



Configuration — BGP Services

Avaya Virtual Services Platform 9000

3.2
NN46250-507, 03.02
April 2012

Notice

While reasonable efforts have been made to ensure that the information in this document is complete and accurate at the time of printing, Avaya assumes no liability for any errors. Avaya reserves the right to make changes and corrections to the information in this document without the obligation to notify any person or organization of such changes.

Documentation disclaimer

"Documentation" means information published by Avaya in varying mediums which may include product information, operating instructions and performance specifications that Avaya generally makes available to users of its products. Documentation does not include marketing materials. Avaya shall not be responsible for any modifications, additions, or deletions to the original published version of documentation unless such modifications, additions, or deletions were performed by Avaya. End User agrees to indemnify and hold harmless Avaya, Avaya's agents, servants and employees against all claims, lawsuits, demands and judgments arising out of, or in connection with, subsequent modifications, additions or deletions to this documentation, to the extent made by End User.

Link disclaimer

Avaya is not responsible for the contents or reliability of any linked Web sites referenced within this site or documentation provided by Avaya. Avaya is not responsible for the accuracy of any information, statement or content provided on these sites and does not necessarily endorse the products, services, or information described or offered within them. Avaya does not guarantee that these links will work all the time and has no control over the availability of the linked pages.

Warranty

Avaya provides a limited warranty on its Hardware and Software ("Product(s)"). Refer to your sales agreement to establish the terms of the limited warranty. In addition, Avaya's standard warranty language, as well as information regarding support for this Product while under warranty is available to Avaya customers and other parties through the Avaya Support Web site: <http://support.avaya.com>. Please note that if you acquired the Product(s) from an authorized Avaya reseller outside of the United States and Canada, the warranty is provided to you by said Avaya reseller and not by Avaya.

Licenses

THE SOFTWARE LICENSE TERMS AVAILABLE ON THE AVAYA WEBSITE, [HTTP://SUPPORT.AVAYA.COM/LICENSEINFO/](http://support.avaya.com/licenseinfo/) ARE APPLICABLE TO ANYONE WHO DOWNLOADS, USES AND/OR INSTALLS AVAYA SOFTWARE, PURCHASED FROM AVAYA INC., ANY AVAYA AFFILIATE, OR AN AUTHORIZED AVAYA RESELLER (AS APPLICABLE) UNDER A COMMERCIAL AGREEMENT WITH AVAYA OR AN AUTHORIZED AVAYA RESELLER. UNLESS OTHERWISE AGREED TO BY AVAYA IN WRITING, AVAYA DOES NOT EXTEND THIS LICENSE IF THE SOFTWARE WAS OBTAINED FROM ANYONE OTHER THAN AVAYA, AN AVAYA AFFILIATE OR AN AVAYA AUTHORIZED RESELLER; AVAYA RESERVES THE RIGHT TO TAKE LEGAL ACTION AGAINST YOU AND ANYONE ELSE USING OR SELLING THE SOFTWARE WITHOUT A LICENSE. BY INSTALLING, DOWNLOADING OR USING THE SOFTWARE, OR AUTHORIZING OTHERS TO DO SO, YOU, ON BEHALF OF YOURSELF AND THE ENTITY FOR WHOM YOU ARE INSTALLING, DOWNLOADING OR USING THE SOFTWARE (HEREINAFTER REFERRED TO INTERCHANGEABLY AS "YOU" AND "END USER"), AGREE TO THESE TERMS AND CONDITIONS AND CREATE A BINDING CONTRACT BETWEEN YOU AND AVAYA INC. OR THE APPLICABLE AVAYA AFFILIATE ("AVAYA").

Copyright

Except where expressly stated otherwise, no use should be made of materials on this site, the Documentation, Software, or Hardware provided by Avaya. All content on this site, the documentation and the Product provided by Avaya including the selection, arrangement and design of the content is owned either by Avaya or its licensors and is protected by copyright and other intellectual property laws including the sui generis rights relating to the protection of databases. You may not modify, copy, reproduce, republish, upload, post, transmit or distribute in any way any content, in whole or in part, including any code and software unless expressly authorized by Avaya. Unauthorized reproduction, transmission, dissemination, storage, and or use without the express written consent of Avaya can be a criminal, as well as a civil offense under the applicable law.

Third-party components

Certain software programs or portions thereof included in the Product may contain software distributed under third party agreements ("Third Party Components"), which may contain terms that expand or limit rights to use certain portions of the Product ("Third Party Terms"). Information regarding distributed Linux OS source code (for those Products that have distributed the Linux OS source code), and identifying the copyright holders of the Third Party Components and the Third Party Terms that apply to them is available on the Avaya Support Web site: <http://support.avaya.com/Copyright>.

Trademarks

The trademarks, logos and service marks ("Marks") displayed in this site, the Documentation and Product(s) provided by Avaya are the registered or unregistered Marks of Avaya, its affiliates, or other third parties. Users are not permitted to use such Marks without prior written consent from Avaya or such third party which may own the Mark. Nothing contained in this site, the Documentation and Product(s) should be construed as granting, by implication, estoppel, or otherwise, any license or right in and to the Marks without the express written permission of Avaya or the applicable third party.

Avaya is a registered trademark of Avaya Inc.

All non-Avaya trademarks are the property of their respective owners, and "Linux" is a registered trademark of Linus Torvalds.

Downloading Documentation

For the most current versions of Documentation, see the Avaya Support Web site: <http://support.avaya.com>.

Contact Avaya Support

Avaya provides a telephone number for you to use to report problems or to ask questions about your Product. The support telephone number is 1-800-242-2121 in the United States. For additional support telephone numbers, see the Avaya Web site: <http://support.avaya.com>.

Contents

Chapter 1: Purpose of this document	5
Chapter 2: New in this release	7
Features.....	7
Other changes.....	7
Chapter 3: BGP fundamentals	9
Autonomous systems.....	9
Routing information consolidation.....	15
BGP communities.....	24
BGP path attributes.....	25
BGP route selection.....	25
BGP and dampened routes.....	27
BGP updates.....	28
Equal Cost Multipath.....	31
MD5 message authentication.....	31
BGP and route redistribution.....	33
BGP+.....	34
Circuitless IP.....	35
BGP configuration considerations and limitations.....	36
Chapter 4: BGP configuration using ACLI	39
Configuring BGP globally.....	39
Configuring 4 byte AS numbers.....	44
Configuring IPv4 and IPv6 aggregate routes.....	46
Configuring allowed IPv4 and IPv6 networks.....	48
Configuring BGP peers or peer groups.....	49
Configuring a BGP peer or peer group password.....	54
Configuring redistribution to BGP.....	55
Configuring redistribution to BGP+ for VRF 0.....	57
Configuring prefix lists.....	60
Configuring an IPv6 prefix list.....	62
Configuring AS path lists.....	63
Configuring community lists.....	64
Configuring extended community lists.....	65
Chapter 5: BGP verification using ACLI	69
Viewing BGP aggregate information.....	69
Viewing IPv6 BGP+ aggregate information.....	70
Viewing CIDR routes.....	70
Viewing BGP configuration.....	72
Viewing BGP confederation.....	73
Viewing flap-dampened routes.....	73
Viewing global flap-dampening configurations.....	75
Viewing imported routes.....	76
Viewing BGP neighbors information.....	78
Viewing BGP network configurations.....	79
Viewing IPv6 BGP+ network configurations.....	80

Viewing BGP peer group information.....	81
Viewing BGP redistributed routes.....	81
Viewing a summary of BGP configurations.....	82
Viewing BGP routes.....	84
Viewing IPv6 BGP+ routes.....	86
Chapter 6: BGP configuration using EDM.....	89
Configuring BGP globally.....	89
Configuring 4-byte AS numbers.....	93
Configuring aggregate routes.....	94
Configuring aggregate IPv6 routes.....	96
Configuring allowed networks.....	97
Configuring allowed IPv6 networks.....	98
Configuring BGP peers.....	99
Configuring peer groups.....	103
Viewing route information.....	105
Viewing community attributes.....	108
Viewing IPv6 community attributes.....	109
Displaying dampened routes information.....	111
Configuring redistribution to BGP.....	112
Configuring a prefix list.....	113
Configuring IPv6 Prefix List.....	114
Viewing IPv6 BGP+ route summary information.....	115
Configuring an AS path list.....	116
Configuring a community access list.....	117
Configuring an extended community list.....	118
Chapter 7: BGP configuration examples.....	121
IPv6 tunnel configurations for BGP+.....	121
Chapter 8: Customer service.....	127
Getting technical documentation.....	127
Getting product training.....	127
Getting help from a distributor or reseller.....	127
Getting technical support from the Avaya Web site.....	127
Glossary.....	129

Chapter 1: Purpose of this document

This document describes the conceptual and procedural information to configure Border Gateway Protocol (BGP) services on Avaya Virtual Services Platform 9000.

Following operations are supported by BGP:

- IPv4
- IPv6
- 4 byte AS
- IPv6 Prefix-list
- Peer groups
- Redistribution

Purpose of this document

Chapter 2: New in this release

The following sections detail what's new in *Avaya Virtual Services Platform 9000 Configuration — BGP Services*, NN46250–507 for Release 3.2.

Features

See the following sections for feature-related changes:

4-byte AS support

In release 3.2, the Virtual Services Platform 9000 supports both 2-byte and 4-byte AS numbers. The 4-byte AS number is a limited deployment for BGP. You can configure a BGP peer to operate in the old 2-byte AS mode or in the new 4-byte mode. For more information about 4-byte AS support for BGP, see [BGP 4 byte AS support](#) on page 13.

BGP+

VSP 9000 Release 3.2 extends the BGPv4 process to support the exchange of IPv6 routes using BGPv4 peering. BGP+ is an extension of BGPv4 for IPv6. For more information on BGP+, see [BGP+](#) on page 34.

Other changes

See the following sections for information about changes that are not feature-related:

Glossary

A glossary of related terms is added to the end of this document.

Grouping of the chapters

The chapters are grouped according to the description of the feature to improve clarity. The Fundamentals chapter is placed in the beginning, followed by the configuration of the feature chapters by using ACLI and EDM interfaces.

ACLI Commands

Examples for ACLI commands exist for most commands in the document.

Introduction chapter and navigation

Introduction chapter is renamed to Purpose of the document that states the purpose of referring to that document. Navigation is removed to reduce the length of the document.

Terminology

Terminology no longer exists in a separate document. Terminology for this document is in a glossary at the end of this document.

Updates

Resulting from WI 988778, document version 03.02 was issued to include

- improvements to the BGP+ tunnel section in the BGP fundamentals chapter
- improvements to the IPv6 tunnel configurations for BGP+ content at the beginning of the BGP configuration examples chapter
- 2 minor changes to the iBGP+ peership on CLIP between two VSP 9000 switches with IPv6 tunneling example in the BGP configuration examples chapter

Chapter 3: BGP fundamentals

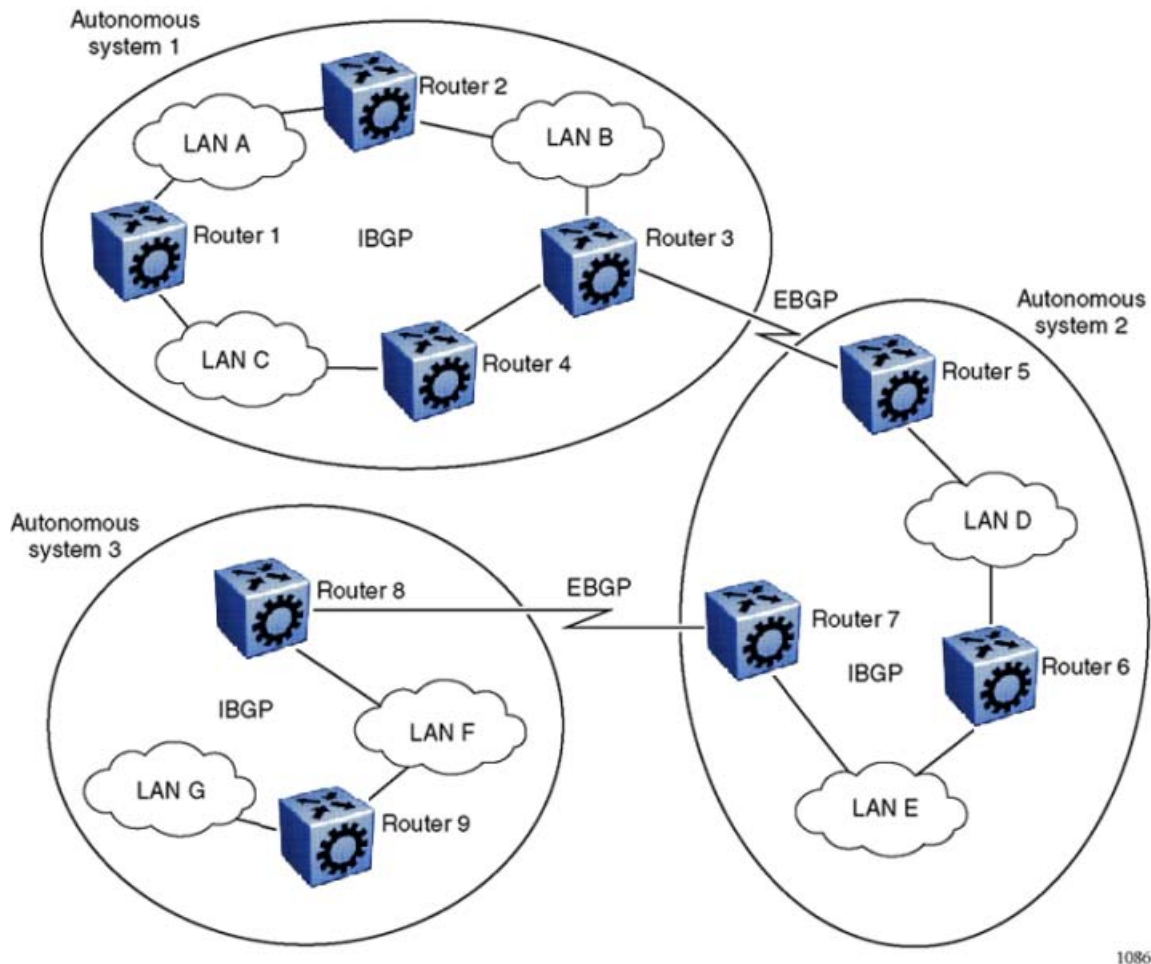
Border Gateway Protocol (BGP) is an inter-domain routing protocol that provides loop-free routing between autonomous systems (AS) or within an AS.

For information about how to use the Avaya command line interface (ACLI), and Enterprise Device Manager (EDM), see *Avaya Virtual Services Platform 9000 User Interface Fundamentals*, NN46250-103.

Autonomous systems

An AS is a group of routers and hosts run by a single technical administrator that has a single, clearly defined routing policy. Each AS uses a unique AS number assigned by the appropriate Internet Registry entity. LANs and WANs that interconnect by IP routers form a group of networks called an internetwork. For administrative purposes, internetworks divide into boundaries known as autonomous systems.

The following figure shows a sample internetwork segmented into three autonomous systems.



10860FA

Figure 1: Internetwork segmented into three autonomous systems

BGP exchanges information between ASs as well as between routers within the same AS. As shown in [Figure 1: Internetwork segmented into three autonomous systems](#) on page 10, routers that are members of the same AS and exchange BGP updates run internal BGP (IBGP), and routers that are members of different ASs and exchange BGP updates run external BGP (EBGP).

Internal and external BGP routing

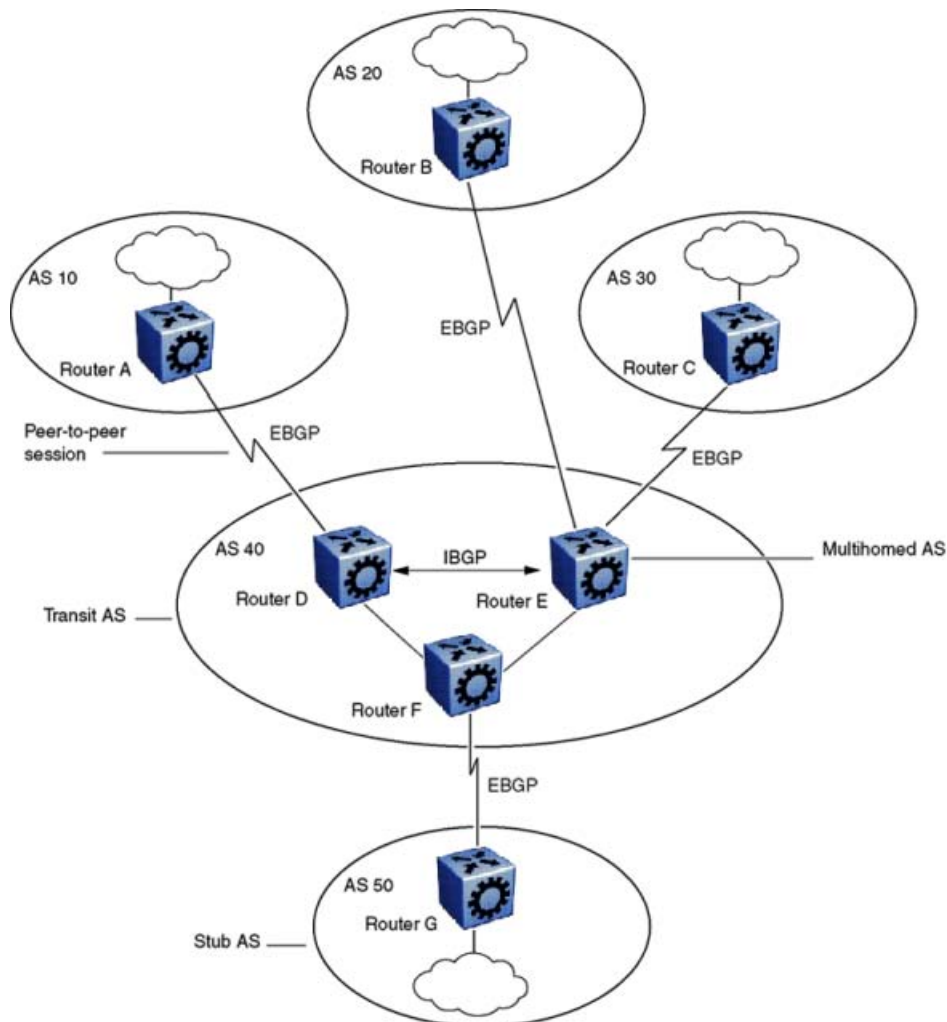
Avaya Virtual Services Platform 9000 supports both IBGP intra-AS routing and EBGP external-AS routing. With IBGP, each router within an AS runs an interior gateway protocol (IGP), such as Routing Information Protocol (RIP) or Open Shortest Path First (OSPF). The IBGP information, along with the IGP route to the originating BGP border router, determines the next hop to use to exchange information with an external AS. Each router uses IBGP exclusively to determine reachability to external ASs. After a router receives an IBGP update destined for an external AS, it passes the update to IP for inclusion in the routing table only if a viable IGP route to the correct border gateway is available.

BGP speakers in different ASs use EBGP communicate routing information.

BGP speaker

BGP routers employ an entity within the router, referred to as a BGP speaker, which transmits and receives BGP messages and acts upon them. BGP speakers establish a peer-to-peer session with other BGP speakers to communicate.

All BGP speakers within an AS must be fully meshed. The following figure shows a BGP network with fully-meshed BGP speakers.



10861FA

Figure 2: BGP networks

Transit AS

An AS with more than one BGP speaker can use IBGP to provide a transit service for networks located outside the AS. An AS that provides this service is a transit AS. As shown in [Figure 2: BGP networks](#) on page 11, AS 40 is the transit AS. AS 40 provides information about the internal networks, as well as transit networks, to the remaining ASs. The IBGP connections between routers D, E, and F provide consistent routing information to the ASs.

Stub and multihomed autonomous systems

As shown in [Figure 2: BGP networks](#) on page 11, an AS can include one or more BGP speakers that establish peer-to-peer sessions with BGP speakers in other ASs to provide external route information for the networks within the AS.

A stub AS has a single BGP speaker that establishes a peer-to-peer session with one external BGP speaker. In this case, the BGP speaker provides external route information only for the networks within its own AS.

A multihomed AS has multiple BGP speakers.

Peers

BGP uses Transmission Control Protocol (TCP) as a transport protocol. When two routers open a TCP connection to each other for the purpose of exchanging routing information, they form a peer-to-peer relationship. In [Figure 2: BGP networks](#) on page 11, Routers A and D are BGP peers, as are Routers B and E, C and E, F and G, and Routers D, E, and F.

Although Routers A and D run EBGP, Routers D, E, and F within AS 40 run IBGP. The EBGP peers directly connect to each other, while the IBGP peers do not. As long as an IGP operates and allows two neighbors to logically communicate, the IBGP peers do not require a direct connection.

Because all BGP speakers within an AS must be fully meshed logically, the IBGP mesh can grow to large proportions and become difficult to manage. You can reduce the number of peers within an AS by creating confederations and route reflectors.

BGP peers exchange complete routing information only after the peers establish a connection. Thereafter, BGP peers exchange routing updates. An update message consists of a network number, a list of autonomous systems that the routing information passed through (the AS path), and other path attributes that describe the route to a set of destination networks. When multiple paths exist, BGP compares the path attributes to choose the preferred path. Even if you disable BGP, the system logs all BGP peer connection requests. For more information about update messages, see [BGP updates](#) on page 28.

Supernet advertisements

BGP has no concept of address classes. Each network listed in the network layer reachability information (NLRI) portion of an update message contains a prefix length field, which describes the length of the mask associated with the network. The prefix length field allows for both supernet and subnet advertisement. The supernet advertisement is what makes classless interdomain routing (CIDR) possible (see [CIDR and aggregate addresses](#) on page 15).

Bandwidth and maintenance reduction

BGP provides two features that reduce the high bandwidth and maintenance costs associated with a large full-mesh topology:

- confederations
- route reflectors

The following section describes confederations and route reflectors.

BGP 4 byte AS support

Each Autonomous System (AS) must have its own unique number. Because the 2-byte AS numbering scheme is unable to meet the increasing demand, Avaya supports 4-byte AS numbers. This feature is enabled by supporting RFC4893, BGP Support for Four-octet AS Number Space.


VSP 9000 supports the following three types of peer relationships as a result of 4 byte AS support:

- Old peer to old peer
- Old peer to new peer
- New peer to new peer

An old peer is the one that supports 2-byte AS numbers only and new peer is the one that supports both 2-byte AS numbers and 4-byte AS numbers.

RFC4893 supports two new path attributes:

- AS4_PATH contains the AS path encoded with a 4-octet AS number.
- AS4-AGGR is a new aggregator attribute that carries a 4-octet AS number.



The preceding figure shows an example of how VSP 9000 uses the AS4_PATH attribute in a mixed environment. The figure illustrates how a 2-byte BGP speaker interoperates with a 4-byte BGP speaker.

Router B is a 2-byte BGP speaker. Router A substitutes AS_PATH with the AS_TRANS, a 2-octet AS number defined by RFC4893 for backward compatibility, and encodes the 4-byte AS into AS4_PATH in BGP updates it sends to router B.

