]  | Displays object tracking information.                                                                                         |
| <b>Step 4</b> | switch(config-track)# <b>copy running-config startup-config</b> | Saves the change persistently through reboots and restarts by copying and running configuration to the startup configuration. |

### Example

The following example shows how to configure object tracking for an IPv4 route in the default VRF:

```
switch # configure terminal
switch(config)# track 4 ip route 192.0.2.0/8 reachability
switch(config-track)# copy running-config startup-config
```

The following example shows how to configure object tracking for an IPv6 route in the default VRF:

```
switch # configure terminal
switch(config)# track 5 ipv6 route 10::10/128 reachability
switch(config-track)# copy running-config startup-config
```

## Configuring an Object Track List with a Boolean Expression

You can configure an object track list that contains multiple tracked objects. A tracked list contains one or more objects. The Boolean expression enables two types of calculation by using either "and" or "or" operators. For example, when tracking two interfaces using the "and" operator, up means that both interfaces are up, and down means that either interface is down.

### Before you begin

Confirm that you are in the correct VDC. To change the VDC, use the **switchto vdc** command.

### Procedure

|               | Command or Action                                                 | Purpose                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|---------------|-------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configure terminal</b>                                 | Enters global configuration mode.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Step 2</b> | switch(config)# <b>track track-number list boolean {and   or}</b> | Configures a tracked list object and enters tracking configuration mode. Specifies that the state of the tracked list is based on a Boolean calculation. The keywords are as follows: <ul style="list-style-type: none"> <li>• <b>and</b>—Specifies that the list is up if all objects are up or down if one or more objects are down. For example, when tracking two interfaces, up means that both interfaces are up, and down means that either interface is down.</li> <li>• <b>or</b>—Specifies that the list is up if at least one object is up. For example, when</li> </ul> |

|               | Command or Action                                                          | Purpose                                                                                                                                                                                                                                                 |
|---------------|----------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|               |                                                                            | tracking two interfaces, up means that either interface is up, and down means that both interfaces are down.<br><br>The track-number range is from 1 to 500.                                                                                            |
| <b>Step 3</b> | switch(config-track)# <b>object</b> <i>object-number</i><br>[not]          | Adds a tracked object to the track list. The object-id range is from 1 to 500. The not keyword optionally negates the tracked object state.<br><br><b>Note</b> The example means that when object 10 is up, the tracked list detects object 10 as down. |
| <b>Step 4</b> | (Optional) switch(config-track)# <b>copy running-config startup-config</b> | Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.                                                                                                                           |
| <b>Step 5</b> | (Optional) switch(config-track)# <b>show track</b>                         | Displays object tracking information.                                                                                                                                                                                                                   |

### Example

The following example shows how to configure a track list with multiple objects as a Boolean “and”:

```
switch # configure terminal
switch(config)# track 1 list boolean and
switch(config-track)# object 10
switch(config-track)# object 20 not
switch(config)# copy running-config startup-config
```

## Configuring an Object Track List with a Percentage Threshold

You can configure an object track list that contains a percentage threshold. A tracked list contains one or more objects. The percentage of up objects must exceed the configured track list up percent threshold before the track list is in an up state. For example, if the tracked list has three objects, and you configure an up threshold of 60 percent, two of the objects must be in the up state (66 percent of all objects) for the track list to be in the up state.

### Before you begin

Confirm that you are in the correct VDC. To change the VDC, use the **switchto vdc** command.

### Procedure

|               | Command or Action                 | Purpose                           |
|---------------|-----------------------------------|-----------------------------------|
| <b>Step 1</b> | switch# <b>configure terminal</b> | Enters global configuration mode. |