Router B does not understand the AS4_PATH but does preserve the information and sends it to router C.

Router C is a 4-byte BGP speaker. Router C merges the information received in AS_PATH and AS4_PATH, and encodes the 4-byte AS when it sends the AS_PATH information to router D.

Old peer to old peer

When the peer relationship between an old peer and another old peer is established, 4 byte AS numbers contained in the AS4_PATH and AS4_AGGREGATOR are transited to other peers.

Important:

Attention: Do not assign 23456 as an AS number. The Internet Assigned Numbers Authority (IANA) reserved this number for the AS_TRANS attribute and BGP uses it to facilitate communication between peer modes. AS_TRANS uses a 2-byte AS format to represent a 4-byte AS number. The VSP 9000 switch interprets the AS_TRANS attribute and propagates it to other peers.

New peer to new peer

The new BGP speaker establishes its 4 byte AS support through BGP capability advertisement. A BGP speaker that announces such capability and receives it from its peer, uses 4 byte AS numbers in AS_PATH and AGGREGATOR attributes and assumes these attributes received from its peer are encoded in 4 byte AS numbers.

The new BGP attributes AS4_PATH and AS4_AGGREGATOR received from the new BGP speaker between the new BGP peers in the update message is discarded.

Old peer to new peer

An old BGP speaker and a new BGP speaker can form peering relationship only if the new BGP speaker is assigned a 2 byte AS number. This 2 byte number can be any global unique AS number or AS_TRANS.

New BGP speaker sends AS path information to the old BGP speaker in AS_PATH attribute as well as AS4_PATH attribute. If the entire AS_PATH consists of only 2 byte AS numbers then the new BGP speaker does not send AS4_PATH information.

The 4-byte AS number feature does not in any way restrict the use or change the way you configure 2-byte AS numbers. You can also configure 2-byte AS or 4 byte AS numbers in AS path lists, community lists, and route policies.

BGP 4-byte AS Number notation

BGP 4-byte AS numbers are represented in two ways: AS Plain and AS dot. The default form of representing the AS numbers is AS Plain while you have an option to configure AS dot. AS Plain form of representation is preferred over AS dot representation as a large amount of network providers find the AS dot notation incompatible with the regular expressions used by

them. In case of any issues, troubleshooting and analyzing also gets difficult with AS dot notation.

BGP AS Number Format – AS Plain

Table 1: Default Asplain 4-Byte Autonomous System Number Format

Format	Configuration Format	Show Command Output and Regular Expression Match Format
asplain	2-byte: 1 to 65535 4-byte: 65536 to 4294967295	2-byte: 1 to 65535 4-byte: 65536 to 4294967295
asdot	2-byte: 1 to 65535 4-byte: 1.0 to 65535.65535	2-byte: 1 to 65535 4-byte: 65536 to 4294967295

BGP AS Number Format - ASdot

Table 2: Asdot 4-Byte Autonomous System Number Format

Format	Configuration Format	Show command output and Regular Expression Match Format
asplain	2-byte: 1 to 65535 4-byte: 65536 to 4294967295	2-byte: 1 to 65535 4-byte: 1.0 to 65535.65535
asdot	2-byte: 1 to 65535 4-byte: 1.0 to 65535.65535	2-byte: 1 to 65535 4-byte: 1.0 to 65535.65535

For more information on configuring 4 byte AS numbers, see [Configuring 4 byte AS numbers](#) on page 44.

Routing information consolidation

Use the information in this section to understand how to reduce the size of routing tables.

CIDR and aggregate addresses

Classless interdomain routing (CIDR) is an addressing scheme (also known as supernetting) that eliminates the concept of classifying networks into class types. Earlier addressing schemes identified five classes of networks: Class A, Class B, Class C, Class D, and Class E. *Avaya Virtual Services Platform 9000 Configuration — BGP Services, NN46250-507* does not discuss Classes D (used for multicast) and E (reserved and currently not used).

Network 195.215.0.0, an illegal Class C network number, becomes a legal supernet when represented in CIDR notation as 195.215.0.0/16. The /16 is the prefix length and expresses the explicit mask that CIDR requires. In this case, the addition of the prefix /16 indicates that the subnet mask consists of 16 bits (counting from the left).

Using this method, supernet 195.215.0.0/16 represents 195.215.0.0 255.255.0.0. The following table shows the conversion of prefix length to subnet mask.

Table 3: CIDR conversion

Prefix	Dotted-decimal	Binary	Network class
/1	128.0.0.0	1000 0000 0000 0000 0000 0000 0000 0000	128 Class A
/2	192.0.0.0	1100 0000 0000 0000 0000 0000 0000 0000	64 Class A
/3	224.0.0.0	1110 0000 0000 0000 0000 0000 0000 0000	32 Class A
/4	240.0.0.0	1111 0000 0000 0000 0000 0000 0000 0000	16 Class A
/5	248.0.0.0	1111 1000 0000 0000 0000 0000 0000 0000	8 Class A
/6	252.0.0.0	1111 1100 0000 0000 0000 0000 0000 0000	4 Class A
/7	254.0.0.0	1111 1110 0000 0000 0000 0000 0000 0000	2 Class A
/8	255.0.0.0	1111 1111 0000 0000 0000 0000 0000 0000	1 Class A or 256 Class B
/9	255.128.0.0	1111 1111 1000 0000 0000 0000 0000 0000	128 Class B
/10	255.192.0.0	1111 1111 1100 0000 0000 0000 0000 0000	64 Class B
/11	255.224.0.0	1111 1111 1110 0000 0000 0000 0000 0000	32 Class B
/12	255.240.0.0	1111 1111 1111 0000 0000 0000 0000 0000	16 Class B
/13	255.248.0.0	1111 1111 1111 1000 0000 0000 0000 0000	8 Class B
/14	255.252.0.0	1111 1111 1111 1100 0000 0000 0000 0000	4 Class B
/15	255.254.0.0	1111 1111 1111 1110 0000 0000 0000 0000	2 Class B
/16	255.225.0.0	1111 1111 1111 1111 0000 0000 0000 0000	1 Class B or 256 Class C
/17	255.255.128.0	1111 1111 1111 1111 1000 0000 0000 0000	128 Class C

Prefix	Dotted-decimal	Binary	Network class
/18	255.255.192.0	1111 1111 1111 1111 1100 0000 0000 0000	64 Class C
/19	255.255.224.0	1111 1111 1111 1111 1110 0000 0000 0000	32 Class C
/20	255.255.240.0	1111 1111 1111 1111 1111 0000 0000 0000	16 Class C
/21	255.255.248.0	1111 1111 1111 1111 1111 1000 0000 0000	8 Class C
/22	255.255.252.0	1111 1111 1111 1111 1111 1100 0000 0000	4 Class C
/23	255.255.254.0	1111 1111 1111 1111 1111 1110 0000 0000	2 Class C
/24	255.255.255.0	1111 1111 1111 1111 1111 1111 0000 0000	1 Class C

Use CIDR to assign network prefixes of arbitrary lengths, as opposed to the obsolete class system, which assigned prefixes as even multiples of an octet.

For example, you can assign a single routing table supernet entry of 195.215.16/21 to represent 8 separate Class C network numbers: 195.215.16.0 through 195.215.23.0.

Supernet addressing

You can create a supernet address that covers an address range.

For example, to create a supernet address that covers an address range of 192.32.0.0 to 192.32.9.255, perform the following steps:

1. Convert the starting and ending address range from dotted-decimal notation to binary notation (see the following figure).

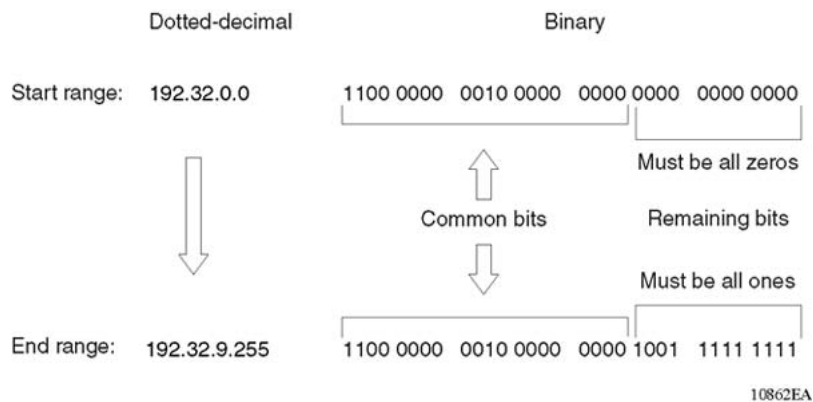
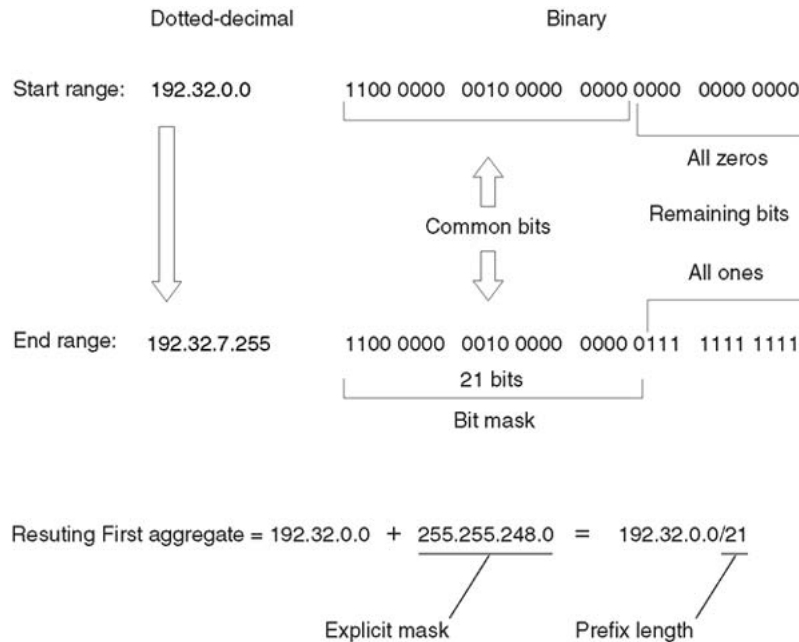


Figure 3: Binary notation conversion

2. Locate the common bits in both ranges. Ensure that the remaining bits in the start range are zeros, and the remaining bits in the end range are all ones. In this example, the remaining bits in the end range are not all ones.

3. If the remaining bits in the end range are not all ones, you must recalculate to find the IP prefix that has only ones in the remaining bits in the end range.
4. Recalculate to find a network prefix that has all ones in the remaining end range bits (see the following figure). In this example, 192.32.7.255 is the closest IP prefix that matches the common bits for the start range.



10863EA

Figure 4: First aggregate and prefix length

5. The 21 bits that match the common bits form the prefix length. The prefix length is the number of binary bits that form the explicit mask (in dotted-decimal notation) for this IP prefix.
6. The remaining aggregate is formed from 192.32.8.0 to the end range, 192.32.9.255.

As shown in [Figure 4: First aggregate and prefix length](#) on page 18, the resulting first aggregate 192.32.0.0/21 represents all of the IP prefixes from 192.32.0.0 to 192.32.7.255.

The following figure shows the results after forming the remaining aggregate from 192.32.9.0 to the end range, 192.32.9.255.

The resulting aggregate 192.32.8.0/23 represents all of the IP prefixes from 192.32.8.0 to 192.32.9.255.

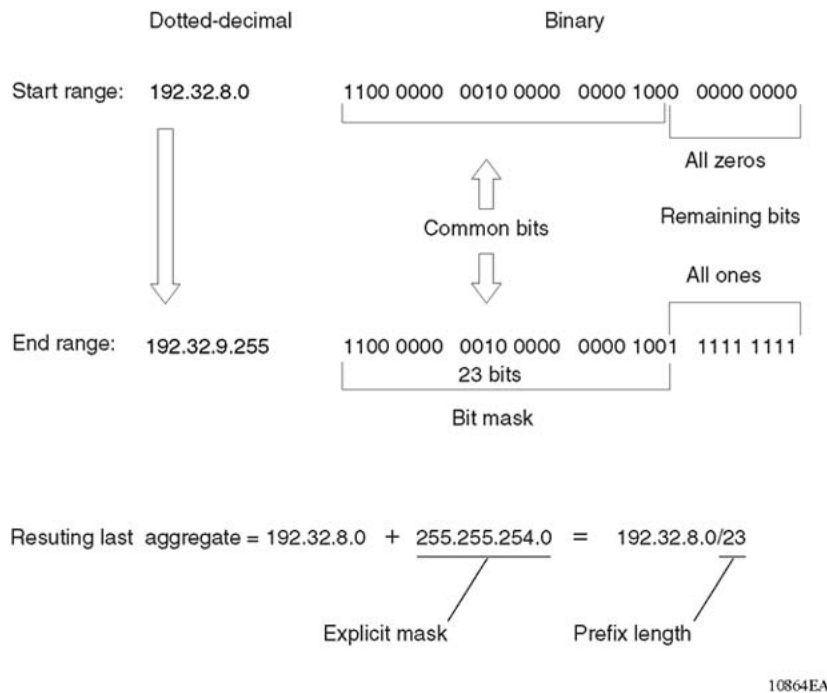


Figure 5: Last aggregate and prefix length

The final result of calculating the supernet address that ranges from 192.32.0.0 to 192.32.9.255 is as follows:

192.32.0.0 (with mask) 255.255.248.0 = 192.32.0.0/21

192.32.8.0 (with mask) 255.255.254.0 = 192.32.8.0/23

Aggregate routes

Eliminating the idea of network classes provides an easy method to aggregate routes. Rather than advertise a separate route for each destination network in a supernet, BGP uses a supernet address to advertise a single route (called an aggregate route) that represents all the destinations. CIDR also reduces the size of the routing tables used to store advertised IP routes.

The following figure shows an example of route aggregation using CIDR. In this example, a single supernet address 195.215.0.0/16 advertises 256 separate Class C network numbers 195.215.0.0 through 195.215.255.0.

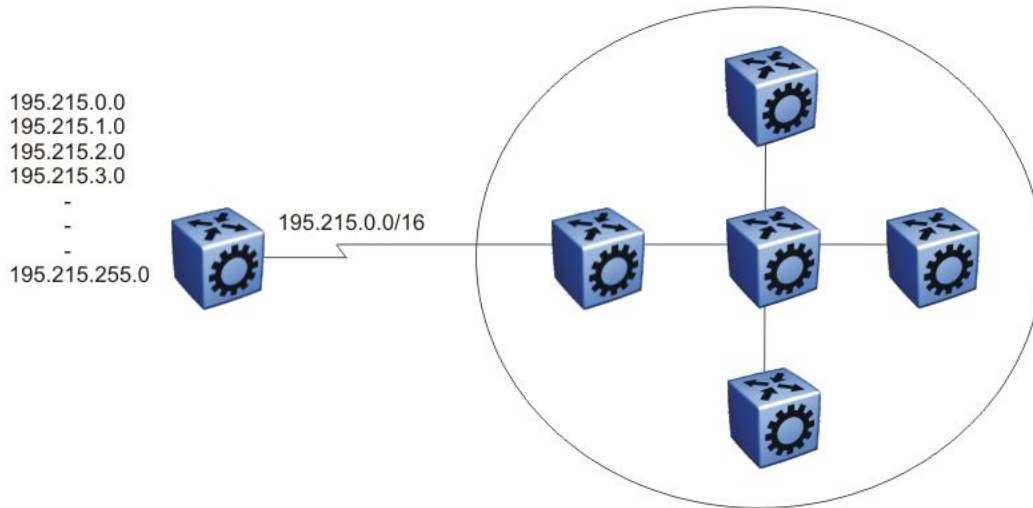


Figure 6: Aggregating routes with CIDR

Confederations

A BGP router configured for IBGP establishes a peer-to-peer session with every other IBGP speaker in the AS. In an AS with a large number of IBGP speakers, this full-mesh topology can result in high bandwidth and maintenance costs.

As shown in the following example, a full-mesh topology for an AS with 50 IBGP speakers requires 1225 internal peer-to-peer connections:

Example:

$$n \times (n-1)/2 = n \text{ IBGP sessions}$$

where:

$$50 \times (50-1)/2 = 1225 \text{ number of unique IBGP sessions}$$

You can reduce the high bandwidth and maintenance costs associated with a large full-mesh topology by dividing the AS into multiple smaller ASs (sub-ASs), and then group them into a single confederation (see the following figure).

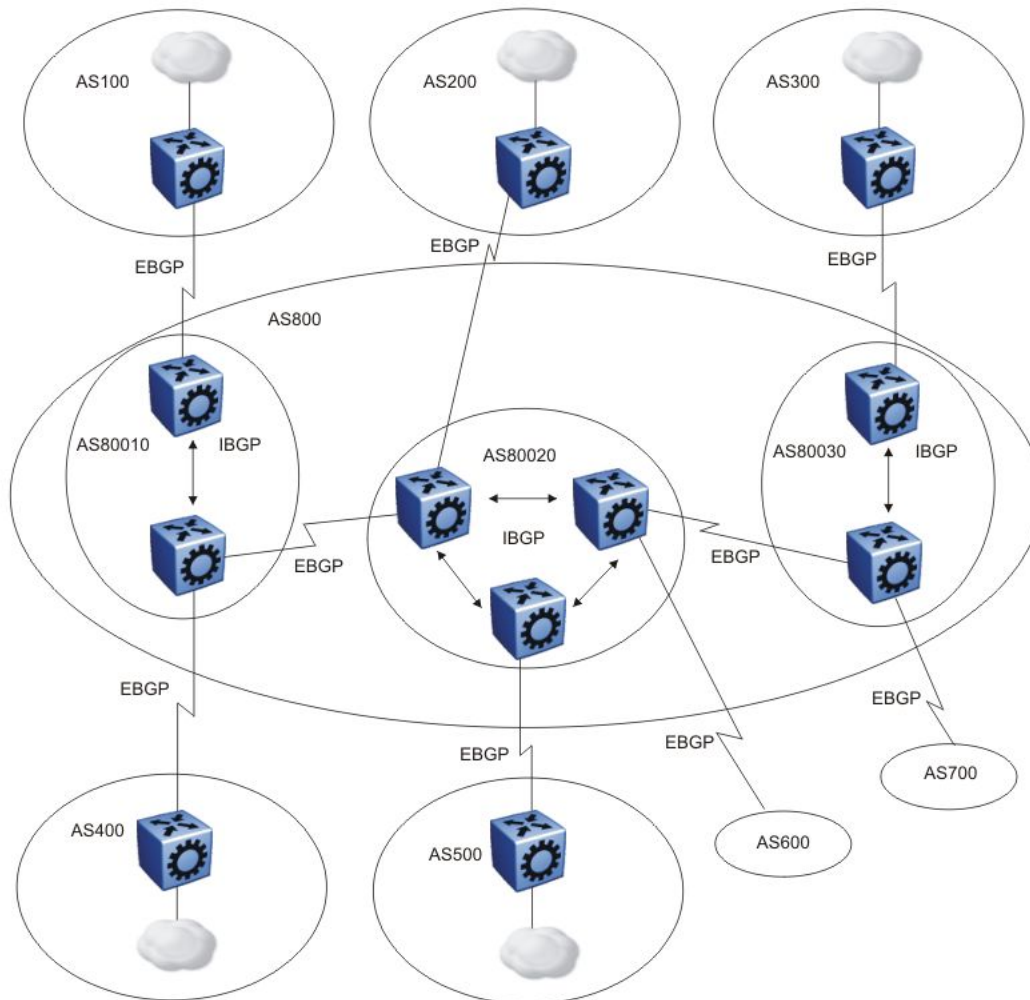


Figure 7: Confederations

As shown in the preceding figure, each sub-AS is fully meshed within itself and has EBGP sessions with other sub-ASs in the same confederation.

Although the peers in different ASs have EBGP sessions with the various sub-AS peers, they preserve the next-hop, Multi-Exit Discriminator (MED), and local preference information and exchange routing updates as if they were IBGP peers. All of the ASs retain a single interior gateway protocol (IGP). When the confederation uses its own confederation identifier, the group of sub-ASs appear as a single AS (with the confederation identifier as the AS number).

Route reflectors

Another way to reduce the IBGP mesh inherent in an AS with a large number of IBGP speakers is to configure a route reflector. Using this method, when an IBGP speaker needs to communicate with other BGP speakers in the AS, the speaker establishes a single peer-to-peer route reflector client session with the IBGP route reflector.

In an AS, more than one route reflector cluster can exist and more than one route reflector in a cluster. When more than one reflector exists in a cluster, take care to prevent route loops.

The following figure shows a simple IBGP configuration with three IBGP speakers (routers A, B, and C). Without route reflectors, after Router A receives an advertised route from an external neighbor, it must advertise the route to Routers B and C.

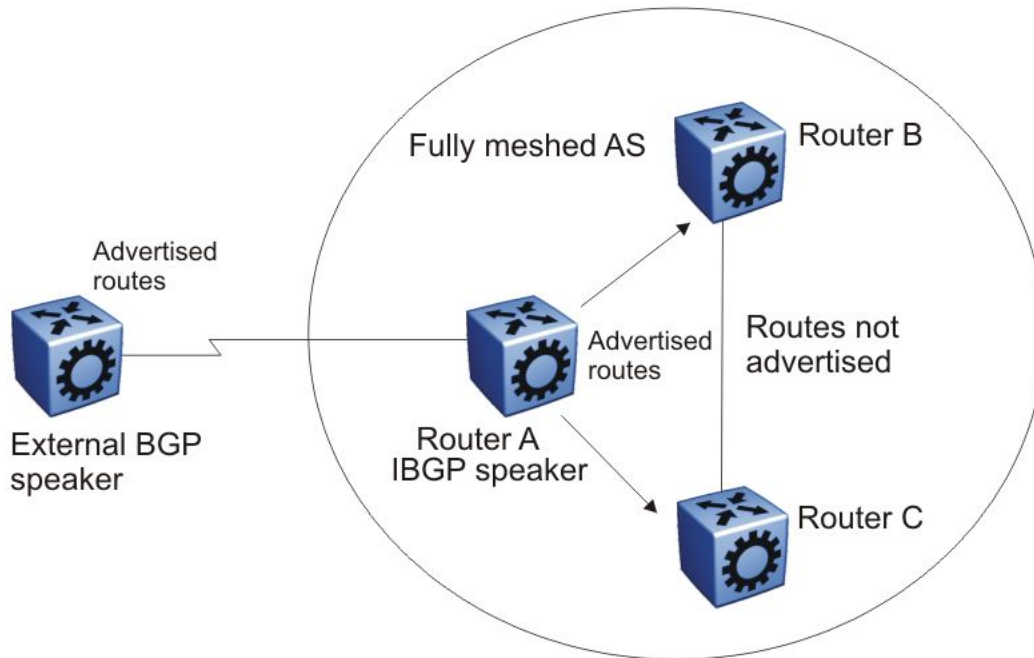


Figure 8: Fully meshed AS with IBGP speakers

Routers B and C do not readvertise the IBGP learned routes to other IBGP speakers. BGP does not allow routers to pass routes learned from internal neighbors on to other internal neighbors, which avoids routing information loops.

As shown in the following figure, when you configure an internal BGP peer (Router B) as a route reflector, all of the IBGP speakers do not need to be fully meshed. In this case, the assigned route reflector passes IBGP learned routes to a set of IBGP neighbors.