|               | Command or Action                                                                                            | Purpose                                                                                                                                                                                                           |
|---------------|--------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 2</b> | switch(config)# <b>track</b> <i>track-number</i> <b>list</b><br><b>threshold percentage</b>                  | Configures a tracked list object and enters tracking configuration mode. Specifies that the state of the tracked list is based on a configured threshold percent.<br><br>The track-number range is from 1 to 500. |
| <b>Step 3</b> | switch(config-track)# <b>threshold percentage</b><br><b>up</b> <i>up-value</i> <b>down</b> <i>down-value</i> | Configures the threshold percent for the tracked list. The range from 0 to 100 percent.                                                                                                                           |
| <b>Step 4</b> | switch(config-track)# <b>object</b> <i>object-number</i>                                                     | Adds a tracked object to the track list. The object-id range is from 1 to 500.<br><br><b>Note</b> The example means that when object 10 is up, the tracked list detects object 10 as down.                        |
| <b>Step 5</b> | (Optional) switch(config-track)# <b>copy</b><br><b>running-config startup-config</b>                         | Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.                                                                                     |
| <b>Step 6</b> | (Optional) switch(config-track)# <b>show track</b>                                                           | Displays object tracking information.                                                                                                                                                                             |

### Example

The following example shows how to configure a track list with multiple objects as a Boolean “and”:

```
switch # configure terminal
switch(config)# track 1 list threshold percentage
switch(config-track)# threshold percentage up 70 down 30
switch(config-track)# object 10
switch(config-track)# object 20
switch(config-track)# object 30
switch(config-track)# copy running-config startup-config
```

## Configuring an Object Track List with a Weight Threshold

You can configure an object track list that contains a weight threshold. A tracked list contains one or more objects. The combined weight of up objects must exceed the configured track list up weight threshold before the track list is in an up state. For example, if the tracked list has three objects with the default weight of 10 each, and you configure an up threshold of 15, two of the objects must be in the up state (combined weight of 20) for the track list to be in the up state.

### Before you begin

Confirm that you are in the correct VDC. To change the VDC, use the **switchto vdc** command.

**Procedure**

|               | <b>Command or Action</b>                                                  | <b>Purpose</b>                                                                                                                                                                                                          |
|---------------|---------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configure terminal</b>                                         | Enters global configuration mode.                                                                                                                                                                                       |
| <b>Step 2</b> | switch(config)# <b>track track-number list threshold weight</b>           | Configures a tracked list object and enters tracking configuration mode. Specifies that the state of the tracked list is based on a configured threshold weight.<br><br>The <i>track-number</i> range is from 1 to 500. |
| <b>Step 3</b> | switch(config-track)# <b>threshold weight up up-value down down-value</b> | Configures the threshold weight for the tracked list. The range is from 1 to 255.                                                                                                                                       |
| <b>Step 4</b> | switch(config-track)# <b>object object-id weight value</b>                | Adds a tracked object to the track list. The <i>object-id</i> range is from 1 to 500. The <i>value</i> range is from 1 to 255. The default weight value is 10                                                           |
| <b>Step 5</b> | switch(config-track)# <b>copy running-config startup-config</b>           | Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.                                                                                           |
| <b>Step 6</b> | (Optional) switch(config-track)# <b>show track</b>                        | Displays object tracking information.                                                                                                                                                                                   |

**Example**

The following example shows how to configure a track list with a up weight threshold of 30 and a down threshold of 10:



**Note** In this example, the track list is up if object 10 and object 20 are up, and the track list goes to the down state if all three objects are down.

```
switch # configure terminal
switch(config)# track 1 list threshold weight
switch(config-track)# threshold weight up 30 down 10
switch(config-track)# object 10 weight 15
switch(config-track)# object 20 weight 15
switch(config-track)# object 30
switch(config-track)# copy running-config startup-config
```

## Configuring an Object Tracking Delay

You can configure a delay for a tracked object or an object track list that delays when the object or list triggers a stage change. The tracked object or track list starts the delay timer when a state change occurs but does not recognize a state change until the delay timer expires. At that point, Cisco NX-OS checks the object state again and records a state change only if the object or list currently has a changed state. Object tracking ignores any intermediate state changes before the delay timer expires.

For example, for an interface line-protocol tracked object that is in the up state with a 20 second down delay, the delay timer starts when the line protocol goes down. The object is not in the down state unless the line protocol is down 20 seconds later.

You can configure independent up delay and down delay for a tracked object or track list. When you delete the delay, object tracking deletes both the up and down delay.