Figure 9: AS with route reflector

After Router B, the route reflector, receives routes from Router A (the IBGP speaker), it advertises them to router C. Conversely, after the route reflector receives routes from internal peers, it advertises those routes to Router A. Routers A and C do not need an IBGP session.

Route reflectors separate internal peers into two groups: client peers and nonclient peers. The route reflector and its clients form a cluster. The client peers in the cluster do not need to be fully meshed, and do not communicate with IBGP speakers outside their cluster. Nonclient peers must be fully meshed with each other.

The following figure shows a cluster, where Router A is the route reflector in a cluster with client routers B, C, and D. Routers E, F, and G are fully meshed, nonclient routers.

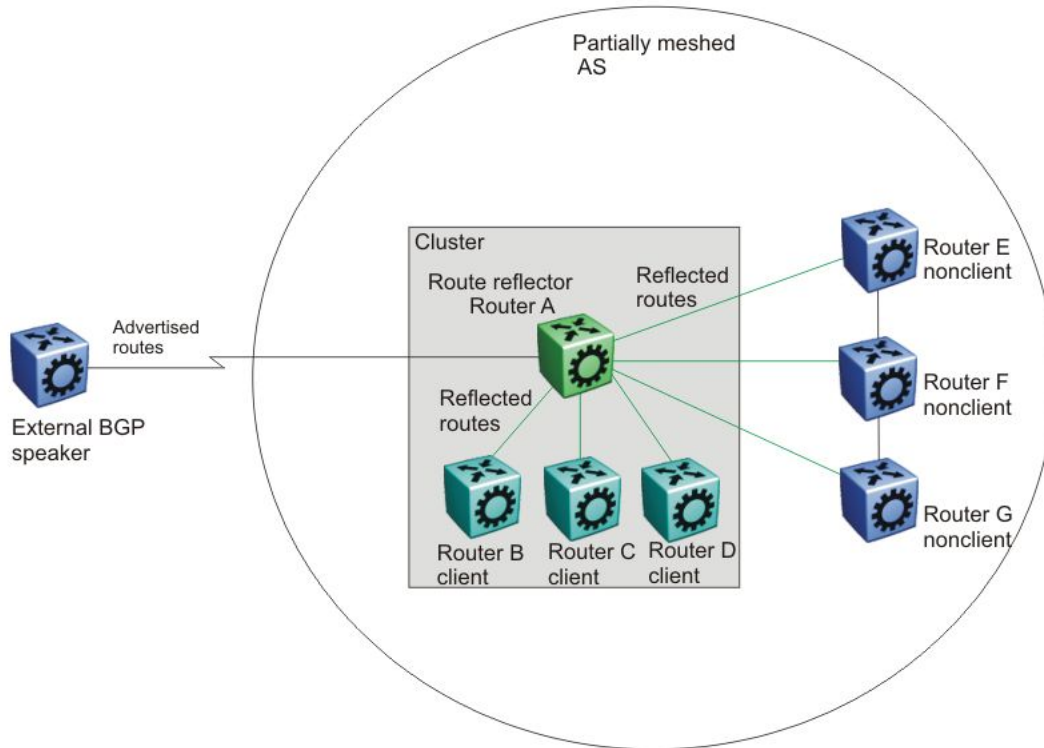


Figure 10: Route reflector with client and nonclient peers

BGP communities

You can group destinations into communities to simplify policy administration. A community is a group of destinations that share a common administrative property.

Use community control routing policies with respect to destinations. Create communities when you have more than one destination and want to share a common attribute.

The following list identifies specific community types:

- Internet—advertise this route to the Internet community
- no advertise—do not advertise to BGP peers including IBGP peers
- no export—do not advertise destinations outside of a BGP confederation
- no export subconfed—do not advertise to external BGP peers, even within the same confederation.

You can use a community to control which routing information to accept, prefer, or distribute to other BGP neighbors. If you specify the append option in the route policy, the router adds the specified community value to the existing value of the community attribute. Otherwise, the specified community value replaces a previous community value.

BGP path attributes

You can create policies that control routes, work with default routing, control specific and aggregated routes, and manipulate BGP path attributes.

Four categories of BGP path attributes exist:

- Well-known mandatory attributes must be in every BGP update message.
- Well-known discretionary attributes can be in a BGP update message.
- Optional transitive attributes are accepted and passed to other BGP peers.
- Optional non-transitive attributes can be either accepted or ignored, but must not pass along to other BGP peers.

Border routers that utilize built-in algorithms or manually configured policies to select paths use path attributes. BGP uses the following path attributes to control the path a BGP router chooses:

- origin (well-known mandatory)
- AS_path (well-known mandatory)
- next hop (well-known mandatory)
- MED attribute (optional non-transitive)
- local preference (well-known discretionary)
- atomic aggregate (well-known discretionary)
- aggregator (optional transitive)
- community (optional transitive)

For more information about path attributes in BGP updates, see [Path attributes](#) on page 29.

BGP route selection

A BGP router determines the best path to a destination network. This path is then eligible for use in the IP forwarding table and the router also advertises the path to its EBGP peers. To choose the best of multiple BGP routes to a destination, the router executes a best path algorithm.

The algorithm chooses a route in the following order:

- highest weight

Weight is a locally significant parameter associated with each BGP peer. You can use the weight to influence which peer paths the router uses.

- highest local preference

The local preference has global significance within an AS. You can manipulate the preference using route policies to influence path selection.

- prefer locally originated paths

The router prefers a path locally originated using the network, redistribution, or aggregate command over a path learned through a BGP update. The router prefers local paths sourced by network or redistribute commands over local aggregates sourced by the aggregate address command.

- shortest AS path

The AS path parameter specifies the ASs that the network prefix traversed. The AS path commonly determines the best path. For example, a router can choose a path based on whether the network passed through a specific AS. You can configure a route policy to match the AS, and then modify the local preference. Also, you can pad the AS path before the AS advertises it to a peer AS, so that downstream routers are less likely to prefer the advertised network path.

- lowest origin type

The order of preference is IGP, EGP, INC (incomplete).

- lowest MED

The MED parameter influences the preferred path from a remote AS to the advertising AS. This parameter applies when there are multiple exit points from the remote AS to the advertising AS. A lower MED value indicates a stronger path preference than a higher MED value. By default, the MED attribute is ignored as specified by the BGP global parameter Always Compare MED except when the routes come from the same AS. This parameter must be enabled for MEDs to be compared (and for this step of the best path algorithm to execute).

The router compares MEDs regardless of what the first (neighboring) AS specified in the AS_PATH. Deterministic MED, when enabled, means that the first AS of the multiple paths must be the same. Paths received with no MED are assigned a MED of 0, unless the global BGP parameter Missing Is Worst is enabled. If so, received paths are assigned a MED of 4 294 967 294. Missing Is Worst is enabled by default. The router changes paths received with a MED of 4 294 967 295 to 4 294 967 294 before insertion into the BGP table.

If you enable the BGP global parameter to compare MED values within a confederation, the system compares MEDs for all paths that consist only of AS_CONFED_SEQUENCE (paths originated within the local confederation). This parameter is disabled by default.

- lowest IGP metric to the BGP next-hop

If multiple paths exist whose BGP next-hop is reachable through an IGP, the path with the lowest IGP metric to the BGP next-hop is chosen.

- prefer external paths (learned by EBGP) over internal paths (IBGP)

The system prefers external paths over internal paths

- if Equal Cost Multipath (ECMP) is enabled, insert up to four paths in the routing table

If you enable ECMP, multiple BGP learned routes that use the same metric to different IP next-hops are installed in the IP forwarding table for traffic load-balancing purposes.

- lowest router ID

The lowest router ID, or Circuitless IP (CLIP) address, is preferred.

BGP and dampened routes

Virtual Services Platform 9000 supports route dampening (route suppression). When you use route dampening, a route accumulates penalties each time the route fails. After the accumulated penalties exceed a threshold, the router no longer advertises the route. The router enters the suppressed routes into the routing table only after the accumulated penalty falls below the reuse threshold.

Route flap dampening suppresses the advertisement of the unstable route until the route becomes stable.

Dampening applies only to routes that are learned through an EBGP. Route flap dampening prevents routing loops and protects IBGP peers from having higher penalties for routes external to the AS.

The following paragraph describes the algorithm that controls route flaps.

After the route flaps the first time

- the router creates a route history entry
- a timer starts (180 seconds)

If the route does not flap again, the router uses this timer to delete the history entry after the 180 seconds expires.

After the route flaps a second time

- The penalty is recalculated based on the decay function.

If the penalty is greater than the cut-off value (1536), the route is suppressed and the reuse time is calculated based on the reuse time function.

- The reuse timer starts.

After the reuse time expires, the suppressed route is announced again (the reuse time is recalculated if the route flaps again). The penalty decays slower for withdrawn routes than

for update routes. The route history entry is kept longer if the route is withdrawn. For update history, the delete time is 90 seconds and the withdrawn history delete time is 180 seconds.

BGP updates

BGP uses update messages to communicate information between two BGP speakers. The update message can advertise a single feasible route to a peer, or withdraw multiple unfeasible routes from service.

The following figure shows the format of an update message.




Figure 11: Update message format

This section describes how BGP uses the update message fields to communicate information between BGP speakers.

Withdrawn routes length

The withdrawn routes length parameter (referred to in RFC1771 as the Unfeasible Routes Length field) indicates the total length of the withdrawn routes field in octets. The withdrawn routes length field calculates the length of the NLRI field. For example, a value of 0 indicates that no routes are withdrawn from service, and that the withdrawn routes field is not present in this update message.

Withdrawn routes

The withdrawn routes parameter is a variable-length parameter that contains a list of IP prefixes for routes that are withdrawn from service. The following figure shows the format of an IP prefix.




Figure 12: IP Prefix format

The length indicates the number of bits in the prefix (also called the network mask).

For example, 195.215.0.0/16 is equivalent to 195.215.0.0 255.255.0.0 (the /16 indicates the number of bits in the length parameter to represent the network mask 255.255.0.0).

The prefix parameter contains the IP address prefix itself, followed by enough trailing bits to make the length of the whole field an integer multiple of 8 bits (1 octet).

Total path attributes length

The total path attributes length parameter indicates the total length of the path attributes parameter in octets.

The total path attributes length calculates the length of the NLRI parameter. For example, a value of 0 indicates that no NLRI field is present in this update message.

Path attributes

The path attributes parameter is a variable-length sequence of path attributes that exists in every BGP update. The path attributes contain BGP attributes associated with the prefixes in the NLRI parameter.

For example, the attribute values allow you to specify the prefixes that the BGP session can exchanged, or which of the multiple paths of a specified prefix to use.

The attributes carry the following information about the associated prefixes:

- the path origin
- the AS paths through which the prefix is advertised
- the metrics that display degrees of preference for this prefix

The following figure shows the encoding used with the path attribute parameter.




Figure 13: Path attribute encoding

Attribute type

As shown in the following figure, the attribute type is a two-octet field that comprises two sub-fields: attribute flags and attribute type code.




Figure 14: Attribute Type fields

The attribute flags parameter is a bit string that contains four binary values that describe the attribute, and four unused bits. The bit descriptions (from the high-order bit to the low-order bit) are

- The high-order bit (bit 0) is the optional bit. When this bit is set (the value is 1), the attribute is optional. When this bit is clear (the value is 0), the attribute is well-known. Well-known attributes must be recognized by all BGP implementations and, when appropriate, passed on to BGP peers. Optional attributes are not required in all BGP implementations.
- The second high-order bit (bit 1) is the transitive bit. For well-known attributes, this bit must be set to 1. For optional attributes, it defines whether the attribute is transitive (when set to 1) or non-transitive (when set to 0).

- The third high-order bit (bit 2) is the partial bit. The partial bit defines whether the information in the optional transitive attribute is partial (when set to 1) or complete (when set to 0). For well-known attributes and for optional non-transitive attributes the partial bit must be set to 0.
- The fourth high-order bit (bit 3) is the extended length bit. The extended length bit defines whether the attribute length is one octet (when set to 0) or two octets (when set to 1). The attribute flag can use the extended length only if the length of the attribute value is greater than 255 octets.
 - If the extended length bit of the attribute flags octet is set to 0, the third octet of the path attribute contains the length of the attribute data in octets.
 - If the extended length bit of the attribute flags octet is set to 1, then the third and the fourth octets of the path attribute contain the length of the attribute data in octets.
- The lower-order four bits of the attribute flags octet are unused. The lower-order four bits must be zero (and must be ignored when received).

The attribute type code parameter contains the attribute type code, as defined by the Internet Assigned Numbers Authority (IANA). The attribute type code uniquely identifies the attribute from all others. The remaining octets of the path attribute represent the attribute value and are interpreted according to the attribute flags and the attribute type code parameters.

The following table shows the supported attribute type codes.

Table 4: BGP mandatory path attributes

Attribute	Type code	Description
Origin	1	Defines the origin of the path information: <ul style="list-style-type: none"> • Value = 0 --- IGP (the path is valid all the way to the IGP of the originating AS) • Value = 1--- EGP (the last AS in the AS path uses an EGP to advertise the path) • Value = 2--- Incomplete (the path is valid only to the last AS in the AS path)
AS path	2	Contains a list of the ASs that packets must traverse to reach the destinations. This code represents each AS path segment as follows: <ul style="list-style-type: none"> • path segment type • path segment length • path segment value
Next hop	3	Specifies the IP address of the border router to use as a next hop for the advertised destinations (destinations listed in the NLRI field of the update message).
Multiexit discriminator	4	Discriminates among multiple exit or entry points to the same neighboring AS on external (internal-AS) links.

Attribute	Type code	Description
Local preference	5	Indicates the preference that AS border routers assign to a chosen route when they advertise it to IBGP peers
Atomic aggregate	6	Ensures that certain NLRI is not deaggregated
Aggregator	7	Identifies which AS performed the most recent route aggregation. This attribute contains the last AS number that formed the aggregate route followed by the IP address of the BGP speaker that formed the aggregate route.

Attribute length

The attribute length can be one or two octets in length, depending on the value of the extended length parameter in the attributes flag field.

This parameter indicates the length of the attribute value field.

Attribute value

The attribute value contains the actual value of the specific attribute. The system implements the attribute value according to the values in the attribute flags and the attribute type code parameters.

NLRI

The NLRI parameter is a variable length field that contains a list of prefixes. The packet size that BGP speakers can exchange limits the number of prefixes in the list.

Equal Cost Multipath

Equal Cost Multipath (ECMP) support allows a BGP speaker to perform route or traffic balancing within an AS by using multiple equal-cost routes submitted to the routing table by OSPF, RIP, or static routes. For more information about ECMP, see *Avaya Virtual Services Platform 9000 Configuration — IP Routing, NN46250-505*.

MD5 message authentication

Authenticate BGP messages by using Message Digest 5 (MD5) signatures. After you enable BGP authentication, the BGP speaker verifies that the BGP messages it receives from its peers are actually from a peer and not from a third party masquerading as a peer.

BGPv4 TCP MD5 message authentication provides the following features:

- A TCP MD5 signature can exist for BGP peers. You can configure authentication and secret keys for each peer. Peers configured with common secret keys can authenticate each other and exchange routing information.
- The switch can concurrently have BGP peers with authentication enabled and other BGP peers with authentication disabled.
- The switch always encrypts the secret keys.

After you enable BGPv4 TCP MD5 authentication, the router computes an MD5 signature for each TCP packet based on the TCP packet and an individual peer secret key. The router adds this MD5 signature to the TCP packet that contains a BGP message and sends it with the packet, but it does not send the secret key.

The receiver of the TCP packet also knows the secret key and can verify the MD5 signature. A third party that tries to masquerade as the sender, however, cannot generate an authentic signature because it does not know the secret key.

In ACLI commands, the term password refers to the secret key. The secret keys provide security. If the keys are compromised, then the authentication itself is compromised. To prevent this, the switch stores the secret keys in encrypted form.

MD5 signature generation

BGP peers calculate MD5 signatures in BGP messages based on the following elements:

- TCP pseudo-header
- TCP header, excluding options
- TCP segment data
- TCP MD5 authentication key

If TCP receives an MD5 authentication key, it reduces its maximum segment size by 18 octets, which is the length of the TCP MD5 option. TCP adds an MD5 signature to each transmitted packet. The peer inserts the resulting 16-byte MD5 signature into the following TCP options: kind=19, length=18.

MD5 signature verification

After the switch receives a packet, it performs three tests. The following table lists the tests and the event message that TCP logs if a test fails.

Table 5: MD5 signature verification rules on BGP TCP packets

Condition tested	Action on success	Failure event message
Is the connection configured for MD5 authentication?	Verify that the packet contains a kind=19 option.	TCP MD5 No Signature
Is MD5 authentication enabled for this TCP connection?	TCP computes the expected MD5 signature.	TCP MD5 Authentication Disabled

Condition tested	Action on success	Failure event message
Does the computed MD5 signature match the received MD5 signature?	TCP sends the packet to BGP.	TCP MD5 Invalid Signature

If a packet passes a test, it proceeds to the next test. After a packet passes all three tests, TCP accepts the packet and sends it to BGP.

If a packet fails a test, the switch logs an event, increments the count of TCP connection errors (wfTcpConnMd5Errors), and discards the packet. The TCP connection remains open.

BGP and route redistribution

Redistribution imports routes from one protocol to another. Redistribution sends route updates for a protocol-based route through another protocol. For example, if OSPF routes exist in a router and they must travel through a BGP network, then configure redistribution of OSPF routes through BGP. This sends OSPF routes to a router that uses BGP.

Virtual Services Platform 9000 can redistribute routes

- on an interface basis
- on a global basis between protocols on a single VRF instance (intraVRF)
- between the same or different protocols on different VRF instances (interVRF)

Configure interface-based redistribution by configuring a route policy and apply it to the interface. Configure the match parameter to the protocol from which to learn the routes.

You can redistribute routes on a global basis, rather than on an interface basis. Use the `ip bgp redistribute` command to accomplish the (intraVRF) redistribution of routes through BGP, so that BGP redistribution occurs globally on all BGP-enabled interfaces. This redistribution does not require a route policy, but you can use one for more control.

If you configure redistribution globally and on an interface, redistribution through the route policy takes precedence.

You can redistribute routes from a protocol in one VRF to BGP in another VRF. You can use a route policy for redistribution control. If you enable route redistribution between VRF instances, ensure that IP addresses do not overlap.

Use caution when you configure redistribution. An improperly configured parameter can cause the router to advertise learned EBGP routes out of your local AS. If this happens, the local AS can route other networks.

Do not use redistribution if you peer to an Internet Service Provider (ISP) and do not want traffic to transit your local AS.

When you redistribute OSPF routes into BGP, route priorities can create routing loops. Because BGP has a higher route preference than OSPF external type 1 and 2 routes, if you

redistribute OSPF external type 1 and 2 routes into BGP, the router uses the BGP routes, which can cause a routing loop.

Virtual Services Platform 9000 supports route redistribution between BGP, RIP, OSPF, direct, static, IPv6–direct, IPv6–static, and OSPFv3 routes. For more information about RIP and OSPF route redistribution, see *Avaya Virtual Services Platform 9000 Configuration — OSPF and RIP, NN46250-506*.

BGP+

VSP 9000 extends the BGPv4 process to support the exchange of IPv6 routes using BGPv4 peering. BGP+ is an extension of BGPv4 for IPV6.

Note that the VSP 9000 BGP+ support is not an implementation of BGPv6. Native BGPv6 peering uses the IPv6 Transport layer (TCPv6) for establishing the BGPv6 peering, route exchanges, and data traffic. Native BGPv6 peering is not supported.

VSP 9000 supports the exchange of BGP+ reachability information over IPv4 transport. To support BGP+, the VSP 9000 switch supports two BGP protocol extensions, standards RFC 4760 (multi-protocol extensions to BGP) and RFC 2545 (MP-BGP for IPv6). These extensions allow BGPv4 peering to be enabled with IPv6 address family capabilities.

The VSP 9000 implementation of BGP+ uses an existing TCPv4 stack to establish a BGPv4 connection. Optional, nontransitive BGP properties are used to transfer IPV6 routes over the BGPv4 connection. Any BGP+ speaker has to maintain at least one IPv4 address to establish a BGPv4 connection.

Different from IPv4, IPv6 introduces scoped unicast addresses, identifying whether the address is global or link-local. When BGP+ is used to convey IPv6 reachability information for interdomain routing, it is sometimes necessary to announce a next hop attribute that consists of a global address and a link-local address. For BGP+, no distinction is made between global and site-local addresses.

The BGP+ implementation includes support for BGPv6 policies, including redistributing BGPv6 into OSPFv3, and advertising OSPFv3, static, and local routes into BGPv6 (through BGP+). It also supports the aggregation of global unicast IPv6 addresses and partial HA.

When configuring BGP+ on the router that is enabled only for IPv6 (the router does not have an IPv4 address), then BGP router ID must be manually configured for the router.

BGP+ does not support confederations. In this release, you can configure confederations for IPv4 routes only.

The basic configuration of BGP+ is the same as BGPv4 with one additional parameter added and some existing commands altered to support IPv6 capabilities. You can enable and disable IPv6 route exchange by specifying the address family attribute as IPv6. Note that an IPv6 tunnel is required for the flow of IPv6 data traffic.

BGP+ is only supported on the global VRF instance.

BGP+ tunnel

When you use BGP+ you must configure an IPv6 tunnel and static routes at BGP+ peers.

When BGP+ peers advertise route information, they use Update messages to advertise route information.

These RTM routes contain next-hop addresses from the BGP peer that the route was learned from.

The static routes correlate the next-hop addresses represented by the IPv4-Mapped IPv6 address to a specific outgoing interface.

Example: static route in an IPv6 tunnel configuration for BGP+:

Following is one way to express a static route in an IPv6–configured tunnel for BGP+:

```
ipv6 route 0:0:0:0:fff:30.1.1.2/28 cost 1 tunnel 10 where 0:0:0:0:fff:30.1.1.2 is the IPv4-Mapped IPv6 address of the BGP peer at 30.1.1.2
```

Circuitless IP

Circuitless IP (CLIP) is a virtual (or loop back) interface that you do not associate with a physical port. You can use a CLIP interface to provide uninterrupted connectivity to your switch as long as an actual path exists to reach the device. For example, as shown in the following figure, a physical point-to-point link exists between R1 and R2 along with the associated addresses (195.39.1.1/30 and 195.39.1.2/30). Note also that an IBGP session exists between two additional addresses 195.39.128.1/32 (CLIP 1) and 195.39.128.2/32 (CLIP 2).




Figure 15: Routers with IBGP connections

The system treats the CLIP interface like an IP interface and treats the network associated with the CLIP as a local network attached to the device. This route always exists and the circuit is always up because no physical attachment exists.

The router advertises routes to other routers in the domain either as external routes using the route-redistribution process or after you enable OSPF in a passive mode to advertise an OSPF internal route. You can configure only the OSPF protocol on the CLIP interface. After you create

a CLIP interface, the system software programs a local route with the CPU as the destination ID. The CPU processes all packets destined to the CLIP interface address. The system treats other packets with destination addresses associated with this network (but not to the interface address) as if they are from an unknown host.