You can change the delay at any point. If the object or list is already counting down the delay timer from a triggered event, the new delay is computed as the following:

- If the new configuration value is less than the old configuration value, the timer starts with the new value.
- If the new configuration value is more than the old configuration value, the timer is calculated as the new configuration value minus the current timer countdown minus the old configuration value.

### Before you begin

Confirm that you are in the correct VDC. To change the VDC, use the **switchto vdc** command.

### Procedure

|               | Command or Action                                                                                                                                           | Purpose                                                                                                                                                                                                                                                                                        |
|---------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configure terminal</b>                                                                                                                           | Enters global configuration mode.                                                                                                                                                                                                                                                              |
| <b>Step 2</b> | switch(config)# <b>track</b> <i>object-id</i> { <i>parameters</i> }                                                                                         | Creates a tracked object for a route and enters tracking configuration mode. The object-id range is from 1 to 500. The prefix format for IP is A.B.C.D/length, where the length range is from 1 to 32. The prefix format for IPv6 is A:B::C:D/length, where the length range is from 1 to 128. |
| <b>Step 3</b> | switch(config-track)# <b>track</b> <i>track-number</i> <b>list</b> { <i>parameters</i> }                                                                    | Configures a tracked list object and enters tracking configuration mode. Specifies that the state of the tracked list is based on a configured threshold weight.<br><br>The <i>track-number</i> range is from 1 to 500.                                                                        |
| <b>Step 4</b> | switch(config-track)# <b>delay</b> { <b>up</b> <i>up-time</i> [ <b>down</b> <i>down-time</i> ]   <b>down</b> <i>down-time</i> [ <b>up</b> <i>up-time</i> ]} | Configures the object delay timers. The range is from 0 to 180 seconds.<br><br>The <i>track-number</i> range is from 1 to 500.                                                                                                                                                                 |
| <b>Step 5</b> | (Optional) switch(config-track)# <b>copy running-config startup-config</b>                                                                                  | Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.                                                                                                                                                                  |
| <b>Step 6</b> | (Optional) switch(config-track)# <b>show track</b>                                                                                                          | Displays object tracking information.                                                                                                                                                                                                                                                          |

### Example

The following example shows how to configure object tracking for a route and use delay timers:

```
switch # configure terminal
switch(config)# track 2 ip route 209.165.201.0/8 reachability
switch(config-track)# delay up 20 down 30
switch(config-track)# copy running-config startup-config
```

The following example shows how to configure a track list with an up weight threshold of 30 and a down threshold of 10 with delay timers:

```
switch# configure terminal
switch(config)# track 1 list threshold weight
switch(config-track)# threshold weight up 30 down 10
switch(config-track)# object 10 weight 15
switch(config-track)# object 20 weight 15
switch(config-track)# object 30
switch(config-track)# delay up 20 down 30
switch(config-track)# copy running-config startup-config
```

The following example shows the delay timer in the **show track** command output before and after an interface is shut down:

```
switch(config-track)# show track
Track 1
 Interface loopback1 Line Protocol
 Line Protocol is UP
 1 changes, last change 00:00:13
 Delay down 10 secs

switch(config-track)# interface loopback 1
switch(config-if)# shutdown
switch(config-if)# show track
Track 1
 Interface loopback1 Line Protocol
 Line Protocol is delayed DOWN (8 secs remaining)<----- delay timer counting down
 1 changes, last change 00:00:22
 Delay down 10 secs
```

## Configuring Object Tracking for a Nondefault VRF

You can configure Cisco NX-OS to track an object in a specific VRF.

### Before you begin

- Confirm that you are in the correct VDC. To change the VDC, use the **switchto vdc** command.
- Ensure that nondefault VRFs are created first.

### Procedure

|               | Command or Action                                                                   | Purpose                                                                                                                                                                                                                                     |
|---------------|-------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configure terminal</b>                                                   | Enters global configuration mode.                                                                                                                                                                                                           |
| <b>Step 2</b> | switch(config)# <b>track object-id {ip   ipv6} route prefix/length reachability</b> | Creates a tracked object for a route and enters tracking configuration mode. The <i>object-id</i> range is from 1 to 500. The prefix format for IP is A.B.C.D/length, where the length range is from 1 to 32. The prefix format for IPv6 is |

|               | Command or Action                                               | Purpose                                                                                                                       |
|---------------|-----------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------|
|               |                                                                 | A::B::C:D/length, where the length range is from 1 to 128.                                                                    |
| <b>Step 3</b> | switch(config-track)# <b>vrf member</b> <i>vrf-name</i>         | Configures the VRF to use for tracking the configured object.                                                                 |
| <b>Step 4</b> | switch(config-track)# <b>copy running-config startup-config</b> | Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration. |
| <b>Step 5</b> | (Optional) switch(config-track)# <b>show track</b>              | Displays object tracking information.                                                                                         |

### Example

The following example shows how to configure object tracking for a route and use VRF Red to look up reachability information for this object:

```
switch # configure terminal
switch(config)# track 2 ip route 209.165.201.0/8 reachability
switch(config-track)# vrf member Red
switch(config-track)# copy running-config startup-config
```

The following example shows how to configure object tracking for an IPv6 route and use VRF Red to look up reachability information for this object:

```
switch# configure terminal
switch(config)# track 3 ipv6 route 1::2/64 reachability
switch(config-track)# vrf member Red
switch(config-track)# copy running-config startup-config
```

The following example how to modify tracked object 2 to use VRF Blue instead of VRF Red to look up reachability information for this object:

```
switch(config-track)# show track
switch(config)# track 2
switch(config-track)# vrf member Blue
switch(config-track)# copy running-config startup-config
```

## Verifying the Object Tracking Configuration

Use one of the following commands to verify the configuration:

| Command                                                         | Purpose                                                           |
|-----------------------------------------------------------------|-------------------------------------------------------------------|
| <b>show track</b> [ <i>object-id</i> ] [brief]                  | Displays the object tracking information for one or more objects. |
| <b>show track</b> [ <i>object-id</i> ] <b>interface</b> [brief] | Displays the interface-based object tracking information.         |



| Command                                                       | Purpose                                                               |
|---------------------------------------------------------------|-----------------------------------------------------------------------|
| <code>show track [object-id] {ip   ipv6} route [brief]</code> | Displays the IPv4 or IPv6 route-based object tracking information.    |
| <code>show trun track</code>                                  | Displays the IP route IPv6 object tracking configuration information. |

## Configuration Example for Object Tracking

This example shows how to configure object tracking for route reachability and use VRF Red to look up reachability information for this route:

```
switch# configure terminal
switch(config)# track 2 ip route 209.165.201.0/8 reachability
switch(config-track)# vrf member Red
```

## Related Documents for Object Tracking

| Related Topic                          | Document Title                                                             |
|----------------------------------------|----------------------------------------------------------------------------|
| Object Tracking CLI commands           | <i>Cisco Nexus 7000 Series NX-OS Unicast Routing Command Reference</i>     |
| Configuring the Embedded Event Manager | <i>Cisco Nexus 7000 Series NX-OS System Management Configuration Guide</i> |

## Standards for Object Tracking

No new or modified standards are supported by this feature, and support for existing standards has not been modified by this feature.

## Feature History for Object Tracking

This table includes only the updates for those releases that have resulted in additions or changes to the feature.

| Feature Name      | Releases | Feature Information                                           |
|-------------------|----------|---------------------------------------------------------------|
| IPv6              | 5.0(2)   | Added support for IPv6.                                       |
| Tracking delay    | 4.2(4)   | Added support for delaying a tracked object update.           |
| Object track list | 4.2(1)   | Added support for object track lists and Boolean expressions. |
| Object tracking   | 4.0(1)   | This feature was introduced.                                  |





## APPENDIX **A**

# IETF RFCs Supported by Cisco NX-OS Unicast Features Release 6.x

- [BGP RFCs, on page 521](#)
- [First-Hop Redundancy Protocols RFCs, on page 522](#)
- [IP Services RFCs, on page 522](#)
- [IPv6 RFCs, on page 523](#)
- [IS-IS RFCs, on page 523](#)
- [OSPF RFCs, on page 524](#)
- [RIP RFCs, on page 524](#)