A circuitless IP or CLIP address is a logical IP address for network management, as well as other purposes. The CLIP is typically a host address (with a 32 bit subnet mask). Avaya recommends that you configure the OSPF router ID to the configured CLIP address. By default, the BGP router ID is automatically equivalent to the OSPF router ID.

For information about how to configure CLIP interfaces, see *Avaya Virtual Services Platform 9000 Configuration — IP Routing, NN46250-505*.

BGP configuration considerations and limitations

Use the information in this section to help you configure BGP on the Avaya Virtual Services Platform 9000. The Virtual Services Platform 9000 supports BGPv4 as described in RFC1771.

BGP implementation guidelines

The following list provides guidelines to successfully implement BGP:

- BGP does not operate with an IP router in nonforwarding (host-only) mode. Make sure that the routers you want BGP to operate with are in forwarding mode.
- If you use BGP for a multihomed AS (one that contains more than a single exit point), Avaya recommends that you use OSPF for your IGP and BGP for your sole exterior gateway protocol, or use intra-AS IBGP routing.
- If OSPF is the IGP, use the default OSPF tag construction. Using EGP or modifying the OSPF tags makes network administration and proper configuration of BGP path attributes difficult.
- For routers that support both BGP and OSPF, the OSPF router ID and the BGP identifier must be the same IP address. The BGP router ID automatically uses the OSPF router ID.
- In configurations where BGP speakers reside on routers that have multiple network connections over multiple IP interfaces (the typical case for IBGP speakers), consider using the address of the circuitless (virtual) IP interface as the local peer address. In this configuration, you ensure that BGP is reachable as long as an active circuit exists on the router.
- By default, BGP speakers do not advertise or inject routes into the IGP. You must configure route policies to enable route advertisement.
- Coordinate routing policies among all BGP speakers within an AS so that every BGP border router within an AS constructs the same path attributes for an external path.
- Configure accept and announce policies on all IBGP connections to accept and propagate all routes. Make consistent routing policy decisions on external BGP connections.

Minimum requirements

On the Virtual Services Platform 9000, you must configure the following minimum parameters:

- router ID
- local AS number
- enable BGP globally
- BGP neighbor peer session: remote IP addresses
- enable BGP peers
- When you use both BGP and OSPF, the OSPF and BGP router ID must be the same.

The router ID must be a valid IP address of an IP interface on the router or a CLIP address. BGP update messages use this IP address. By default, the BGP router ID automatically uses the OSPF router ID.

You cannot configure the BGP router ID if you configure BGP before you configured the OSPF router ID. You must first disable BGP, configure the OSPF route ID, and then enable BGP globally.

You can add BGP policies to the BGP peer configuration to influence route decisions. BGP policies apply to the peer through the ACLI soft-reconfiguration commands.

After you configure the Virtual Services Platform 9000 for BGP, some parameter changes can require you to enable or disable the BGP global state or the neighbor admin-state.

You can dynamically modify BGP policies. On the global level, the BGP redistribution command has an apply parameter that causes the policy to take effect after you issue the command.

BGP neighbor maximum prefix configuration

By default, the maximum prefix parameter limits 12 000 NLRI messages for each neighbor. The maximum prefix parameter limits the number of routes that the Virtual Services Platform 9000 can accept.

The maximum prefix parameter prevents large numbers of BGP routes from flooding the network if you implement an incorrect configuration. You can assign a value to the maximum prefix limit, including 0 (0 means unlimited routes). When you configure the maximum prefix value, consider the maximum number of active routes that your equipment configuration can support.

BGP and OSPF interaction

RFC1745 defines the interaction between BGP and OSPF when OSPF is the IGP within an autonomous system. For routers that use both protocols, the OSPF router ID and the BGP ID must be the same IP address. You must configure a BGP route policy to allow BGP advertisement of OSPF routes.

Interaction between BGPv4 and OSPF can advertise supernets to support CIDR. BGPv4 supports interdomain supernet advertisements; OSPF can carry supernet advertisements within a routing domain.

Chapter 4: BGP configuration using ACLI

About this task

Configure the Border Gateway Protocol (BGP) using ACLI to create and maintain an interdomain routing system that guarantees loop-free routing information between autonomous systems (AS).

For information about how to configure route policies for BGP, see *Avaya Virtual Services Platform 9000 Configuration — IP Routing, NN46250-505*.

Configuring BGP globally

Before you begin

- To configure the suppress-map, advertise-map, or attribute-map options, the route policy for those options must exist.
- For initial BGP configuration, you must know the AS number.
- You must log on to Global Configuration mode in ACLI to configure the GlobalRouter.
- You configure BGP on a VRF instance the same way you configure the GlobalRouter, except that you must use VRF Router Configuration mode and the prefix `ip bgp`. The VRF must have an RP Trigger of BGP. Not all parameters are configurable on non0 VRFs.

About this task

Configure BGP globally to enable BGP on the switch and determine how BGP operates.

Procedure

1. Specify the AS number and enable BGP:

```
router bgp [WORD <0-11>] [enable]
```

The AS number parameter only applies to VRF 0.
2. Access Router BGP Configuration mode:

```
router bgp
```
3. Configure BGP variables or accept the default values.

Example

Specify the AS number and enable BGP:

```
VSP-9012:1(config)#router bgp 3 enable
```

Access Router BGP Configuration mode:

```
VSP-9012:1(config)#router bgp
```

```
VSP-9012:1(router-bgp)#
```

Variable definitions

Use the data in the following table to use the `router bgp` command.

Variable	Value
<i>WORD</i> <0-11>	Specifies the AS number. You cannot enable BGP until you change the local AS to a value other than 0.
enable	Enables BGP on the router.

Use the data in the following table to use the BGP variables in BGP Router Configuration mode.

Table 6: Variable definitions

Variable	Value
auto-peer-restart enable	Enables the process that automatically restarts a connection to a BGP neighbor. The default value is enable.
auto-summary	When enabled, BGP summarizes networks based on class limits, for example, Class A, B, and C networks. The default value is enable.
bgp always-compare-med	Enables the comparison of the multi-exit discriminator (MED) parameter for paths from neighbors in different autonomous systems. The system prefers a path with a lower MED over a path with a higher MED. The default value is disable.
bgp client-to-client reflection	Enables or disables route reflection between two route reflector clients. This variable applies only if the route reflection value is enable. The default value is enable. You can enable route reflection even when clients are fully meshed. This variable only applies to VRF 0. Example: <code>VSP-9012:1(router-bgp)#bgp client-to-client reflection</code> System Response: Restart or soft-restart BGP for the change to take effect.

Variable	Value
bgp cluster-id {A.B.C.D}	Configures a cluster ID. This variable applies only if the route reflection value is enable, and if multiple route reflectors are in a cluster. {A.B.C.D} is the IP address of the reflector router. This variable only applies to VRF 0. Example: VSP-9012:1(router-bgp)#bgp cluster-id 0.0.0.0
bgp confederation identifier <0-65535> [peers WORD<0-255>]	Configures a BGP confederation. identifier <0-65535> specifies the confederation identifier. peers WORD<0-255> lists adjoining ASs that are part of the confederation in the format (5500,65535,0,10,...,...). Use quotation marks (") around the list of ASs. Example: VSP-9012:1(router-bgp)#bgp confederation identifier 1 peers "20 30 40"
bgp default local-preference <0-2147483647>	Specifies the default value of the local preference attribute. The default value is 0. You must disable BGP before you can change the default value. Example: VSP-9012:1(router-bgp)#bgp default local-preference 2-12
bgp deterministic-med enable	Enables deterministic MED. Example: VSP-9012:1(router-bgp)#bgp deterministic-med enable
bgp multiple-paths <1-8>	Configures the maximum number of equal-cost-paths that are available to a BGP router by limiting the number of equal-cost-paths the routing table can store. The default value is 1. Example: VSP-9012:1(router-bgp)#bgp multiple-paths 4
comp-bestpath-med-confed enable	When enabled, compares MED attributes within a confederation. The default value is disable. This variable only applies to VRF 0. Example: VSP-9012:1(router-bgp)#comp-bestpaht-med-confed enableRestart or soft-restart BGP for the change to take effect
debug-screen <off on>	Displays debug messages on the console, or saves them in a log file. Disable BGP screen logging (off) or enable BGP screen logging (on). Example: VSP-9012:1(router-bgp)#debug-screen onSystem Response: BGP Screen Logging is On

Variable	Value
default-information originate	Enables the advertisement of a default route to peers, if the route exists in the routing table. The default value is disable.
default-metric <-1-2147483647>	Configures a value to send to a BGP neighbor to determine the cost of a route a neighbor uses. A default metric value solves the problems associated with redistributing routes that use incompatible metrics. For example, whenever metrics do not convert, using a default metric provides a reasonable substitute and redistribution proceeds. Use this option in conjunction with the redistribute commands so the current routing protocol uses the same metric for all redistributed routes. The default value is 0.
flap-dampening enable	Enables route suppression for routes that flap on and off. The default value is disable.
global-debug mask WORD<1-100>	<p>Displays specified debug information for BGP global configurations. The default value is none.</p> <ul style="list-style-type: none"> • <WORD 1-100> is a list of mask choices separated by commas with no space between choices. <p>Mask choices are:</p> <ul style="list-style-type: none"> none disables all debug messages. all enables all debug messages. error enables display of debug error messages. packet enables display of debug packet messages. event enables display of debug event messages. trace enables display of debug trace messages. warning enables display of debug warning messages. state enables display of debug state transition messages. init enables display of debug initialization messages. filter enables display of debug messages related to filtering. update enables display of debug messages related to sending and receiving updates. <p>Example: VSP-9012:1(router-bgp)#global-debug mask event, trace, warning, state</p>
ibgp-report-import-rt enable	Configures BGP to advertise imported routes to an interior BGP (IBGP) peer. This variable enables or disables advertisement of nonBGP imported routes to other IBGP neighbors. The default value is enable.

Variable	Value
ignore-illegal-rtrid enable	When enabled, BGP overlooks an illegal router ID. For example, you can configure this variable to enable or disable the acceptance of a connection from a peer that sends an open message using a router ID of 0 (zero). The default value is enable.
neighbor-debug-all mask WORD<1-100>	Displays specified debug information for BGP neighbors. The default value is none. For mask options, see the global-debug mask WORD<1-100> variable. Example: VSP-9012:1(router-bgp)#neighbor-debug-all mask error, packet, event.trace, state, filter
no-med-path-is-worst enable	Enables BGP to treat an update without a MED attribute as the worst path. The default value is disable.
quick-start enable	Enables the quick-start flag for exponential backoff.
route-reflector enable	Enables the reflection of routes from IBGP neighbors. The default value is enable. This variable only applies to VRF 0.
route-refresh	Enables or disables route refresh. If enabled, a route refresh request received by a BGP speaker causes the speaker to resend all route updates it contains in its database that are eligible for the peer that issues the request. This variable only applies to VRF 0.
router-id {A.B.C.D}	Specifies the BGP router ID in IP address format. This variable only applies to VRF 0.
synchronization	Enables the router to accept routes from BGP peers without waiting for an update from the IGP. The default value is enable.
traps enable	Enables BGP traps.

Job aid

Use debug command values to control debug messages for global BGP message types, and for message types associated with a specified BGP peer or peer group.

The following tips can help you use the debug commands:

- Display debug commands for multiple mask choices by entering the mask choices separated by commas, with no space between choices.
- To end (disable) the display of debug messages, use the mask choice of none.
- You can save debug messages in a log file, or you can display the messages on your console using the debug-screen command.

For more information about the logged debug messages, see *Avaya Virtual Services Platform 9000 Logs Reference, NN46250-702*.

Configuring 4 byte AS numbers

Configure Autonomous System (AS) numbers using the 4-byte format and represent the numbers in octets.

Before you begin

- You cannot modify the global BGP configuration unless BGP is disabled.
- Configure the local AS number at Global Router (VRF0) only.
- Access Global Configuration mode.
- Make sure that you define AS numbers in policies the same way that you configure them for the router. The AS list for the route policies accepts AS number only in the asplain format. If you create policies using asplain and configure the switch with asdot, the match will not occur.

About this task

Use BGP 4 byte AS numbers to ensure the continuity of loop-free inter-domain routing information between ASs and to control the flow of BGP updates as 2 byte AS numbers will deplete soon. AS Plain notation format is the default and the preferred form of representing 4-byte AS numbers over the AS dot notation format. You have an option to configure AS dot notation format as well. With AS dot notation, analyzing and troubleshooting any issues encountered becomes difficult as it is incompatible with the regular expressions used by most of the network providers.

Procedure

1. Disable BGP to change the AS number format.
`no router bgp enable`
2. Enable the 4-byte AS numbering format.
`router bgp as-4-byte enable`
3. To use the dotted octet notation, enable as-dot.
`router bgp as-dot enable`

4. Configure the 4-byte AS number and enable BGP. If you have enabled as-dot, enter the AS number in octet.

```
router bgp WORD<0-11> enable
```

Example

Disable BGP to change the AS number format.

```
VSP-9012:1(config)#no router bgp enable
```

Enable the 4-byte AS numbering format.

```
VSP-9012:1(config)#router bgp as-4-byte enable
```

To use the dotted octet notation, enable as-dot.

```
VSP-9012:1(config)#router bgp as-dot enable
```

Configure the 4-byte AS number and enable BGP.

```
VSP-9012:1(config)#router bgp 65536 enable
```

Variable definitions

Use the data in the following table to use the `router bgp` command.

Variable	Value
as-4-byte <enable>	Enables the switch for using 4 byte numbers for Autonomous systems. The default value is disable.
as-dot <enable>	Enables or disables representing AS numbers in octets. The default is disable so the switch uses the plain notation format. If you enable the 4-byte-as and as-dot parameters, enter numbers in the range of 1.0 to 65535.65535. The default value is disable. Attention: This parameter is now supported on VSP 9000 with BGP+.
WORD <0-11> enable	Sets the local autonomous system (AS) number. You cannot change local-as when BGP is set to enable. <ul style="list-style-type: none"> To set a 2-byte local AS number, enter a local-as number in the range of 0 to 65535. To set a 4-byte local-as number, enable the 4-byte as variable and enter a number in the range of 0 to 4294967295.

Variable	Value
	<p>Note:</p> <p>If as-4-byte is set to false, the range for AS number is 0–65535 and if as-4-byte is set to true, the range is 0–4294967295.</p> <p>If you enable as-dot, enter the AS number in octets in the range of 1.0 to 65535.65535.</p> <p>Attention: This parameter is not supported with BGP +.</p>

Configuring IPv4 and IPv6 aggregate routes

Before you begin

- Disable BGP before you enable aggregation.
- You need the appropriate aggregate address and mask.
- If required, policies exist.
- You must log on to BGP Router Configuration mode in ACLI to configure the GlobalRouter.
- You configure BGP on a VRF instance the same way you configure the GlobalRouter, except that you must use VRF Router Configuration mode and the prefix `ip bgp`. The VRF must have an RP Trigger of BGP. Not all parameters are configurable on non0 VRFs.

About this task

Configure aggregate routes so that the router advertises a single route (aggregate route) that represents all destinations. Aggregate routes also reduce the size of routing tables.

Procedure

1. Enable BGP aggregation:


```
bgp aggregation enable
```
2. Add an aggregate route to the routing table:


```
aggregate-address WORD<1-256> {advertise-map WORD<0-1536>}
[as-set] [attribute-map WORD<0-1536>] [summary-only]
[suppress-map WORD<0-1536>]
```
3. Exit to Global Configuration mode:


```
exit
```
4. Enable BGP:

```
router bgp [<0-65535>] [enable]
```

Example

Add an aggregate route to the routing table:

```
VSP-9012:1(router-bgp)#aggregate-address 60:60:60:40::/64 advertise-  
map map1 attribute-map map2
```

Enable BGP:

```
router bgp 4 enable
```

Variable definitions

Use the data in the following table to use the `aggregate-address` command.

Table 7: Variable definitions

Variable	Value
advertise-map WORD<0-1536>	Specifies the route map name for route advertisements.
as-set	Enables autonomous system information. The default value is disable.
attribute-map WORD<0-1536>	Specifies the route map name.
WORD <1-256>	Specifies an IPv4 or an IPv6 address and its length in the form {a.b.c.d/len or ipv6 addr/len}.
summary-only	Enables the summarization of routes not included in routing updates. This variable creates the aggregate route and suppresses advertisements of more specific routes to all neighbors. The default value is disable.
suppress-map WORD<0-1536>	Specifies the route map name for the suppressed route list.

Use the data in the following table to use the `router bgp` command.

Variable	Value
<0-65535>	Specifies the AS number. You cannot enable BGP until you change the local AS to a value other than 0.
enable	Enables BGP on the router.

Configuring allowed IPv4 and IPv6 networks

Before you begin

- You must log on to BGP Router Configuration mode in ACLI to configure the GlobalRouter.
- You configure BGP on a VRF instance the same way you configure the GlobalRouter, except that you must use VRF Router Configuration mode and the prefix `ip bgp`. The VRF must have an RP Trigger of BGP. Not all parameters are configurable on non0 VRFs.

About this task

Configure network addresses to determine the network addresses that BGP advertises. The allowed addresses determine the BGP networks that originate from the Avaya Virtual Services Platform 9000.

Procedure

Specify IGP network prefixes for BGP to advertise:
`network <WORD 1-256> [metric <0-65535>]`

Example

Specify IGP network prefixes for BGP to advertise:

```
VSP-9012:1(router-bgp)#network 60:60:60:40::/64 metric 32
```

Variable definitions

Use the data in the following table to use the `network` command.

Table 8: Variable definitions

Variable	Value
WORD <1-256>	Specifies the IPv4 or the IPv6 network address and mask as {a.b.c.d/len or ipv6 addr/len}.
metric <0-65535>	Specifies the metric to use when the system sends an update for the routes in the network table. The metric configures the MED for the routes advertised to EBGP peers. The range is 0-65535.

Configuring BGP peers or peer groups

Before you begin

- If required, route policies exist.
- You must log on to BGP Router Configuration mode in ACLI to configure the GlobalRouter.
- You configure BGP on a VRF instance the same way you configure the GlobalRouter, except that you must use VRF Router Configuration mode and the prefix `ip bgp`. The VRF must have an RP Trigger of BGP. Not all parameters are configurable on non0 VRFs.

About this task

Configure peers and peer groups to simplify BGP configuration and makes updates more efficient.

BGP speakers can have many neighbors configured with similar update policies, for example, many neighbors use the same distribute lists, filter lists, outbound route maps, and update source. Group the neighbors that use the same update policies into peer groups and peer associations.

Many of the command variables in this procedure use default values. You can accept the default values or change them to customize the configuration.

Procedure

1. Create a peer or peer group:
`neighbor <nbr_ipaddr|peer-group-name>`
2. Apply a route policy to all incoming routes:
`neighbor <nbr_ipaddr|peer-group-name> in-route-map
WORD<0-256>`
3. Apply a route policy to all outgoing routes:
`neighbor <nbr_ipaddr|peer-group-name> out-route-map
WORD<0-256>`
4. Configure the source IP address:
`neighbor <nbr_ipaddr|peer-group-name> update-source
{A.B.C.D}`
5. Enable MD5 authentication:
`neighbor <nbr_ipaddr|peer-group-name> MD5-authentication
enable`
6. Specify an MD5 authentication password:
`neighbor password <nbr_ipaddr|peer-group-name> WORD<0-1536>`

7. Change the default values for other command variables as required.
8. Enable the configuration:


```
neighbor <nbr_ipaddr|peer-group-name> enable
```

Example

Create a peer or a peer group:

```
VSP-9012:1(router-bgp)#neighbor peergroupa
```

Apply a route policy (in-route-map or out-route-map) to all incoming or outgoing routes:

```
VSP-9012:1(router-bgp)#neighbor peergroupa in-route-map map1 out-
route-map map2
```

Configure the source IP address:

```
VSP-9012:1(router-bgp)# neighbor peergroupa update-source
47.10.17.31
```

Enable MD5 authentication:

```
VSP-9012:1(router-bgp)# neighbor peergroupa MD5-authentication
enable
```

Specify an MD5 authentication password:

```
VSP-9012:1(router-bgp)# neighbor password peergroupa password
```

Enable the configuration:

```
VSP-9012:1(router-bgp)# neighbor peergroupa enable
```

Variable definitions

Use the data in the following table to use the `neighbor` command.

Table 9: Variable definitions

Variable	Value
address-family <ipv6 vpn4>	Enables the VPNv4 address family on BGP neighbor. VSP9012:1(router-bgp)#neighbor peergroupa address-family ipv6
advertisement-interval <5-120>	Specifies the time interval, in seconds, that transpires between each transmission of an advertisement from a BGP neighbor. The default value is 5 seconds. VSP9012:1(router-bgp)#neighbor peergroupa advertisement-interval 26 enable

Variable	Value
default-ipv6-originate	Enables IPv6 BGP neighbor default originate. VSP9012:1(router-bgp)#neighbor peergroupa default-ipv6-originate
default-originate	Enables the switch to send a default route advertisement to the specified neighbor. A default route does not need to be in the routing table. The default value is disable. Do not use this command if default-information originate is globally enabled. VSP9012:1(router-bgp)#neighbor peergroupa default-originate enable peer-group test
ebgp-multihop	Enables a connection to a BGP peer that is more than one hop away from the local router. The default value is disable. VSP9012:1(router-bgp)#neighbor peergroupa ebgp-multihop retry-interval 3 timers 4 5
enable	Enables the BGP neighbor.
in-route-map WORD<0-256>	Applies a route policy rule to all incoming routes that are learned from, or sent to, the peers or peer groups of the local router. The local BGP router is the BGP router that allows or disallows routes and configures attributes in incoming or outgoing updates. WORD<0-256> is an alphanumeric string length (0–256 characters) that indicates the name of the route map or policy. VSP9012:1(router-bgp)#neighbor peergroupa in-route-map map1 address-family ipv6
ipv6-in-route-map WORD <0–256>	Creates IPv6 in route map. Word <0–256> Specifies the route map name in the range of 0 to 256 characters. VSP9012:1(router-bgp)#neighbor peergroupa ipv6-in-route-map map1
ipv6-out-route-map WORD <0–256>	Creates IPv6 out route map. Word <0–256> Specifies the route map name in the range of 0 to 256 characters. VSP9012:1(router-bgp)#neighbor peergroupa ipv6-out-route-map map2
max-prefix <0-2147483647>	Configures a limit on the number of routes that the router can accept from a neighbor. The default value is 12 000 routes. A value of 0 (zero) indicates that no limit exists.

Variable	Value
	<pre>VSP9012:1(router-bgp)#neighbor peergroupa max-prefix 158 in-route- map map1 out-route-map map2 address- family ipv6</pre>
MD5-authentication enable	<p>Enables TCP MD5 authentication between two peers. The default value is disable.</p>
neighbor-debug mask WORD<1-100>	<p>Displays specified debug information for a BGP peer. The default value is none.</p> <p><WORD 1-100> is a list of mask choices separated by commas with no space between choices. For example: {<mask>,<mask>,<mask>...}.</p> <p>Mask choices are:</p> <ul style="list-style-type: none"> none disables all debug messages. all enables all debug messages. error enables display of debug error messages. packet enables display of debug packet messages. event enables display of debug event messages. trace enables display of debug trace messages. warning enables display of debug warning messages. state enables display of debug state transition messages. init enables display of debug initialization messages. filter enables display of debug messages related to filtering. update enables display of debug messages related to sending and receiving updates. <pre>VSP9012:1(router-bgp)#neighbor peergroupa neighbor-debug-mask event,trace,warning,state</pre>
next-hop-self	<p>When enabled, specifies that the next-hop attribute in an IBGP update is the address of the local router or the router that generates the IBGP update. The default value is disable.</p> <p>You can only configure this variable if the neighbor is disabled.</p> <pre>VSP9012:1(router-bgp)#neighbor peergroupa next-hop-self out-route- map map2 peer-group peergroupb</pre>
<nbr_ipaddr peer-group-name>	<p>Specifies the peer IP address or the peer group name.</p>
out-route-map WORD<0-256>	<p>Applies a route policy rule to all outgoing routes that are learned from, or sent to, the peers or peer groups of the local router. The local BGP router is the BGP router that allows or disallows routes and configures attributes in incoming or outgoing updates.</p>

Variable	Value
	<i>WORD</i> <0-256> is an alphanumeric string length (0–256 characters) that indicates the name of the route map or policy.
peer-group <WORD 0-1536>	Adds a BGP peer to the specified subscriber group. You must create the specified subscriber group before you use this command.
remote-as <WORD 0-11>	Configures the remote AS number of a BGP peer or a peer-group. You must disable the admin-state before you can configure this variable. VSP9012:1(router-bgp)#neighbor peergroupa remote-as test
remove-private-as enable	Strips private AS numbers when an update is sent. This function is especially useful within a confederation. The default value is enable.
retry-interval <1-65535>	Configures the time interval, in seconds, for the ConnectRetry timer. The default value is 120 seconds. VSP9012:1(router-bgp)#neighbor peergroupa retry-interval 34
route-reflector-client	Configures the specified neighbor or group of neighbors as a route reflector client. The default value is disable. All configured neighbors become members of the client group and the remaining IBGP peers become members of the nonclient group for the local route reflector. VSP9012:1(router-bgp)#neighbor
route-refresh	Enables route refresh for the BGP peer. If enabled, a route refresh request received by a BGP speaker causes the speaker to resend all route updates it contains in its database that are eligible for the peer that issues the request.
send-community	Enables the switch to send the update message community attribute to the specified peer. The default value is disable.
soft-reconfiguration-in enable	Enables the router to relearn routes from the specified neighbor or group of neighbors without restarting the connection after the policy changes in the inbound direction. The default value is disable.
timers <0-21845> <0-65535>	Configures timers, in seconds, for the BGP speaker for this peer. <0-21845> is the keepalive time. The default is 60. Avaya recommends that you configure a value of 30 seconds. <0-65535> is the hold time. The default is 180. VSP9012:1(router-bgp)#neighbor peergroupa timers 4 6

Variable	Value
update-source {A.B.C.D}	Specifies the source IP address to use when the system sends BGP packets to this peer or peer group. You must disable the admin-state before you can configure this variable. VSP9012:1(router-bgp)#neighbor peergroupa update-source 47.17.10.32 weight 560
weight <0-65535>	Specifies the weight of a BGP peer or peer group, or the priority of updates the router can receive from that BGP peer. The default value is 0. If you have particular neighbors that you want to use for most of your traffic, you can assign a higher weight to all routes learned from that neighbor.

Use the data in the following table to use the `neighbor password` command.

Table 10: Variable definitions

Variable	Value
<nbr_ipaddr peer-group-name>	Specifies the peer IP address or the peer group name.
WORD<0-1536>	Specifies a password for TCP MD5 authentication between two peers.

Configuring a BGP peer or peer group password

Use this procedure to configure a BGP peer or peer group password for Transmission Control Protocol (TCP) MD5 authentication between two peers.

Before you begin

- Log on to the BGP Router Configuration mode in the ACLI.

Procedure

Assign a BGP peer or peer group password:

```
neighbor password <nbr_ipaddr|peer-group=name> WORD <0-1536>
```

Example

Assign a BGP peer or peer group password:

```
VSP-9012:1<router-bgp>#neighbor password peergroupa password1
```

Variable definitions

Use the data in the following table to use `neighbor password <nbr_ipaddr | peer-group-name> WORD <0-1536>` command.

Variable	Value
<code>password <nbr_ipaddr peer-group-name> WORD <0-1536></code>	<p>Specifies a password for TCP MD5 authentication between two peers.</p> <p><code>WORD <0-1536></code> is an alphanumeric string length from 0 to 1536 characters.</p> <p>To disable this option, use no operator with the command.</p> <p>To configure this option to the default value, use default operator with the command.</p>

Configuring redistribution to BGP

Before you begin

- If required, a route policy exists.
- You must log on to BGP Router Configuration mode in ACLI to configure the GlobalRouter.
- You configure BGP on a VRF instance the same way you configure the GlobalRouter, except that you must use VRF Router Configuration mode and the prefix `ip bgp`. The VRF must have an RP Trigger of BGP. Not all parameters are configurable on non0 VRFs.

About this task

Configure a redistribute entry to announce routes of a certain source protocol type into the BGP domain, for example, static, Routing Information Protocol (RIP), or direct routes. Use a route policy to control the redistribution of routes.

Procedure

1. Create a redistribution instance:


```
redistribute <direct|ospf|rip|static>
```
2. If required, specify a route policy to govern redistribution:


```
redistribute <direct|ospf|rip|static> route-map WORD<0-64>
[vrf-src WORD<0-16>]
```
3. If required, configure the route metric:

```
redistribute <direct|ospf|rip|static> metric <0-65535> [vrf-  
src WORD<0-16>]
```

4. Enable the instance:

```
redistribute <direct|ospf|rip|static> enable [vrf-src  
WORD<0-16>]
```

5. Exit BGP Router Configuration mode:

```
exit
```

6. Apply the redistribution instance configuration:

```
ip bgp apply redistribute <direct|ospf|rip|static> [vrf  
WORD<0-16>] [vrf-src <WORD 0-16>]
```

7. Apply BGP redistribution to a specific VRF:

```
ip bgp apply redistribute vrf WORD<0-16>
```

Example

Create a redistribution instance:

```
VSP-9012:1(router-bgp)#redistribute direct
```

If required, specify a route policy to govern redistribution:

```
VSP-9012:1(router-bgp)#redistribute direct route-map policy1 vrf-src  
source1
```

If required, configure the route metric:

```
VSP-9012:1(router-bgp)#redistribute direct metric 4 vrf-src source1
```

Enable the instance:

```
VSP-9012:1(router-bgp)#redistribute direct enable vrf-src source1
```

Exit BGP Router Configuration mode:

```
VSP-9012:1(router-bgp)#exit
```

Apply the redistribution instance configuration:

```
VSP-9012:1(config)#ip bgp apply redistribute direct vrf-src source1
```

Apply BGP redistribution to a specific VRF:

```
VSP-9012:1(config)#ip bgp apply redistribute vrf test
```

Variable definitions

Use the data in the following table to use the `redistribute` and `ip bgp apply redistribute` commands.

Table 11: Variable definitions

Variable	Value
<direct ipv6–direct ipv6–static ospf ospfv3 rip static>	Specifies the type of routes to redistribute (the protocol source).
enable	Enables the BGP route redistribution instance.
metric <0-65535>	Configures the metric to apply to redistributed routes.
route-map WORD<0-64>	Configures the route policy to apply to redistributed routes.
vrf WORD<0-16>	Specifies the name of a VRF instance.
vrf-src WORD<0-16>	Specifies the source VRF instance by name for route redistribution.

Configuring redistribution to BGP+ for VRF 0

Configure an IPv6 redistribute entry to announce IPv6 routes of a certain source protocol type into the BGP domain, for example, static, OSPF, or direct routes. Use a route policy to control the redistribution of routes.

Before you begin

- If required, a route policy exists.
- Access BGP Router Configuration mode.

Procedure

1. Create a redistribution instance:

```
redistribute <ipv6-direct|ipv6-static|ospfv3>
```
2. If required, specify a route policy to govern redistribution:

```
redistribute <ipv6-direct|ipv6-static|ospfv3> route-map WORD  
<0-64>
```
3. If required, configure a route metric:

```
redistribute <ipv6-direct|ipv6-static|ospfv3> metric <0-  
65535>
```
4. Enable the instance:

```
redistribute <ipv6-direct|ipv6-static|ospfv3> enable
```

Unlike IPv4 redistribution, you do not need to manually apply the IPv6 redistribution instance. Once you enable the IPv6 redistribution instance, it is automatically applied.

Example

Specify a route policy to govern redistribution by using the following command:

```
VSP-9012:1(router-bgp)#redistribute ipv6-direct route-map policy2
```

Variable definitions

Use the information in the following table to help you use the `redistribute <ipv6-direct | ipv6-static | ospfv3>` command.

Table 12: Variable definitions

Variable	Value
enable	Enables the BGP route redistribution instance. The default value is none. To configure this option to the default value, use default operator with the command. To disable this option, use no operator with the command.
metric<0-65535>	Configures the metric to apply to redistributed routes. The default value is 0. To configure this option to the default value, use default operator with the command.
route-map <Word 0-64>	Configures the route policy to apply to redistributed routes. The default value is none. To configure this option to the default value, use default operator with the command.

Job aid

Use the data in the following table to know how route policies are used for BGP from IPv6 perspective:

	REDISTRIBUTE			ACCEPT	ANNOUNCE
	IPv6 Direct	IPv6 Static	OSPFv3	BGP	BGP
MATCH					
as-path				Yes	Yes

	REDISTRIBUTE			ACCEPT	ANNOUNCE
	IPv6 Direct	IPv6 Static	OSPFv3	BGP	BGP
community	Yes	Yes	Yes	Yes	Yes
community-exact				Yes	Yes
extcommunity				Yes	Yes
interface					
local-preference					
metric	Yes	Yes	Yes	Yes	Yes
network	Yes	Yes	Yes	Yes	Yes
next-hop		Yes	Yes	Yes	Yes
protocol					
route-source				Yes	
route-type			Yes		Yes
tag					
vrf					
vrfids					
SET					
as-path				Yes	Yes
as-path-mode				Yes	Yes
automatic-tag					
community				Yes	Yes
community-mode				Yes	Yes
injectlist	Yes	Yes	Yes		
ip-preference					
local-preference				Yes	Yes
mask					
metric	Yes	Yes	Yes	Yes	Yes
metric-type					
metric-type-internal					
next-hop				Yes	Yes
nssa-pbit					
origin					Yes

	REDISTRIBUTE			ACCEPT	ANNOUNCE
	IPv6 Direct	IPv6 Static	OSPFv3	BGP	BGP
origin-egp-as					
tag					
weight				Yes	

Configuring prefix lists

Before you begin

- You must log on to the Global Configuration mode in ACLI.
- You configure prefix lists on a VRF instance the same way you configure the Global Router, except that you must use VRF Router Configuration mode.

Important:

When you configure a prefix list for a route policy, add the prefix as a.b.c.d/32. You must enter the full 32-bit mask to exact a full match of a specific IP address.

About this task

Configure a prefix list to allow or deny specific route updates. A prefix list policy specifies route prefixes to match. After a match occurs, the system uses the route.

The prefix list contains a set of contiguous or noncontiguous routes. Reference prefix lists by name from within a routing policy.

Procedure

1. Configure a prefix list:

```
ip prefix-list WORD<1-64> {A.B.C.D/X} [ge <0-32>] [le <0-32>]
```
2. Optionally, rename an existing prefix list:

```
ip prefix-list WORD<1-64> name WORD<1-64>
```
3. Display the prefix list:

```
show ip prefix-list [prefix {A.B.C.D}] [vrf WORD<0-16>]
[vrfids WORD<0-512>] [WORD <1-64>]
```

Example

Configure a prefix list:

```
VSP-9012:1(config)#ip prefix-list list1 15.15.15.34/32 ge 32 le 32
```

Optionally, rename an existing prefix list:

```
VSP-9012:1(config)#ip prefix-list list1 name list2
```

Variable definitions

Use the data in the following table to use the `ip prefix-list` command.

Table 13: Variable definitions

Variable	Value
{A.B.C.D/X}	Specifies the IP address and the mask in one of the following formats: <ul style="list-style-type: none"> • a.b.c.d/x • a.b.c.d/x.x.x.x • default
ge <0–32>	Specifies the minimum length to match. Lower bound and higher bound mask lengths together can define a range of networks.
le <0–32>	Specifies the maximum length to match. Lower bound and higher bound mask lengths together can define a range of networks.
name WORD<1-64>	Renames the specified prefix list. The name length is 1–64 characters.
WORD<1-64>	Specifies the name for a new prefix list.

Use the data in the following table to use the `show ip prefix-list` command.

Table 14: Variable definitions

Variable	Value
{A.B.C.D}	Specifies the prefix to include in the command output.
vrf WORD<0-16>	Specifies the name of the VRF.
vrfids WORD<0-512>	Specifies the ID of the VRF and is an integer in the range of 0–512.
WORD<1-64>	Specifies a prefix list, by name, to use for the command output.

Use the following table to use the `show ip prefix-list` command output.

Table 15: Variable definitions

Variable	Value
PREFIX	Indicates the member of a specific prefix list.
MASKLEN	Indicates the prefix mask length in bits.
FROM	Indicates the prefix mask starting point in bits.
TO	Indicates the prefix mask endpoint in bits.

Configuring an IPv6 prefix list

Use IPv6 prefix lists to allow or deny specific IPv6 route updates. A prefix list policy specifies route prefixes to match. When there is a match, the route is used.

Before you begin

About this task

Access Global Configuration mode.

Procedure

1. Create an IPv6 prefix list:

```
ipv6 prefix-list <WORD 1-64> <WORD 1-256> [<ge|le> <0-128>]
```

Use the same command to add additional prefixes to the list.

2. To rename the list:

```
ipv6 prefix-list <WORD 1-64> name <WORD 1-64>
```

Example

Create an IPv6 prefix list:

```
VSP-9012:1<config>#ipv6 prefix-list list4 4717:0:0:0:0:0:7933:6/64
ge 32 le 64
```

To rename the list:

```
VSP-9012:1<config>#ipv6 prefix-list list4 name list5
```

Variable definitions

Use the data in the following table to use the ipv6 prefix-list command..

Table 16: Variable definitions

Variable	Value
<WORD 1–256> [<ge le> <0–128>]	Creates or adds a prefix to the list. The default value is none. <ul style="list-style-type: none"> • <WORD 1–256> specifies the IP prefix and length. • <ge le> specifies greater than or equal to or less than or equal to. • <0–128> specifies the mask length in the range 0 to 128. To disable this option, use no operator with the command
name <WORD 1–64>	Names the prefix list. The default value is none.

Configuring AS path lists

Before you begin

- You must log on to BGP Router Configuration mode in ACLI to configure BGP for the GlobalRouter.
- You configure BGP on a VRF instance the same way you configure the GlobalRouter, except that you must use VRF Router Configuration mode and the prefix `ip bgp`. The VRF must have an RP Trigger of BGP. Not all parameters are configurable on non0 VRFs.

About this task

Configure an AS path list to restrict the routing information a router learns or advertises to and from a neighbor. The AS path list acts as a filter that matches AS paths.

Procedure

Create the path list:

```
ip as-list <1-1024> memberid <0-65535> <permit|deny> as-path
WORD<0-1536>
```

Use this command for each member by specifying different member IDs.

Example

Create the path list:

```
VSP-9012:1(config)#ip as-list 234 memberid 3456 permit as-path "5"
```

Variable definitions

Use the data in the following table to use the `ip as-list` command.

Table 17: Variable definitions

Variable	Value
<0-65535>	Specifies an integer value between 0–65535 that represents the regular expression entry in the AS path list.
<1-1024>	Specifies an integer value from 1–1024 that represents the AS-path list ID you want to create or modify.
<permit deny>	Permits or denies access for matching conditions.
WORD<0–1536>	Specifies the AS number as an integer value between 0–1536. Place multiple AS numbers within quotation marks ("").

Configuring community lists

Before you begin

- You must log on to BGP Router Configuration mode in ACLI to configure BGP for the GlobalRouter.
- You configure BGP on a VRF instance the same way you configure the GlobalRouter, except that you must use VRF Router Configuration mode and the prefix `ip bgp`. The VRF must have an RP Trigger of BGP. Not all parameters are configurable on non0 VRFs.

About this task

Configure community lists to specify permitted routes by using their BGP community. This list acts as a filter that matches communities or AS numbers.

Procedure

Create a community list:


```
ip community-list <1-1024> memberid <0-65535> <permit|deny>
community-string WORD<0-256>
```

Example

Create a community list:

```
VSP-9012:1(config)#ip community-list 1 memberid 4551 permit
community-string internet
```

Variable definitions

Use the data in the following table to use the `ip community-list` command.

Table 18: Variable definitions

Variable	Value
<0-65535>	Specifies an integer value from 0–65535 that represents the member ID in the community list.
<1-1024>	Specifies an integer value from 1–1024 that represents the community list ID.
<permit deny>	Configures the access mode, which permits or denies access for matching conditions.
WORD<0-256>	Specifies the community as an alphanumeric string value with a string length from 0–256 characters. Enter this value in one of the following formats: <ul style="list-style-type: none"> • (AS num:community-value) • (well-known community string) Well known communities include: internet, no-export, no-advertise, local-as (known as NO_EXPORT_SUBCONFED).

Configuring extended community lists

Before you begin

- You must log on to BGP Router Configuration mode in ACLI to configure BGP for the GlobalRouter.
- You configure BGP on a VRF instance the same way you configure the GlobalRouter, except that you must use VRF Router Configuration mode and the prefix `ip bgp`. The

VRF must have an RP Trigger of BGP. Not all parameters are configurable on non0 VRFs.

About this task

Configure community lists to specify permitted routes by BGP extended community attributes, including route targets and sites of origin (SOO). This list acts as a filter that matches route targets and SOO.

Procedure

1. Create an extended community list based on the route target attribute:

```
ip extcommunity-list <1-1024> memberId <0-65535> rt {<0-65535> <0-2147483647>|<A.B.C.D> <0-65535>} [soo {<0-65535> <0-2147483647>|<A.B.C.D> <0-65535>}]
```

You can optionally configure the SOO attributes at the end of the same command or you can configure the SOO separately using the syntax in the following step.

2. Create an extended community list based on the SOO attribute:

```
ip extcommunity-list <1-1024> memberId <0-65535> soo {<0-65535> <0-2147483647>|<A.B.C.D> <0-65535>}
```

Example

Create an extended community list based on the route target attribute:

```
vsp-9012:1(config)#ip extcommunity-list 1 memberId 234 rt 16.16.16.34
5 soo 32 45
```

Variable Definitions

Use the data in the following table to use the `ip extcommunity-list` command.

Table 19: Variable definitions

Variable	Value
<1-1024>	Specifies an integer value from 1–1024 that represents the community list ID you want to create or modify.
memberId <0-65535>	Specifies an integer value from 0–65535 that represents the member ID in the community list.
rt <0-65536> <0-2147483647> rt <A.B.C.D> <0-65535>	Specifies the route target in the format {AS number:assigned number} (that is, {0–65535}:{0–2147483647}) or {ipaddress:assigned number} (that is, {a.b.c.d}:{0–65535}).
soo <0-65535> <0-2147483647> soo <A.B.C.D> <0-65535>	Specifies the site of origin in the format {AS number:assigned number} (that is, {0–65535}:{0–

Variable	Value
	2147483647}) or {ipaddress:assigned number} (that is, {a.b.c.d}:{0-65535}).

Chapter 5: BGP verification using ACLI

About this task

Use show commands to verify Border Gateway Protocol (BGP) configuration and to monitor or troubleshoot BGP operation.

Viewing BGP aggregate information

About this task

Display information about current aggregate addresses.

Procedure

Display information about current aggregates:

```
show ip bgp aggregates [<prefix/len>] [vrf WORD <0-16>] [vrfids  
WORD<0-255>]
```

Variable definitions

Use the data in the following table to use the `show ip bgp aggregates` command.

Table 20: Variable definitions

Variable	Value
<prefix/len>	Specifies the IP address and the mask length (the length can be 0–32).
vrf WORD<0–16>	Specifies a VRF instance by name (the string length ranges from 0–16 characters).
vrfids WORD<0–255>	Specifies a range of VRFs by ID number (the ID ranges from 0–512).

Viewing IPv6 BGP+ aggregate information

Display information about current IPv6 aggregate addresses.

About this task

Use BGP 4 byte AS numbers to ensure the continuity of loop-free inter-domain routing information between ASs and to control the flow of BGP updates as 2 byte AS numbers will deplete soon.

Procedure

Display information about current IPv6 aggregates:

```
show bgp ipv6 aggregates [<WORD 1-256>]
```

Variable definitions

Use the data in the following table to use the `show bgp ipv6 aggregates` command.

Variable	Value
WORD <1-256>	Specifies the IPv6 prefix and the prefix length (the length can be 0 to 128).

Viewing CIDR routes

About this task

Display information about classless interdomain routing (CIDR) routes.

Procedure

Display information about CIDR routes:

```
show ip bgp cidr-only [<prefix/len>] [vrf WORD<0-16>] [vrfs  
WORD<0-512>]
```

Variable definitions

Use the data in the following table to use the `show ip bgp cidr-only` command.

Table 21: Variable definitions

Variable	Value
<prefix/len>	Specifies an exact match of the prefix. This variable is an IP address and an integer value from 0–32 in the format a.b.c.d/xx.
vrf WORD<0–16>	Specifies a VRF instance by name (the string length ranges from 0–16 characters).
vrfids WORD<0–512>	Specifies a range of VRFs by ID number (the ID ranges from 0–512).

Job aid

Use the data in the following table to understand the `show ip bgp cidr-only` command output.

Table 22: show ip bgp cidr-only field descriptions

Field	Description
NETWORK/MASK	Specifies the network IP address and exact mask length (must be an integer value from 0–32).
PEER REM ADDR	Specifies the IP address of the remote peer.
NEXTHOP ADDRESS	Specifies the IP address of the next hop.
ORG	Specifies the source of a route: <ul style="list-style-type: none"> • IGP — the route is interior to the originating AS that inserts this route into the BGP table (0 = IGP). • EGP — the route is learned through an Exterior Gateway Protocol (EGP) before it is inserted into the BGP table (1 = BGP). • Incomplete — the origin of the route is unknown or learned by some other means. For example, the router learns these routes through RIP, OSPF, or static routes (2 = Incomplete).
LOC PREF	Specifies the local preference.

Viewing BGP configuration

About this task

View information about the BGP configuration.