## BGP RFCs

| RFCs     | Title                                                                      |
|----------|----------------------------------------------------------------------------|
| RFC 1997 | <i>BGP Communities Attribute</i>                                           |
| RFC 1998 | <i>An Application of the BGP Community Attribute in Multi-home Routing</i> |
| RFC 2385 | <i>Protection of BGP Sessions via the TCP MD5 Signature Option</i>         |
| RFC 2439 | <i>BGP Route Flap Damping</i>                                              |
| RFC 2519 | <i>A Framework for Inter-Domain Route Aggregation</i>                      |
| RFC 2545 | <i>Use of BGP-4 Multiprotocol Extensions for IPv6 Inter-Domain Routing</i> |
| RFC 2858 | <i>Multiprotocol Extensions for BGP-4</i>                                  |
| RFC 3065 | <i>Autonomous System Confederations for BGP</i>                            |
| RFC 3107 | <i>Carrying Label Information in BGP-4</i>                                 |
| RFC 3392 | <i>Capabilities Advertisement with BGP-4</i>                               |
| RFC 4271 | <i>A Border Gateway Protocol 4 (BGP-4)</i>                                 |
| RFC 4273 | <i>Definitions of Managed Objects for BGP-4</i>                            |

| RFCs                                  | Title                                                                                |
|---------------------------------------|--------------------------------------------------------------------------------------|
| RFC 4456                              | <i>BGP Route Reflection: An Alternative to Full Mesh Internal BGP (IBGP)</i>         |
| RFC 4486                              | <i>Subcodes for BGP Cease Notification Message</i>                                   |
| RFC 4724                              | <i>Graceful Restart Mechanism for BGP</i>                                            |
| RFC 4893                              | <i>BGP Support for Four-octet AS Number Space</i>                                    |
| RFC 5004                              | <i>Avoid BGP Best Path Transitions from One External to Another</i>                  |
| RFC 5396 <sup>3</sup>                 | <i>Textual Representation of Autonomous System (AS) Numbers</i>                      |
| RFC 5549                              | <i>Advertising IPv4 Network Layer Reachability Information with an IPv6 Next Hop</i> |
| RFC 5668                              | <i>4-Octet AS Specific BGP Extended Community</i>                                    |
| draft-ietf-idr-add-paths-08.txt       | <i>Advertisement of Multiple Paths in BGP</i>                                        |
| draft-ietf-idr-bgp4-mib-15.txt        | <i>BGP4-MIB</i>                                                                      |
| draft-kato-bgp-ipv6-link-local-00.txt | <i>BGP4+ Peering Using IPv6 Link-local Address</i>                                   |

<sup>3</sup> RFC 5396 is partially supported. The asplain and asdot notations are supported, but the asdot+ notation is not.

## First-Hop Redundancy Protocols RFCs

| RFCs     | Title                                     |
|----------|-------------------------------------------|
| RFC 2281 | <i>Hot Standby Redundancy Protocol</i>    |
| RFC 3768 | <i>Virtual Router Redundancy Protocol</i> |

## IP Services RFCs

| RFCs    | Title       |
|---------|-------------|
| RFC 768 | <i>UDP</i>  |
| RFC 791 | <i>IP</i>   |
| RFC 792 | <i>ICMP</i> |
| RFC 793 | <i>TCP</i>  |
| RFC 826 | <i>ARP</i>  |

| RFCs     | Title                          |
|----------|--------------------------------|
| RFC 1027 | <i>Proxy ARP</i>               |
| RFC 1591 | <i>DNS Client</i>              |
| RFC 1812 | <i>IPv4 routers</i>            |
| RFC 4022 | <i>TCP-MIB</i>                 |
| RFC 4292 | <i>IP-FORWARDING-TABLE-MIB</i> |
| RFC 4293 | <i>IP-MIB</i>                  |