Procedure

Display information about the current BGP configuration:

```
show ip bgp conf [vrf WORD<0-16>] [vrfrids WORD<0-512>]
```

Example

```
VSP-9012:1(config)#show ip bgp conf
```

```

=====
                        BGP Configuration - GlobalRouter
=====
                        BGP version - 4
                          local-as - 2147483647
                          Identifier - 1.100.10.0
                          BGP on/off - OFF
                          as-4-byte - enable
                          as-dot - disable
                          aggregation - enable
                          always-cmp-med - disable
                          auto-peer-restart - enable
                          auto-summary - enable
                          comp-bestpath-med-confed - disable
                          default-local-preference - 100
                          default-metric - -1
                          deterministic-med - disable
                          flap-dampening - disable
                          debug-screen - Off
                          global-debug - none
                          ibgp-report-import-rt - enable
                          ignore-illegal-rtrid - enable
                          max-equalcost-routes - 1
                          no-med-path-is-worst - enable
                          route-refresh - disable
                          orig-def-route - disable
                          quick-start - disable
                          synchronization - enable
                          vrfId - 0

route-reflection config state - enable
  route-reflection oper state - disable
    cluster-id - 0.0.0.0
    cl-to-cl-reflection - enable
    decision state - Idle
  confederation identifier - 0
    traps - disable

```

Variable definitions

Use the data in the following table to use the `show ip bgp conf` command.

Table 23: Variable definitions

Variable	Value
vrf WORD<0–16>	Specifies a VRF instance by name (the string length ranges from 0–16 characters).
vrfids WORD<0–512>	Specifies a range of VRFs by ID number (the ID ranges from 0–512).

Viewing BGP confederation

About this task

Display information about BGP confederations.

Procedure

Display information about current BGP confederations:

```
show ip bgp confederation
```

Example

```
CB-SWA:1(config)#show ip bgp confederation
confederation identifier 0
confederation peer as
```

Viewing flap-dampened routes

About this task

Display information about flap-dampened routes to determine unreliable routes.

Procedure

Display information about flap-dampened routes:

```
show ip bgp dampened-paths {A.B.C.D} [<prefix/len>] [longer-
prefixes] [vrf WORD<0-16>] [vrfids WORD<0-512>]
```

Variable definitions

Use the data in the following table to use the `show ip bgp dampened-paths` command.

Table 24: Variable definitions

Variable	Value
{A.B.C.D}	Specifies the source IP address in the format a.b.c.d.
longer-prefixes	Shows long prefixes. The longer-prefixes indicate the mask length from a specified prefix to 32 (for example, show from prefix a.b.c.d/len to a.b.c./32).
<prefix/len>	Shows paths with this prefix. The prefix is the IP address and exact mask length (must be an integer value from 0–32).
vrf <value>	Specifies a VRF instance by name (the string length ranges from 0–16 characters).
vrfids <value>	Specifies a range of VRFs by ID number (the ID ranges from 0–512).

Job aid

Use the data in the following table to understand the `show ip bgp dampened-paths` command output.

Table 25: show ip bgp dampened-paths field descriptions

Field	Description
NETWORK/MASK	Specifies the network IP address and exact mask length (must be an integer value from 0–32).
PEER REM ADDR	Specifies the IP address of the remote peer.
NEXTHOP ADDRESS	Specifies the IP address of the next hop.
ORG	Specifies the source of a route:

Field	Description
	<ul style="list-style-type: none"> • IGP — the route is interior to the originating AS that inserts this route into the BGP table (0 = IGP). • EGP — the route is learned through an Exterior Gateway Protocol (EGP) before it is inserted into the BGP table (1 = BGP). • Incomplete — the origin of the route is unknown or learned by some other means. For example, the router learns these routes through RIP, OSPF, or static routes (2 = Incomplete).
LOC PREF	Specifies the local preference.

Viewing global flap-dampening configurations

About this task

Display global information about flap-dampening.

Procedure

Display global information about flap-dampening:

```
show ip bgp flap-damp-config [prefix/len] [vrf WORD<0-16>]
[vrfids WORD<0-512>]
```

Variable definitions

Use the data in the following table to use the `show ip bgp flap-damp-config` command.

Table 26: Variable definitions

Variable	Value
<prefix/len>	Shows paths with this prefix. The prefix is the IP address and exact mask length (must be an integer value from 0–32).
vrf WORD<0–16>	Specifies a VRF instance by name (the string length ranges from 0–16 characters).
vrfids WORD<0–512>	Specifies a range of VRFs by ID number (the ID ranges from 0–512).

Job aid

Use the data in the following table to understand the `show ip bgp flap-damp-config` command output.

Table 27: show ip bgp flap-damp-config field descriptions

Field	Description
Status	Indicates the global state of the route flap dampening feature. Valid values are enable or disable.
PolicyName	This field does not apply for this release.
CutoffThreshold	Indicates the penalty level that causes route suppression.
ReuseThreshold	Specifies the system-configured time for route reuse.
Decay	Indicates the decay rate based on the decay algorithm.
MaxHoldDown	Indicates the maximum length of time, in seconds, to suppress the route.

Viewing imported routes

About this task

Display information about BGP imported routes.

Procedure

Display information about BGP imported routes:

```
show ip bgp imported-routes [<prefix/len>] [longer-prefixes]
[vrf WORD<0-16>] [vrfids WORD<0-512>]
```

Variable definitions

Use the data in the following table to use the `show ip bgp imported-routes` command.

Table 28: Variable definitions

Variable	Value
longer-prefixes	Shows long prefixes. The longer-prefixes indicate the mask length from a specified prefix to 32 (for example, show from prefix a.b.c.d/len to a.b.c./32).
<prefix/len>	Shows paths with this prefix. The prefix is the IP address and exact mask length (must be an integer value from 0–32).
vrf WORD<0–16>	Specifies a VRF instance by name (the string length ranges from 0–16 characters).
vrfids WORD<0–512>	Specifies a range of VRFs by ID number (the ID ranges from 0–512).

Job aid

Use the data in the following table to understand the `show ip bgp imported-routes` command output.

Table 29: show ip bgp imported-routes field descriptions

Field	Description
ROUTE	Specifies the IP address of the route.
METRIC	Specifies the route metric.
COMMUNITY	Specifies the BGP community.
LOCALPREF	Specifies the local preference.
NEXTHOP	Specifies the IP address of the next hop.
ORIGIN	Specifies the source of a route: <ul style="list-style-type: none"> • IGP — the route is interior to the originating AS that inserts this route into the BGP table (0 = IGP). • EGP — the route is learned through an Exterior Gateway Protocol (EGP) before it is inserted into the BGP table (1 = BGP). • Incomplete — the origin of the route is unknown or learned by some other means. For example, the router learns these routes through RIP, OSPF, or static routes (2 = Incomplete).

Viewing BGP neighbors information

About this task

Display information about BGP neighbors.

Procedure

1. Display information about BGP neighbors:

```
show ip bgp neighbors [{A.B.C.D}] [vrf WORD<0-16>] [vrfids WORD<0-512>]
```
2. Display information about BGP peer advertised routes:

```
show ip bgp neighbors {A.B.C.D} advertised-routes [<prefix/len>] [longer-prefixes] [vrf WORD<0-16>] [vrfids WORD<0-512>]
```
3. Display information about BGP peer routes:

```
show ip bgp neighbors {A.B.C.D} routes [<prefix/len>] [community <enable|disable>] [longer-prefixes] [vrf WORD<0-16>] [vrfids WORD<0-512>]
```
4. Display statistics for BGP peers:

```
show ip bgp neighbors {A.B.C.D} stats [vrf WORD<0-16>] [vrfids WORD<0-512>]
```

Example

```
VSP-9012:1(router-bgp)#show ip bgp neighbors

=====
BGP Neighbor Info - GlobalRouter
=====

BGP neighbor is 200.200.200.63 remote AS 63, Internal Peer, MP-BGP-capable, BGP
state [Established] remote router ID 63.1.1.1

      vrf instance - 0
      admin-state - BGP ON
      connect-retry-interval - 120
      ebgp-multihop - disable
      hold-time - 30
      keepalive-time - 10
      hold-time-configured - 180
      keepalive-time-configured - 60
      max-prefix - 12000
      nexthop-self - disable
      originate-def-route - disable
      originate-v6-def-route - disable
      MD5-authentication - disable
      neighbor-debug - all
      remove-private-as - disable
      route-advertisement-interval - 5
```

```

route-reflector-client - disable
  send-community - disable
soft-reconfiguration-in - disable
  updt-source-interface - 0.0.0.0
    weight - 100
    Route Policy In -
    Route Policy Out -
    IPv6Route Policy In -
    IPv6Route Policy Out -
  address-family ipv6 - enable
  address-family vpnv4 - disable
    route-refresh - disable
negotiated-session-capabilites - IPv6
  Total bgp neighbors - 1

```

Variable definitions

Use the data in the following table to use the `show ip bgp neighbors` command.

Table 30: Variable definitions

Variable	Value
{A.B.C.D}	Specifies the IP address.
community <enable disable>	Enables or disables the display of community attributes.
longer-prefixes	Shows long prefixes. The longer-prefixes indicate the mask length from a specified prefix to 32 (for example, show from prefix a.b.c.d/len to a.b.c./32).
prefix/len	Shows paths with this prefix. The prefix is the IP address and exact mask length (must be an integer value from 0–32).
WORD<0–16>	Specifies a VRF instance by name.
vrfids WORD<0–512>	Specifies a range of VRFs by ID number.

Viewing BGP network configurations

About this task

Display information about BGP network configurations.

Procedure

Display information about BGP network configurations:

```
show ip bgp networks [<prefix/len>] [vrf WORD<0-16>] [vrfids
WORD<0-255>]
```

Variable definitions

Use the data in the following table to use the `show ip bgp networks` command.

Table 31: Variable definitions

Variable	Value
<prefix/len>	Shows paths with this prefix. The prefix is the IP address and exact mask length (must be an integer value from 0–32).
vrf WORD<0–16>	Specifies a VRF instance by name (the string length ranges from 0–16 characters).
vrfids WORD<0–512>	Specifies a range of VRFs by ID number (the ID ranges from 0–512).

Viewing IPv6 BGP+ network configurations

Display information about BGP+ network configurations.

Procedure

Display information about BGP+ network configurations:

```
show bgp ipv6 networks <WORD 1-256>
```

Variable definitions

Use the data in the following table to use the `show bgp ipv6 networks` command.

Table 32: Variable definitions

Variable	Value
<WORD 1–256>	Shows networks with this prefix. The prefix is the IP address and exact mask length (must be an integer value between 0 and 32).

Viewing BGP peer group information

About this task

Display information about BGP peer groups.

Procedure

Display information about BGP peer groups:

```
show ip bgp peer-group [WORD<0-1536>] [vrf WORD<0-16>] [vrfids
WORD<0-512>]
```

Variable definitions

Use the data in the following table to use the `show ip bgp peer-group` command.

Table 33: Variable definitions

Variable	Value
vrf WORD<0-16>	Specifies a VRF instance by name (the string length ranges from 0-16 characters).
vrfids WORD<0-512>	Specifies a range of VRFs by ID number (the ID ranges from 0-512).
WORD<0-1536>	Specifies the name of the peer group (the string length ranges from 0-1536 characters).

Viewing BGP redistributed routes

About this task

Display information about BGP redistributed routes.

Procedure

Display information about BGP redistributed routes:

```
show ip bgp redistributed-routes [<prefix/len>] [vrf WORD<0-
16>] [vrfids WORD<0-512>]
```

Variable definitions

Use the data in the following table to use the `show ip bgp redistributed-routes` command.

Table 34: Variable definitions

Variable	Value
<prefix/len>	Shows paths with this prefix. The prefix is the IP address and exact mask length (must be an integer value from 0–32).
vrf WORD<0–16>	Specifies a VRF instance by name (the string length ranges from 0–16 characters).
vrfids WORD<0–512>	Specifies a range of VRFs by ID number (the ID ranges from 0–255).

Job aid

Use the data in the following table to understand the `show ip bgp redistributed-routes` command output.

Table 35: show ip bgp redistributed-routes field descriptions

Field	Description
SRC-VRF	Indicates the redistribution source VRF instance.
SRC	Indicates the redistribution source: RIP, Local, Static, or OSPF.
MET	Indicates the metric value
ENABLE	Indicates whether the redistribution policy is enabled (T) true or disabled (F) false.
RPOLICY	The route policy currently assigned to the redistribution.

Viewing a summary of BGP configurations

About this task

Display summarized information about BGP.

Procedure

Display summarized information about BGP:

```
show ip bgp summary [vrf WORD<0-16>] [vrfids WORD<0-512>]
```

Example

```
VSP-9012:1(config)#show ip bgp summary

=====
BGP Summary - GlobalRouter
=====
                BGP version - 4
                local-as - 63
                Identifier - 45.1.1.1
                Decision state - Idle
                The total number of routes is 9

BGP NEIGHBOR INFO :

  NEIGHBOR      RMTAS      STATE      HLDTM  KPALV  HLDCFG  KPCFG   WGHT  CONRTY  ADVINT
-----
200.200.200.63  63          Established  30     10     180     60     100    120    5
Total bgp neighbors: 1

BGP CONFEDERATION INFO :
confederation identifier 0
confederation peer as

BGP NETWORK INFO :

=====
BGP Networks - GlobalRouter
=====
```

Variable definitions

Use the data in the following table to use the `show ip bgp summary` command.

Table 36: Variable definitions

Variable	Value
vrf WORD<0-16>	Specifies a VRF instance by name.
vrfids WORD<0-512>	Specifies a range of VRFs by ID number.

Job aid

Use the data in the following table to understand the `show ip bgp summary` command output.

Table 37: Variable definitions

Field	Description
BGP version	Specifies the version of BGP that runs on the router.
local-as	Specifies the local autonomous system number.
Identifier	Specifies the BGP identifier.
Decision state	Specifies the BGP process state.
NEIGHBOR	Specifies the IP address of the remote peer.
RMTAS	Specifies the AS number of the remote peer.
STATE	Specifies the peer operating state: Idle, Accept, Connect, Open, Open-sent, and Established.
HLDTM	Specifies the negotiated hold time timer.
KPALV	Specifies the keep alive timer.
HLDCFG	Specifies the configured hold time timer.
KPCFG	Specifies the configured keep alive timer.
WGHT	Specifies the weight value assigned to the peer.
CONRTY	Specifies the retry timer.
ADVINT	Specifies the advertisement interval.

Viewing BGP routes

About this task

Display information about BGP routes.

Procedure

Display information about BGP routes:

```
show ip bgp route [<prefix/len>] [community <enable|disable>]
[ip {A.B.C.D}] [longer-prefixes] [vrf WORD<0-16>] [vrfids
WORD<0-512>]
```

Variable definitions

Use the data in the following table to use the `show ip bgp route` command.

Table 38: Variable definitions

Variable	Value
community <enable disable>	Enables or disables the display of community attributes.
ip {A.B.C.D}	Specifies an IP address.
longer-prefixes	Shows long prefixes. Longer-prefixes indicates the mask length from a specified prefix to 32 (for example, show from prefix a.b.c.d/len to a.b.c./32).
<prefix/len>	Shows paths with this prefix. The prefix is the IP address and exact mask length (must be an integer value from 0–32).
vrf WORD<0–16>	Specifies a VRF instance by name.
vrfids WORD<0–512>	Specifies a range of VRFs by ID number.

Job aid

Use the data in the following table to understand the `show ip bgp route` command output.

Table 39: show ip bgp route

Field	Description
NETWORK/MASK	Specifies the path prefix address.
PEER REM ADDR	Specifies the remote peer address.
NEXTHOP ADDRESS	Specifies the BGP next hop address.
ORG	Specifies the source of a route:

Field	Description
	<ul style="list-style-type: none"> • IGP — the route is interior to the originating AS that inserts this route into the BGP table (0 = IGP). • EGP — the route is learned through an Exterior Gateway Protocol (EGP) before it is inserted into the BGP table (1 = BGP). • Incomplete — the origin of the route is unknown or learned by some other means. For example, the router learns these routes through RIP, OSPF, or static routes (2 = Incomplete).
LOCAL PREF	Specifies the local preference.

Viewing IPv6 BGP+ routes

Display information about BGP routes.

Procedure

Display information about BGP routes:

```
show bgp ipv6 route [<WORD 1-256>] [longer-prefixes] [community
<enable|disable>] [ipv6 <ipv6-addr>]
```

Variable definitions

Use the data in the following table to use the `show ip bgp route` command.

Table 40: Variable definitions

Variable	Value
[<WORD 1-256>]	Shows paths with this prefix. The prefix is the IP address and exact mask length (must be an integer value between 0 and 32).
community <enable disable>	Enables or disables the display of community attributes.
ipv6 <ipv6-addr>]	Specifies an IPv6 address.
longer-prefixes	Shows long prefixes. The longer-prefixes indicate the mask length from any specified prefix to 128

Variable	Value
	(for example, show from prefix X:X::X:X/len to X:X::X:X/ 128).

Chapter 6: BGP configuration using EDM

About this task

Configure Border Gateway Protocol (BGP) to create an interdomain routing system that guarantees loop-free routing information between autonomous systems.

For information about how to configure route policies, see *Avaya Virtual Services Platform 9000 Configuration — IP Routing, NN46250-505*.

Configuring BGP globally

Before you begin

- Change the VRF instance as required to configure BGP on a specific VRF instance. The VRF must have an RP trigger of BGP. Not all parameters are configurable on non0 VRFs.

About this task

Enable BGP so that BGP runs on the router. Configure general BGP parameters to define how BGP operates on the system.

If you must configure the BGP router ID, use ACLI. You cannot configure the BGP router ID using EDM.

Procedure

1. In the navigation tree, open the following folders: **Configuration > IP**.
2. Click **BGP**.
3. Click the **Generals** tab.
4. In AdminStatus, select **enable**.
5. Configure the local autonomous system (AS) ID.
6. In the **Aggregate** area, enable or disable route aggregation as required.
7. Configure the BGP options as required.
8. In the **DebugMask** area, select the check box for the type of information to show for BGP debugging purposes.
9. Configure BGP confederations as required.
10. Configure BGP route reflectors as required.

11. Click **Apply**.

Generals field descriptions

Use the data in the following table to use the **Generals** tab.

Name	Description
bgpVersion	Specifies the version of BGP that operates on the router.
bgpIdentifier	Specifies the BGP router ID number.
AdminStatus	Enables or disables BGP on the router. The default is disable. You cannot enable AdminStatus until you change the LocalAS value to a nonzero value.
4ByteAs	Enables or disables 4-byte AS numbers. The default is disable.
LocalAs	Configures the local AS number in the range of 0 – 65535. You cannot change the LocalAS if AdminStatus is enable.
AsDot	Enables or disable the AS dot notation format for the 4-byte AS number. The default is disable. The AS dot notation is easier to read and remember than the AS plain notation, but it can be difficult to convert from AS plain to AS dot. The IETF prefers the AS plain notation.
Aggregate	Enables or disables aggregation. The default is enable.
DefaultMetric	Configures the metric sent to BGP neighbors. The default metric determines the cost of a route a neighbor uses. Use this parameter in conjunction with the redistribute parameters so that BGP uses the same metric for all redistributed routes. The default is -1. The range is -1–2147483647.
DefaultLocalPreference	Specifies the default local preference. The local preference indicates the preference that AS border routers assign to a chosen route when they advertise it to IBGP peers. The default is 100. The range is 0–2147483647.
DefaultInformationOriginate	Enables or disables the redistribution of network 0.0.0.0 into BGP. The default is disable.
AlwaysCompareMed	Enables or disables the comparison of the multi-exit discriminator (MED) parameter for paths from

Name	Description
	neighbors in different ASs. The system prefers a path with a lower MED over a path with a higher MED. The default is disable.
DeterministicMed	Enables or disables deterministic MED. Deterministic MED compares the MEDs after routes advertised by different peers in the same AS are chosen. The default is disable.
AutoPeerRestart	Enables or disables the process that automatically restarts a connection to a BGP neighbor. The default is enable.
AutoSummary	Enables or disables automatic summarization. If you enable this variable, BGP summarizes networks based on class limits (for example, Class A, B, or C networks). The default is enable.
NoMedPathsWorst	Enables or disables NoMedPathsWorst. If you enable this variable, BGP treats an update without a MED attribute as the worst path. The default is enabled.
BestPathMedConfed	Enables or disables the comparison of MED attributes within a confederation. The default is disable.
DebugMask	<p>Displays the specified debug information for BGP global configurations. The default value is none. Other options are</p> <ul style="list-style-type: none"> • none disables all debug messages. • event enables the display of debug event messages. • state enables display of debug state transition messages. • update enables display of debug messages related to updates transmission and reception. • error enables the display of debug error messages. • trace enables the display of debug trace messages. • init enables the display of debug initialization messages. • all enables all debug messages. • packet enables the display of debug packet messages.

Name	Description
	<ul style="list-style-type: none"> • warning enables the display of debug warning messages. • filter enables the display of debug messages related to filtering.
IgnoreIllegalRouterId	Enables BGP to overlook an illegal router ID. For example, this variable enables the acceptance of a connection from a peer that sends an open message using a router ID of 0. The default is enable.
Synchronization	Enables or disables the router to accept routes from BGP peers without waiting for an update from the IGP. The default is enable.
MaxEqualCostRoutes	Configures the maximum number of equal-cost-paths that are available to a BGP router by limiting the number of equal-cost-paths the routing table can store. The default value is 1; the range is 1–8.
IbgpReportImportRoute	Configures BGP to report imported routes to an interior BGP (IBGP) peer. This variable also enables or disables reporting of non-BGP imported routes to other IBGP neighbors. The default is enable.
FlapDampEnable	Enables or disables route suppression for routes that go up and down (flap). The default is disable.
QuickStart	Enables or disables the Quick Start feature, which forces the BGP speaker to begin establishing peers immediately, instead of waiting for the auto-restart timer to expire. The default is disable.
ConfederationIdentifier	Specifies a BGP confederation identifier in the range of 0–65535.
ConfederationPeers	Lists adjoining ASs that are part of the confederation in the format (5500,65535,0,10,...,..). This value can use 0–255 characters.
RouteReflectionEnable	Enables or disables the reflection of routes from IBGP neighbors. The default is enable.
RouteReflectorClusterId	Configures a reflector cluster ID IP address. This variable applies only if you enable RouteReflectionEnable, and if multiple route reflectors are in a cluster.
ReflectorClientToClientReflection	Enables or disables route reflection between two route reflector clients. This variable applies only if RouteReflectionEnable is enable. The default is enable.

Name	Description
RouteRefresh	Enables or disables route refresh. If enabled, a route refresh request received by a BGP speaker causes the speaker to resend all route updates it contains in its database that are eligible for the peer that issues the request. This parameter only applies to VRF 0.

Configuring 4-byte AS numbers

Configure AS numbers using the 4-byte format and represent the numbers in octets. This procedure does not apply for BGP+.

Before you begin

- You cannot modify the global BGP configuration unless BGP is disabled.
- Make sure that you define AS numbers in policies the same way that you configure them for the router. The choices are `asplain` (regular expression) or `asdot` (dot notation). If you create policies using `asplain` and configure the switch with `asdot`, the match will not occur.

About this task

Use BGP 4 byte AS numbers to ensure the continuity of loop-free inter-domain routing information between ASs and to control the flow of BGP updates as 2 byte AS numbers will deplete soon. AS Plain notation format is the default and the preferred form of representing 4-byte AS numbers over the AS dot notation format. You have an option to configure AS dot notation format as well. With AS dot notation, analyzing and troubleshooting any issues encountered becomes difficult as it is incompatible with the regular expressions used by most of the network providers.

Procedure

1. In the navigation tree, open the following folders: **Configuration > IP**.
 2. Click **BGP**.
 3. To change the AS number format, select **disable** for **AdminStatus**.
 4. Click **Apply**.
 5. In **4-byteAs** , select **enable**.
 6. In **AsDot**, select **enable**.
 7. Enter the 4-byte AS number in octets in the **LocalAs** field.
 8. In **AdminStatus** , select **enable**.
 9. Click **Apply**.
-

4–byte AS field descriptions

Use the data in the following table to use the 4–byte AS related fields on the **Generals** tab.

Name	Description
LocalAs	Configures the local autonomous system (AS) number. You cannot change this field when AdminStatus is set to enable. This field sets a 2-byte local AS number in the range from 0 to 65535. To set a 4-byte local AS number, click enable in the 4ByteAs field and enter a number in the NewLocalAs field. Attention: This parameter is not supported with BGP +
4byteAs	Enables or disables the switch from using 4 byte numbers for autonomous systems.
AsDot	Enables or disables representing AS numbers in octects. The default is disable so the switch uses the plain notation format. If you enable this field and the 4ByteAs field, enter the AS number in the NewLocalAs field. Attention: This parameter is not supported with BGP +

Configuring aggregate routes

Before you begin

- Enable aggregate routes globally.
- You need the appropriate aggregate address and mask.
- If required, policies exist.
- Change the VRF instance as required to configure BGP on a specific VRF instance. The VRF must have an RP trigger of BGP. Not all parameters are configurable on non0 VRFs.

About this task

Configure aggregate routes so that the router advertises a single route (aggregate route) that represents all destinations. Aggregate routes also reduce the size of routing tables.

Procedure

1. In the navigation tree, open the following folders: **Configuration > IP**.
 2. Click **BGP**.
 3. Click the **Aggregates** tab.
 4. Click **Insert**.
 5. Configure the aggregate **Address** and **Mask**.
 6. Select **enable** for **AsSetGenerate** and **SummaryOnly** as required.
 7. Configure policies for the aggregate route.
 8. Click **Insert**.
-

Aggregates field descriptions

Use the data in the following table to use the **Aggregates** tab.

Name	Description
Address	Specifies the aggregate IP address.
Mask	Specifies the aggregate subnet mask.
AsSetGenerate	Enables or disables AS-set path information generation. The default is disable.
SummaryOnly	Enables or disables the summarization of routes in routing updates. Enable this parameter to create the aggregate route and suppress advertisements of more-specific routes to all neighbors. The default is disable.
SuppressPolicy	Specifies the route policy (by name) used for the suppressed route list. Enable this parameter to create the aggregate route and suppress advertisements of the specified routes.
AdvertisePolicy	Specifies the route policy (by name) used for route advertisements. The route policy selects the routes that create AS-set origin communities.
AttributePolicy	Specifies the route policy (by name) used to determine aggregate route attributes.

Configuring aggregate IPv6 routes

Configure IPv6 aggregate routes so that the router advertises a single route (aggregate route) that represents all destinations. Aggregate routes also reduce the size of routing tables.

Before you begin

- Aggregate routes are enabled.
- You have determined the appropriate aggregate prefix and length.
- If required, policies exist.

Procedure

1. In the navigation tree, open the following folders: **Configuration > IPv6**.
 2. Click **BGP+**.
 3. Click the **Aggregates** tab.
 4. Click **Insert**.
 5. Specify the **aggregate Address** and **PrefixLen**
 6. Configure **AsSetGenerate** and **SummaryOnly** as required.
 7. Configure policies for the aggregate route.
 8. Click **Insert**.
-

Aggregates field descriptions

Use the data in the following table to use the **Aggregates** tab.

Name	Description
Address	Specifies the aggregate address. The default is none.
PrefixLen	Specifies the length of the prefix (in bits).
AsSetGenerate	Enables or disables AS-set path information generation. The default is disable.
SummaryOnly	Enables or disables the summarization of routes in routing updates. Enable this parameter to create the aggregate route and suppress advertisements of more-specific

Name	Description
	routes to all neighbors. The default is disable.
SuppressPolicy	Specifies the route policy (by name) used for the suppressed route list. Enable this parameter to create the aggregate route and suppress advertisements of the specified routes.
AdvertisePolicy	Specifies the route policy (by name) used for route advertisements. The route policy selects the routes that create AS-set origin communities.
AttributePolicy	Specifies the route policy (by name) used to determine aggregate route attributes.

Configuring allowed networks

Before you begin

- Change the VRF instance as required to configure BGP on a specific VRF instance. The VRF must have an RP trigger of BGP. Not all parameters are configurable on non0 VRFs.

About this task

Configure network addresses to determine the network addresses that BGP advertises. The allowed addresses determine the BGP networks that originate from the Avaya Virtual Services Platform 9000.

Procedure

1. In the navigation tree, open the following folders: **Configuration > IP**.
 2. Click **BGP**.
 3. Click the **Network** tab.
 4. Click **Insert**.
 5. Configure the network address, mask, and metric.
 6. Click **Insert**.
-

Network field descriptions

Use the data in the following table to use the **Network** tab.

Name	Description
NetworkAddr	Specifies the network prefix that BGP advertises.
NetworkMask	Specifies the network subnet mask.
NetworkMetric	Specifies the metric to use when the system sends an update for the routes in the network table. The metric configures the MED for the routes advertised to EBGP peers. The range is 0–65535.

Configuring allowed IPv6 networks

Configure IPv6 network addresses to determine the network addresses that BGP advertises. The allowed addresses determine the BGP networks that originate from the Avaya Virtual Services Platform 9000.

Procedure

1. In the navigation tree, open the following folders: **Configuration > IPv6**.
 2. Click **BGP+**.
 3. Click the **Network** tab.
 4. Click **Insert**.
 5. Configure the **network address**, **prefix length** and **metric**.
 6. Click **Insert**.
-

Network field descriptions

Use the data in the following table to use the **Network** tab.

Name	Description
NetworkAfPrefixLen	Specifies the network prefix length. The default is none.

Name	Description
NetworkAfMetric	Specifies the metric used when an update is sent for the routes in the network table. The metric configures the MED for the routes advertised to EBGP peers. The range is 0 to 65535. The default is 0.
NetworkAfAddr	Specifies the network prefix that BGP advertises. The default is none.

Configuring BGP peers

Before you begin

- Change the VRF instance as required to configure BGP on a specific VRF instance. The VRF must have an RP trigger of BGP. Not all parameters are configurable on non0 VRFs.

About this task

Configure BGP peers to connect two routers to each other for the purpose of exchanging routing information. BGP peers exchange complete routing information only after they establish the peer connection.

Procedure

1. In the navigation tree, open the following folders: **Configuration > IP**.
 2. Click **BGP**.
 3. Click the **Peers** tab.
 4. Click **Insert**.
 5. Configure the peer as required.
 6. Click **Insert**.
 7. In the **Enable** column, double-click the value, and then select **enable**.
By default, new peer configuration parameters are disabled.
 8. Click **Apply**.
 9. To modify a peer configuration, double-click the value, and then select a new value.
-

Peers field descriptions

Use the data in the following table to use the **Peers** tab.

Name	Description
RemoteAddr	Specifies the remote IP address of the entered BGP peer.
GroupName	Specifies the peer group name to which the peer belongs (optional).
PeerState	Specifies the BGP peer connection state.
RemoteAs	Configures a remote AS number for the peer or peer-group in the range 0–65535.
NewRemoteAs	Specifies the new remote Autonomous System number of the peer.
Enable	Controls whether the peer connection is enabled or disabled. The default is disabled.
RoutePolicyIn	Specifies the policy (by name) that applies to all network routes learned from this peer.
RoutePolicyOut	Specifies the policy (by name) that applies to all outgoing route updates.
DefaultOriginate	When enabled, specifies that the current route originated from the BGP peer. This parameter enables or disables sending the default route information to the specified neighbor or peer. The default value is false.
DefaultOriginateIPv6	When enabled, it allows a BGP speaker (the local router) to send the default ipv6 route to a neighbor for use as a default route. The default value is false.
ConnectRetryInterval	Specifies the time interval, in seconds, for the connect retry timer. The suggested value for this timer is 120 seconds. The range is 1 to 65535.
HoldTimeConfigured	Specifies the time interval, in seconds, for the hold time for this BGP speaker with this peer. This value is in an open message sent to this peer by this BGP speaker. To determine the hold time with the peer, the switch compares this value with the HoldTime value in an open message received from the peer. The HoldTime must be at least three seconds. If the value is zero, the hold time does not establish with the peer. The suggested value for this timer is 180 seconds. The range is 0 to 65535.
KeepAliveConfigured	Specifies the time interval, in seconds, for the KeepAlive timer configured for this BGP speaker with this peer. KeepAliveConfigured determines the keep alive message

Name	Description
	frequency relative to HoldTimeConfigured; KeepAlive indicates the actual time interval for the keep alive messages. The recommended maximum value for this timer is one-third of HoldTimeConfigured. If KeepAliveConfigured is zero, no periodic keep alive messages are sent to the peer after the peers establish a BGP connection. Avaya recommends that you configure a value of 60 seconds. The range is 0 to 21845.
MD5Authentication	Enables and disables MD5 authentication.
EbgpMultiHop	Enables or disables a connection to a BGP peer that is more than one hop away from the local router. The default value is disable.
RemovePrivateAs	Strips (when enabled) private AS numbers when the switch sends an update. The default is enable.
UpdateSourceInterface	Specifies the source IP address to use when the switch sends EBGP packets to this peer or peer group.
NextHopSelf	Specifies that the next-hop attribute in an IBGP update is the address of the local router or the router that generates the IBGP update. The default is disable.
DebugMask	<p>Displays the specified debug information for the BGP peer. The default value is none.</p> <ul style="list-style-type: none"> • none disables all debug messages. • event enables the display of debug event messages. • state enables display of debug state transition messages. • update enables display of debug messages related to updates transmission and reception. • error enables the display of debug error messages. • trace enables the display of debug trace messages. • init enables the display of debug initialization messages. • all enables all debug messages. • packet enables the display of debug packet messages. • warning enables the display of debug warning messages. • filter enables the display of debug messages related to filtering.
AdvertisementInterval	Specifies the time interval, in seconds, that elapses between each transmission of an advertisement from a

Name	Description
	BGP neighbor. The default value is 30 seconds and the range is 5–120 seconds.
Weight	Specifies the peer or peer group weight, or the priority of updates the system can receive from this BGP peer. The default value is 100 and the range is 0–65535.
MaxPrefix	Configures a limit on the number of routes accepted from a neighbor. The default value is 12000 routes and the range is 0–2147483647. A value of 0 means no limit exists.
RouteReflectorClient	Specifies that this peer is a route reflector client.
SoftReconfigurationIn	When enabled, the router relearns routes from the specified neighbor or group of neighbors without restarting the connection after the policy changes in the inbound direction. The default value is disable. Enabling SoftReconfigurationIn stores all BGP routes in local memory (even non-best routes).
SendCommunity	Enables or disables sending the community attribute of the update message to the specified peer. The default value is disable.
IPv6Cap	Enable or disable the ipv6 capability on the bgp neighbor peer. The default value is disable.
RouteRefresh	Enables or disables route refresh. If enabled, a route refresh request received by a BGP speaker causes the speaker to resend all route updates it contains in its database that are eligible for the peer that issues the request. This parameter only applies to VRF 0.
AsOverride	Specifies that the AS Override variable can be enabled or disabled for the BGP peer.
AllowAsIn	Specifies the number of AS-in allowed for the BGP peer. The range is 0–10.
IPv6RoutePolicyIn	Specifies the Ipv6 Route Policy Name that is applied to all networks learned from the entered peer.
IPv6RoutePolicyOut	Specifies the Ipv6 Route policy name that is applied to all outgoing updates.

Configuring peer groups

Before you begin

- Change the VRF instance as required to configure BGP on a specific VRF instance. The VRF must have an RP trigger of BGP. Not all parameters are configurable on non0 VRFs.

About this task

Configure or edit peer groups to create update policies for neighbors in the same group.

Procedure

1. In the navigation tree, open the following folders: **Configuration > IP**.
 2. Click **BGP**.
 3. Click the **Peer Groups** tab.
You can modify an existing parameter by double-clicking the value.
 4. Click **Insert**.
 5. Configure the peer group as required.
 6. Click **Insert**.
-

Peer Groups field descriptions

Use the data in the following table to use the **Peer Groups** tab.

Name	Description
Index	Specifies the index of this peer group. The range is 1–1024.
GroupName	Specifies the peer group to which this neighbor belongs (optional).
Enable	Enables or disables the peer group.
RemoteAs	Configures a remote AS number for the peer-group in the range 0–65535.
DefaultOriginate	When enabled, the BGP speaker (the local router) sends the default route 0.0.0.0 to a group of neighbors for use as a default route. The default is disabled.

Name	Description
EbgpMultiHop	When enabled, the switch accepts and attempts BGP connections to external peers that reside on networks that do not directly connect. The default is disabled.
AdvertisementInterval	Specifies the time interval, in seconds, that elapses between BGP routing updates. The default value is 30 seconds and the range is 1–120 seconds.
KeepAlive	Specifies the time interval, in seconds, between sent BGP keep alive messages to remote peers. The range is 1–21845 and the default value is 60.
HoldTime	Configures the hold time for the group of peers in seconds. Avaya recommends that you use a value that is three times the value of the KeepAlive time. The range is 3–65535 and the default value is 180.
Weight	Assigns an absolute weight to a BGP network. The default value is 100 and the range is 0–65535.
MaxPrefix	Limits the number of routes accepted from this group of neighbors. A value of zero indicates no limit. The default value is 12000 routes and the range is 0–2147483647.
NextHopSelf	Specifies that the switch must set the NextHop attribute to the local router address before it sends updates to remote peers.
RoutePolicyIn	Specifies the route policy that applies to all networks learned from this group of peers.
RoutePolicyOut	Specifies the route policy that applies to all outgoing updates to this group of peers.
RouteReflectorClient	Specifies that this peer group is a route reflector client.
SoftReconfigurationIn	When enabled, the router relearns routes from the specified neighbor or group of neighbors without restarting the connection after the policy changes in the inbound direction. The default value is enable. Enabling SoftReconfigurationIn stores all BGP routes in local memory (even non-best routes).
MD5Authentication	Enables and disables MD5 authentication. The default is disable.
RemovePrivateAs	Strips (when enabled) private AS numbers when the switch sends an update. The default is enable.
SendCommunity	Enables or disables sending the community attribute of the update message to the specified peer group. The default value is disable.
AfUpdateSourceInterfaceType	Specifies if the interface is IPv4 or IPv6.

Name	Description
AfUpdateSourceInterface	Specifies the IP address used for circuitless IP (CLIP) for this peer group.
Vpvn4Address	Enables BGP address families for IPv4 (BGP) and L3 VPN (MP-BGP) support. Enable this parameter for VPN/VRF Lite routes.
IpvpnLiteCap	Specifies (when enabled) that IP VPN Lite capability can be enabled or disabled on the BGP neighbor peer. The default is disable.
RouteRefresh	Enables or disables route refresh. If enabled, a route refresh request received by a BGP speaker causes the speaker to resend all route updates it contains in its database that are eligible for the peer that issues the request. This parameter only applies to VRF 0.
AsOverride	Specifies that the AS Override parameter can be enabled or disabled for the BGP peer group. The default is disable.
AllowedAsIn	Specifies the number of AS-in allowed for the BGP peer group. The range is 1–10.
Ipv6RoutePolicyIn	Specifies the IPv6 route policy that applies to all networks learned from this group of peers.
Ipv6RoutePolicyOut	Specifies the IPv6 route policy that applies to all outgoing updates to this group of peers.

Viewing route information

Before you begin

- Change the VRF instance as required to view BGP information about a specific VRF instance. The VRF must have an RP trigger of BGP. Not all parameters are configurable on non0 VRFs.

About this task

View information about BGP routes in the routing table to determine connectivity or configuration problems.

Procedure

1. In the navigation tree, open the following folders: **Configuration > IP**.
2. Click **BGP**.
3. Click the **Bgp Routes** tab.

4. Click **Range**.
 5. Choose the type of information you want to see by configuring the options in the BGP Route Range dialog box.
 6. Click **Set Range**.
 7. Click **Close**.
 8. In the **BGP Routes** tab, click **Refresh**.
-

Bgp Routes field descriptions

Use the data in the following table to use the **Bgp Routes** tab.

Name	Description
Type	This variable appears only in the BGP Route Range dialog box, not on the Bgp Routes tab. Specifies the route types to query: <ul style="list-style-type: none"> • All = all routes in the route table • Exact Prefix = only routes that exactly match the prefix/len • Longer Prefix = all routes that match the prefix with longer masks.
IpAddrPrefix	Specifies the IP address prefix in the Network Layer Reachability Information (NLRI) field. This variable is an IP address that contains the prefix with a length specified by IpAddrPrefixLen. Bits beyond the length specified by IpAddrPrefixLen are set to zero.
IpAddrPrefixLen	Specifies the length, in bits, of the IP address prefix in the NLRI field.
Peer	Specifies the IP address of the peer from which the router learns path information.
NextHop	Specifies the address of the border router to use to reach the destination network.
Origin	Specifies the ultimate origin of the path information: <ul style="list-style-type: none"> • igp = networks are interior • egp = networks learned through EGP • incomplete = undetermined

Name	Description
LocalPref	Specifies the originating BGPv4 speaker preference for an advertised route. A value of -1 means that the speaker has no preference.
ASPathSegment	<p>Specifies the sequence of AS path segments. Each AS path segment is represented by a triple <type, length, value>. The type is a one-octet field that has the following possible value:</p> <ul style="list-style-type: none"> • AS_SET: specifies an unordered set of ASs that a route in the update message traversed. • AS_SEQUENCE: specifies an ordered set of ASs that a route in the update message traversed. • AS_CONFED_SET: specifies an unordered set of sub-ASs in the confederation that a route in the update message traversed. • AS_CONFED_SEQUENCE: specifies an ordered set of sub-ASs in the confederation that a route in the update message traversed. The length is a one-octet field that contains the number of ASs in the value field. The value field contains one or more AS numbers, where each AS is represented in the octet string as a pair of octets according to the following algorithm: first-byte-of-pair = ASNumber / 256 second-byte-of-pair = ASNumber & 255
MultiExitDisc	Specifies the MED. The MED discriminates between multiple exit points to an adjacent AS. A value of -1 indicates no MED.
AtomicAggregate	Specifies whether the local system selected a less-specific route without selecting a more-specific route.
AggregatorAS	Specifies the AS number of the last BGPv4 speaker that performed route aggregation. A value of zero (0) indicates no aggregation.
AggregatorAddr	Specifies the IP address of the last BGPv4 speaker that performed route aggregation. A value of zero (0) indicates no aggregation.
Best	Indicates whether the router selected this route as the best BGPv4 route.
Unknown	Specifies one or more path attributes not understood by this BGPv4 speaker. Size zero (0) indicates the absence of unknown attribute. Octets beyond the maximum size are not recorded.

Viewing community attributes

Before you begin

- Change the VRF instance as required to view BGP information about a specific VRF instance. The VRF must have an RP trigger of BGP. Not all parameters are configurable on non0 VRFs.

About this task

View community attributes to determine connectivity or configuration problems for a group of destinations.

Procedure

1. In the navigation tree, open the following folders: **Configuration > IP**.
2. Click **BGP**.
3. Click the **Bgp Community Attributes** tab.
4. Click **Range**.
5. Choose the type of information you want to see by configuring the options in the **BGP Route Range** dialog box.
6. Click **Set Range**.
7. Click **Close**.
8. In the **Bgp Community Attributes** tab, click **Refresh**.

Bgp Community Attributes field descriptions

Use the data in the following table to use the **Bgp Community Attributes** tab.

Name	Description
Type	This variable appears only in the BGP Route Range dialog box, not the Bgp Community Attributes tab. Specifies the route types to query: <ul style="list-style-type: none"> • All = all routes in the route table • Exact Prefix = only routes that exactly match the prefix/len • Longer Prefix = all routes that match the prefix with longer masks.

Name	Description
IpAddrPrefix	Specifies the IP address prefix in the NLRI field. This variable is an IP address that contains the prefix with a length specified by IpAddrPrefixLen. Bits beyond the length specified by IpAddrPrefixLen are set to zero.
IpAddrPrefixLen	Specifies the length, in bits, of the IP address prefix in the NLRI field.
Peer	Specifies the IP address of the peer from which the router learns the path information.
Communities	Represents community path attributes as <path, value>. The length is a one-octet field that contains the number of community numbers in the value field. The value field contains one or more community numbers, where each community is 4 octets in the octet string. Community numbers beyond a maximum size of 255 are not recorded.
Originateld	Specifies the IP address of the router that originates the routes.
ClusterId	Represents cluster ID path attributes as <length, value>. The length is a one-octet field that contains the cluster ID numbers in the value field. The value field contains one or more cluster IDs, where each cluster ID is 4 octets in the octet string, each four octets is the IP address of the router that acts as a route reflector. The cluster ID numbers beyond a maximum size 255 are not recorded.

Viewing IPv6 community attributes

About this task

View IPv6 community attributes for specific routes to utilize the update message fields to communicate information between BGP speakers. The Path Attribute values allow you to specify the prefixes that the BGP session can exchanged, or which of the multiple paths of a specified prefix to use.

Procedure

1. In the navigation tree, open the following folders: **Configuration > IPv6**.
2. Click **BGP+**.
3. Select a route for which you want to view the route summary information.
4. Click the **Bgp Route Summary** tab.

5. Click the **Route Comm Attr** option on the menu.
The **BGP Path Attributes** tab opens with the BGP IPv6 community attribute information.

BGP Path Attributes field descriptions

Use the data in the following table to use the **Community List** tab.

Name	Description
Origin	Specifies the ultimate origin of the path information.
NextHopAddr	Specifies the address of the border router that is used to access the destination network. This address is the nexthop address received in the UPDATE packet associated with this prefix.
Med	This metric is used to discriminate between multiple exit points to an adjacent autonomous system. When the MED value is absent but has a calculated default value, this object will contain the calculated value.
LocalPref	Specifies the value used during route decision process in the BGP protocol. Applicable to BGP only.
AggregatorAS	Specifies the AS number of the last BGP4 speaker that performed route aggregation. If the AGGREGATOR path attribute is absent, this object will not be present in the conceptual row.
AggregatorAddr	Specifies the IP address of the last BGP4 speaker that performed route aggregation. If the AGGREGATOR path attribute is absent, this object will not be present in the conceptual row.
String	<p>This is a string representing the autonomous system path to the network which was received from the peer which advertised it. The format of the string is implementation-dependent, and is designed for operator readability.</p> <p>Note: SnmpAdminString is only capable of representing a maximum of 255 characters. This may lead to the string being truncated in the presence of a large AS Path.</p>

Displaying dampened routes information

Before you begin

- Change the VRF instance as required to view BGP information about a specific VRF instance. The VRF must have an RP trigger of BGP. Not all parameters are configurable on non0 VRFs.
- Enable dampened routes.

About this task

Display dampened path information to see which routes are suppressed.

Procedure

1. In the navigation tree, open the following folders: **Configuration > IP**.
 2. Click **BGP**.
 3. Click the **Dampened Routes** tab.
-

Dampened Routes field descriptions

Use the data in the following table to use the **Dampened Routes** tab.