## IPv6 RFCs

| RFCs     | Title                                                                                               |
|----------|-----------------------------------------------------------------------------------------------------|
| RFC 1981 | Path MTU Discovery for IP version 6                                                                 |
| RFC 2373 | IP Version 6 Addressing Architecture                                                                |
| RFC 2374 | An Aggregatable Global Unicast Address Format                                                       |
| RFC 2460 | Internet Protocol, Version 6 (IPv6) Specification                                                   |
| RFC 2461 | Neighbor Discovery for IP Version 6 (IPv6)                                                          |
| RFC 2462 | IPv6 Stateless Address Autoconfiguration                                                            |
| RFC 2463 | Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification |
| RFC 2464 | Transmission of IPv6 Packets over Ethernet Networks                                                 |
| RFC 3152 | Delegation of IP6.ARPA                                                                              |
| RFC 3162 | RADIUS and IPv6                                                                                     |
| RFC 3513 | Internet Protocol Version 6 (IPv6) Addressing Architecture                                          |
| RFC 3596 | DNS Extensions to Support IP version 6                                                              |
| RFC 4193 | Unique Local IPv6 Unicast Addresses                                                                 |

## IS-IS RFCs

| RFCs     | Title                                                                                              |
|----------|----------------------------------------------------------------------------------------------------|
| RFC 1142 | <i>OSI 10589 intermediate system to intermediate system intro-domain routing exchange protocol</i> |

| RFCs                                    | Title                                                                                              |
|-----------------------------------------|----------------------------------------------------------------------------------------------------|
| RFC 1195                                | <i>Use of OSI IS-IS for routing in TCP/IP and dual environment</i>                                 |
| RFC 2763, RFC 5301                      | <i>Dynamic Hostname Exchange Mechanism for IS-IS</i>                                               |
| RFC 2966, RFC 5302                      | <i>Domain-wide Prefix Distribution with Two-Level IS-IS</i>                                        |
| RFC 2972                                | <i>IS-IS Mesh Groups</i>                                                                           |
| RFC 3277                                | <i>IS-IS Transient Blackhole Avoidance</i>                                                         |
| RFC 3373, RFC 5303                      | <i>Three-Way Handshake for IS-IS Point-to-Point Adjacencies</i>                                    |
| RFC 3567, RFC 5304                      | <i>IS-IS Cryptographic Authentication</i>                                                          |
| RFC 3847, RFC 5306                      | <i>Restart Signaling for IS-IS</i>                                                                 |
| RFC 4205, RFC 5307                      | <i>IS-IS Extensions in Support of Generalized Multi-Protocol Label Switching</i>                   |
| RFC 5120                                | <i>M-ISIS: Multi Topology (MT) Routing in Intermediate System to Intermediate Systems (IS-ISs)</i> |
| draft-ietf-isis-igp-p2p-over-lan-06.txt | <i>Internet Draft Point-to-point operation over LAN in link-state routing protocols</i>            |

## OSPF RFCs

| RFCs                                           | Title                                            |
|------------------------------------------------|--------------------------------------------------|
| RFC 2328                                       | <i>OSPF Version 2</i>                            |
| RFC 2740                                       | <i>OSPF for IPv6</i>                             |
| RFC 3623                                       | <i>Graceful OSPF Restart</i>                     |
| RFC 3101                                       | <i>The OSPF Not-So-Stubby Area (NSSA) Option</i> |
| RFC 2370                                       | <i>The OSPF Opaque LSA Option</i>                |
| RFC 3137                                       | <i>OSPF Stub Router Advertisement</i>            |
| draft-ietf-ospf-ospfv3-graceful-restart-04.txt | <i>OSPFv3 Graceful Restart</i>                   |

## RIP RFCs

| RFCs     | Title                |
|----------|----------------------|
| RFC 2453 | <i>RIP Version 2</i> |

| RFCs     | Title                           |
|----------|---------------------------------|
| RFC 2082 | <i>RIP-2 MD5 Authentication</i> |







## APPENDIX **B**

# Configuration Limits for Cisco NX-OS Layer 3 Unicast Features

---

The configuration limits are documented in the *Cisco Nexus 7000 Series NX-OS Verified Scalability Guide*.