Name	Description
IpAddrPrefix	Specifies the IP address prefix in the NLRI field. This variable is an IP address that contains the prefix with a length specified by IpAddrPrefixLen. Bits beyond the length specified by IpAddrPrefixLen are set to zero.
IpAddrPrefixLen	Specifies the length, in bits, of the IP address prefix in the NLRI field.
Peer	Specifies the IP address of the peer from which the router learns the path information.
FlapPenalty	Specifies the penalty based on number of route flaps.
FlapCount	Specifies the number of times a route flapped (went down and up) since the last time the penalty was reset to zero.
RouteDampened	Indicates whether this route is suppressed or announced.
ReuseTime	Specifies the system-configured time for route reuse.

Configuring redistribution to BGP

Before you begin

- If required, a route policy exists.

About this task

Configure redistribute entries for BGP to announce routes of a certain source type to BGP, for example, direct, static, Routing Information Protocol (RIP), and Open Shortest Path First (OSPF). If you do not configure a route policy, then the switch uses the default action based on metric, metric type, and subnet. Use a route policy to perform detailed redistribution.

Procedure

1. In the navigation tree, open the following folders: **Configuration > IP**.
 2. Click **BGP**.
 3. Click the **Redistribute** tab.
 4. Click **Insert**.
 5. Configure the source protocol.
 6. If required, choose a route policy.
 7. Configure the metric to apply to redistributed routes.
 8. Enable the redistribution instance.
 9. Click **Insert**.
-

Redistribute field descriptions

Use the data in the following table to use the **Redistribute** tab.

Name	Description
DstVrfId	Specifies the destination VRF instance (read-only).
Protocol	Specifies the protocols that receive the redistributed routes.
SrcVrfId	Specifies the source VRF instance (read-only).
RouteSource	Specifies the source protocol for the route redistribution entry.
Enable	Enables (or disables) a BGP redistribute entry for a specified source type.

Name	Description
RoutePolicy	Configures the route policy to use for the detailed redistribution of external routes from a specified source into the BGP domain.
Metric	Configures the metric for the redistributed route. The value can be a range between 0–65535. The default value is 0. Avaya recommends that you use a value that is consistent with the destination protocol.

Configuring a prefix list

Before you begin

- Change the VRF instance as required to configure a prefix list on a specific VRF instance.

About this task

Configure a prefix list to allow or deny specific route updates. A prefix list policy specifies route prefixes to match. After a match occurs, the system uses the route.

The prefix list contains a set of contiguous or noncontiguous routes. Reference prefix lists by name from within a routing policy.

Procedure

1. In the navigation tree, open the following folders: **Configuration > IP**.
 2. Click **Policy**.
 3. Click the **Prefix List** tab.
 4. Click **Insert**.
 5. In the **ID** box, type an ID for the prefix list.
 6. In the **Prefix** box, type an IP address for the route.
 7. In the **PrefixMaskLength** box, type the length of the prefix mask.
 8. Configure the remaining parameters as required.
 9. Click **Insert**.
-

Prefix List field descriptions

Use the data in the following table to use the **Prefix List** tab.

Name	Description
ID	Configures the list identifier.
Prefix	Configures the IP address of the route.
PrefixMaskLen	Configures the specified length of the prefix mask. You must enter the full 32-bit mask to exact a full match of a specific IP address, for example, if you create a policy to match on the next hop.
Name	Names a specified prefix list during the creation process or renames the specified prefix list. The name length can use from 1–64 characters.
MaskLenFrom	Configures the lower bound of the mask length. The default is the mask length. Lower bound and higher bound mask lengths together can define a range of networks.
MaskLenUpto	Configures the higher bound mask length. The default is the mask length. Lower bound and higher bound mask lengths together can define a range of networks.

Configuring IPv6 Prefix List

Use IPv6 prefix lists to allow or deny specific route updates. A prefix list policy specifies route prefixes to match. When there is a match, the route is used.

Procedure

1. In the navigation tree, open the following folders: **Configuration > IPv6**.
2. Click **Policy**.
3. In the **Prefix List** tab, click **Insert**.
4. Edit the parameters as required.
5. Click **Insert**.

Ipv6–Prefix list field descriptions

Use the data in the following table to use the **Ipv6–Prefix List** tab.

Name	Description
Id	Specifies the prefix list. The range is 0 to 65535.

Name	Description
Prefix	Specifies the prefix IPv6 address.
PrefixMaskLen	Specifies the length of the prefix mask. You must enter the full 128-bit mask to exact a full match of a specific IPv6 address (for example, when creating a policy to match the next-hop).
Name	Names a specified prefix list during the creation process or renames the specified prefix list. The name can be from 1 to 64 characters in length.
MaskLenFrom	Specifies the lower bound on the mask length. The default is the mask length. Lower bound and higher bound mask lengths together can define a range of networks.
MaskLenUpto	Specifies the higher bound mask length. The default is the mask length. Lower bound and higher bound mask lengths together can define a range of networks.

Viewing IPv6 BGP+ route summary information

You can display current IPv6 BGP+ route information.

Procedure

1. In the navigation tree, open the following folders: **Configuration > IPv6**.
 2. Click **BGP+**.
 3. Click the **Bgp Route Summary** tab to view the BGP route summary information
-

Bgp Route Summary field descriptions

Use the data in the following table to use the **Bgp Route Summary** tab.

Name	Description
Prefix	Specifies the IP address prefix in the Network Layer Reachability Information (NLRI) field. This is an IP address that contains the prefix with a length specified by

Name	Description
	IpAddrPrefixLen. Any bits beyond the length specified by IpAddrPrefixLen are set to zero.
PrefixLen	Specifies the length, in bits, of the IP address prefix in the NLRI field.
LocalAddr	The local address of this entry's BGP connection.
RemoteAddr	Specifies the IP address of the peer from which path information was learned.

Configuring an AS path list

Before you begin

- Change the VRF instance as required to configure BGP on a specific VRF instance. The VRF must have an RP trigger of BGP. Not all parameters are configurable on non0 VRFs.

About this task

Configure an AS path list to restrict the routing information a router learns or advertises to and from a neighbor. The AS path list acts as a filter that matches AS paths.

Procedure

1. In the navigation tree, open the following folders: **Configuration > IP**.
2. Click **Policy**.
3. Click the **As Path List** tab.
4. Click **Insert**.
5. Enter the appropriate information for your configuration.
6. Click **Insert**.

As Path List field descriptions

Use the data in the following table to use the **As Path List** tab.

Name	Description
Id	Specifies the AS path list. The range is 0–1024.

Name	Description
MemberId	Specifies the AS path access list member ID. The range is 0–65535.
Mode	Specifies the action to take if the system selects a policy for a specific route. Select permit (allow the route) or deny (ignore the route).
AsRegularExpression	Specifies the expression to use for the AS path.

Configuring a community access list

Before you begin

- Change the VRF instance as required to configure BGP on a specific VRF instance. The VRF must have an RP trigger of BGP. Not all parameters are configurable on non0 VRFs.

About this task

Configure community lists to specify permitted routes by using their BGP community. This list acts as a filter that matches communities or AS numbers.

Procedure

1. In the navigation tree, open the following folders: **Configuration > IP**.
2. Click **Policy**.
3. Click the **Community List** tab.
4. Click **Insert**.
5. Configure the list as required.
6. Click **Insert**.

Community List field descriptions

Use the data in the following table to use the **Community List** tab.

Name	Description
Id	Specifies the community list. The range is 0–1024.
MemberId	Specifies the community list member ID. The range is 0–65535.

Name	Description
Mode	Specifies the action to take if the system selects a policy for a specific route. Select permit (allow the route) or deny (ignore the route).
Community	Specifies the community access list community string.

Configuring an extended community list

Before you begin

- Change the VRF instance as required to configure BGP on a specific VRF instance. The VRF must have an RP trigger of BGP. Not all parameters are configurable on non0 VRFs.

About this task

Configure community lists to specify permitted routes by BGP extended community attributes, including route targets and sites of origin (SOO). This list acts as a filter that matches route targets and SOO.

Procedure

1. In the navigation tree, open the following folders: **Configuration > IP**.
 2. Click **Policy**.
 3. Click the **Ext Community List** tab.
 4. Click **Insert**.
 5. Configure the list as required.
 6. Click **Insert**.
-

Ext Community List field descriptions

Use the data in the following table to use the **Ext Community List** tab.

Name	Description
Id	Specifies the extended community list. The range is 0–1024.
MemberId	Specifies the community list member ID. The range is 0–65535.

Name	Description
RTType	Specifies the route target type for the extended community rule. The type can be: none, AS number, or IP address.
RTAddress	Specifies the extended community IP address for the route target entry.
RTAsNumber	Specifies the extended community AS number for the route target entry. The range is 0–65535.
RTAssignedNum	Specifies the extended community assigned number. The range is 0–2147483647 or 0–65535 for an IP address.
SooType	Specifies the SOO type for the extended community rule.
SooAddress	Specifies the site of origin.
SooAsNumber	Specifies the extended community AS number for the site of origin. The range is 0–65535.
SooAssignedNum	Specifies the extended community assigned number for the site of origin. The range is 0–2147483647 or 0–65535 for an IP address..
Vrflid	Specifies the VRF instance.

Chapter 7: BGP configuration examples

This chapter shows configuration examples for BGP deployment options.

IPv6 tunnel configurations for BGP+

You must configure an IPv6 tunnel and static routes at BGP+ peers when you use BGP+.

When BGP+ peers advertise route information, they use Update messages to advertise route information. And, when route information is encapsulated in Update messages, BGP+ peers convert their own IPv4 peer addresses to IPv4-mapped IPv6 addresses and insert them into the next-hop field in the Update message.

When the BGP+ software module receives Update messages, it adds route information to the IPv6 Routing Manager (RTM). These RTM routes contain next-hop addresses from the BGP peer that the route was learned from. The next-hop addresses are represented as IPv4-mapped IPv6 addresses.

But, because the IPv6 RTM cannot correlate the IPv4-mapped IPv6 address to a specific outgoing interface, you must create a manually-configured static route to make the link between the BGP peer and the IPv6 tunnel interface so that traffic can reach networks advertised by the peer.

Following is one way to express a static route in an IPv6-configured tunnel for BGP+

```
ipv6 route 0:0:0:0:0:ffff:200.1.1.1/128 cost 1 tunnel 10
```

Avaya recommends that the IPv6-configured tunnel endpoint and the BGP peer reside on the same switch.

If the IPv6 tunnel endpoint and the BGP peer must reside on different switches you can terminate the tunnel on a different switch, but you must consider the following:

- Because the IPv6 tunnel endpoint does not reside on the same switch as the BGP peer, the BGP device cannot use the tunnel as the outgoing interface. That is, to reach the IPv6-configured tunnel endpoint, if the BGP peer resides on a different switch from the IPv6 tunnel endpoint, the next-hop for the manually-configured IPv4-mapped IPv6 static route is the native IPv6 interface next-hop address.
- The node where the tunnel terminates must contain all of the information needed to route the packets between the remote IPv6 network clouds.

Note:

In order for the tunnel endpoint switch to be aware of all of the necessary IPv6 routes, you may need to redistribute the BGP routes into OSPFv3.

IPv4-mapped IPv6 addresses

IPv4-mapped IPv6 addresses are IPv4 addresses that the system has mapped into the IPv6 address space.

The system uses these IPv4-mapped IPv6 addresses for devices that are only IPv4-capable.

These IPv4-mapped IPv6 addresses start with an initial string of 80 zeroes followed by a set of 16 ones, and then the IPv4 address.

Example IPv4-mapped IPv6 address:

When converted to an IPv4-mapped IPv6 address, an IPv4 device address of 222.1.41.90 would be represented as one of the following:

- 0:0:0:0:0:FFFF:222.1.41.90
- ::FFFF:222.1.41.90

The following figure illustrates the components in an IPv4-mapped IPv6 address.

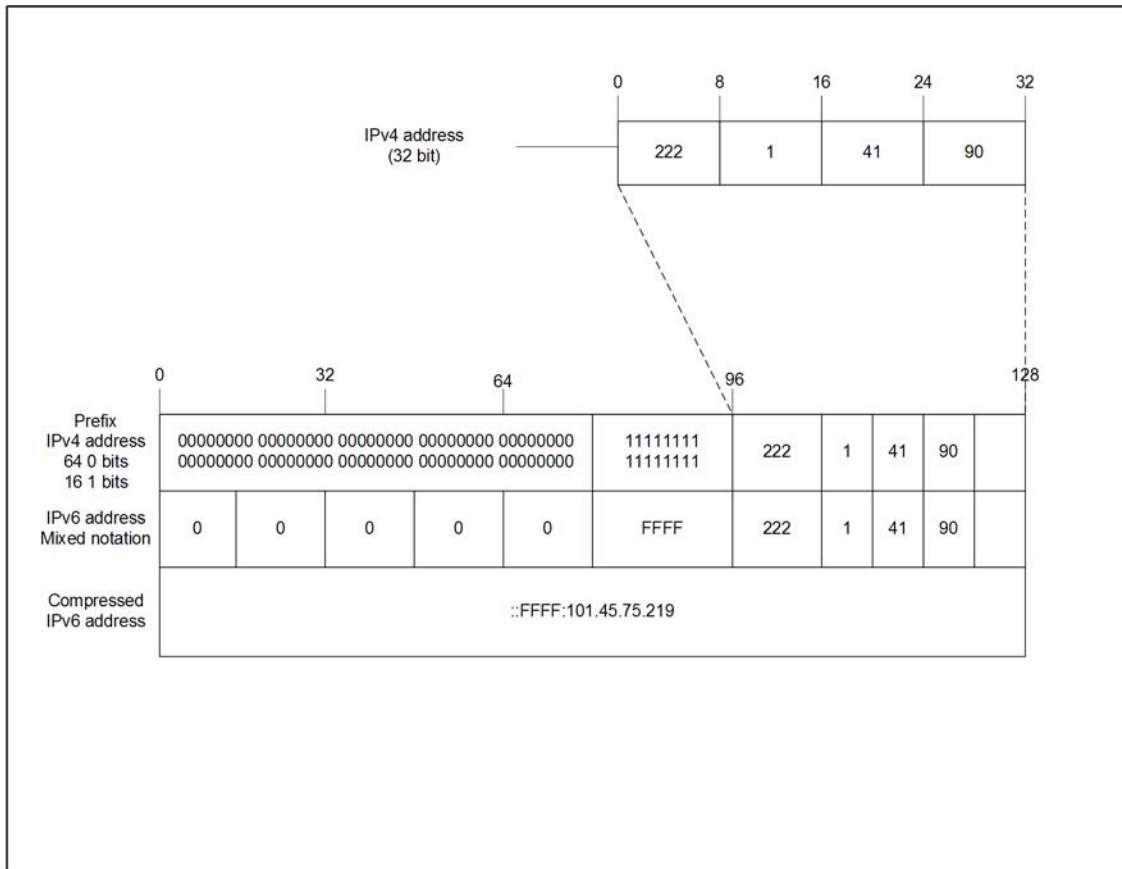


Figure 16: IPv4-mapped IPv6 address components

eBGP+ peership between two VSP 9000 switches with IPv6 tunneling

The following figure shows a sample network that contains eBGP+ peers using IPv6 tunneling.




Figure 17: eBGP+ peers with IPv6 tunneling

The configuration in the figure, *eBGP+ peers with IPv6 tunneling*, assumes that the BGP peer IP address is the next hop.

When you configure the static route for the BGP+ tunnel, you must designate the BGP peer IP address as the next hop in most cases.

You can configure multiple static routes, using the same tunnel, but you must ensure reachability when you create the static routes.

R1 configuration:

```
interface GigabitEthernet 3/2
brouter port 3/2 vlan 2090 subnet 200.1.1.1/255.255.255.0 mac-offset 2
exit
# BGP CONFIGURATION - GlobalRouter
#

router bgp as-dot enable
router bgp 65000 enable
router bgp
neighbor "200.1.1.2"
neighbor 200.1.1.2 remote-as 65001
neighbor 200.1.1.2 address-family ipv6
neighbor 200.1.1.2 enable
exit
# IPV6 CONFIGURATION
#

ipv6 forwarding

# IPV6 TUNNEL CONFIGURATION
#

ipv6 tunnel 10 source 200.1.1.1 address 300:0:0:0:0:0:0:1/126 destination 200.1.1.2

#
# IPV6 STATIC ROUTE CONFIGURATION
#

ipv6 route 0:0:0:0:0:ffff:200.1.1.2/128 cost 1 tunnel 10
#
```

R2 configuration:

```
interface GigabitEthernet 4/32
brouter port 4/32 vlan 2090 subnet 200.1.1.2/255.255.255.0 mac-offset 2
exit
```

BGP configuration examples

```
# BGP CONFIGURATION - GlobalRouter
#

router bgp as-dot enable
router bgp 65000 enable
router bgp
neighbor "200.1.1.1"
neighbor 200.1.1.1 remote-as 65000
neighbor 200.1.1.1 address-family ipv6
neighbor 200.1.1.1 enable
exit
# IPV6 CONFIGURATION
#

ipv6 forwarding

# IPV6 TUNNEL CONFIGURATION
#

ipv6 tunnel 10 source 200.1.1.2 address 300:0:0:0:0:0:2/126 destination
200.1.1.1

#
# IPV6 STATIC ROUTE CONFIGURATION
#

ipv6 route 0:0:0:0:0:ffff:200.1.1.1/128 cost 1 tunnel 10
#
```

iBGP+ peership on CLIP between two VSP 9000 switches with IPv6 tunneling

The following figure shows a sample network that contains iBGP+ peers using IPv6 tunneling.




Figure 18: iBGP+ peers on CLIP interfaces with IPv6 tunneling

You must enable OSPF on the interface and globally.

If you cannot enable OSPF, you must configure static routes to provide reachability to the BGP+ peer.

The static route must point to the next hop for the routes to be installed in the IPv6 RTM.

The next hop must be the BGP peer IP address.

The IPv4 interfaces do not need to connect directly, but the routing table on each switch must include the IPv4 interface of the other switch.

iBGP between the CLIP interfaces needs to run OSPF as a routing protocol so that the BGP neighbor can remain reachable.

eBGP connections cannot use a CLIP interface as an end point.

R1 configuration:

```

interface GigabitEthernet 3/2
brouter port 3/2 vlan 2090 subnet 200.1.1.1/255.255.255.0 mac-offset
2
exit
# OSPF CONFIGURATION - GlobalRouter
#

router ospf enable

# OSPF PORT CONFIGURATION
#

interface gigabitethernet 3/2
ip ospf enable
exit

# CIRCUITLESS IP INTERFACE CONFIGURATION - GlobalRouter
#

interface loopback 1
ip address 1 1.1.1.1/255.255.255.255
ip ospf 1

# BGP CONFIGURATION - GlobalRouter
#

router bgp
no synchronization
exit
router bgp as-dot enable
router bgp 65000 enable
router bgp
neighbor "2.2.2.2"
neighbor 2.2.2.2 remote-as 65000
neighbor 2.2.2.2 next-hop-self
neighbor 2.2.2.2 update-source 1.1.1.1
neighbor 2.2.2.2 address-family ipv6
neighbor 2.2.2.2 enable
exit
# IPV6 CONFIGURATION
#

ipv6 forwarding

# IPV6 TUNNEL CONFIGURATION
#

ipv6 tunnel 10 source 200.1.1.1 address 300:0:0:0:0:0:1/126 destination
200.1.1.2

#
# IPV6 STATIC ROUTE CONFIGURATION
#

ipv6 route 0:0:0:0:ffff:2.2.2.2/128 cost 1 tunnel 10
#

```

R2 configuration:

```
interface GigabitEthernet 4/32
brouter port 4/32 vlan 2090 subnet 200.1.1.2/255.255.255.0 mac-offset
2
exit
# OSPF CONFIGURATION - GlobalRouter
#

router ospf enable

# OSPF PORT CONFIGURATION
#

interface gigabitethernet 4/32
ip ospf enable
exit

# CIRCUITLESS IP INTERFACE CONFIGURATION - GlobalRouter
#

interface loopback 1
ip address 1 2.2.2.2/255.255.255.255
ip ospf 1

# BGP CONFIGURATION - GlobalRouter
#

router bgp
no synchronization
exit
router bgp as-dot enable
router bgp 65000 enable
router bgp
neighbor "1.1.1.1"
neighbor 1.1.1.1 remote-as 65000
neighbor 1.1.1.1 next-hop-self
neighbor 1.1.1.1 update-source 2.2.2.2
neighbor 1.1.1.1 address-family ipv6
neighbor 1.1.1.1 enable
exit
# IPV6 CONFIGURATION
#

ipv6 forwarding

# IPV6 TUNNEL CONFIGURATION
#

ipv6 tunnel 10 source 200.1.1.2 address 300:0:0:0:0:0:2/126 destination
200.1.1.1

#
# IPV6 STATIC ROUTE CONFIGURATION
#

ipv6 route 0:0:0:0:0:ffff:1.1.1.1/128 cost 1 tunnel 10
#
```

Chapter 8: Customer service

Visit the Avaya Web site to access the complete range of services and support that Avaya provides. Go to www.avaya.com or go to one of the pages listed in the following sections.

Getting technical documentation

To download and print selected technical publications and release notes directly from the Internet, go to www.avaya.com/support.

Getting product training

Ongoing product training is available. For more information or to register, you can access the Web site at www.avaya.com/support. From this Web site, you can locate the Training contacts link on the left-hand navigation pane.

Getting help from a distributor or reseller

If you purchased a service contract for your Avaya product from a distributor or authorized reseller, contact the technical support staff for that distributor or reseller for assistance.

Getting technical support from the Avaya Web site

The easiest and most effective way to get technical support for Avaya products is from the Avaya Technical Support Web site at www.avaya.com/support.

Glossary

AS confederation	A single logical AS that comprises multiple sub-ASs to ensure scalability.
AS_TRANS	RFC4893 defines a 4-octet Autonomous System number 23456 to facilitate backward compatibility. This 23456 AS number is also known as AS_TRANS (AS_TRANS=23456).
4-byte AS	4-byte Autonomous System (AS) numbers is the solution to the soon depleting 2-byte AS numbers. It provides a theoretical 4,294,967,296 unique AS numbers in BGP. 4-byte AS numbers are backward compatible with 2-byte AS numbers.
attribute	A unit of data BGP used to describe the prefixes, such as AS-PATH, LOCAL-PREF, NEXT-HOP.
Autonomous System (AS)	A set of routers under a single technical administration, using a single IGP and common metrics to route packets within the AS, and using an EGP to route packets to other ASs.
Autonomous System Number (ASN)	A two-byte number that is used to identify a specific AS.
Border Gateway Protocol (BGP)	An inter-domain routing protocol that provides loop-free inter-domain routing between Autonomous Systems (AS) or within an AS.
BGP+	BGP+ is an extension of BGPv4 to support IPv6. BGP+ carries routing information for multiple network layer protocol address families, for example, IPv6 address family and for IP multicast routes.
Border Gateway Protocol neighbor	Border Gateway Protocol routers that have interfaces to a common network.
Border Gateway Protocol peer	A relationship that is formed between two routers that open a TCP connection to each other for the purpose of exchanging routing information.
Border Gateway Protocol session	An active connection between two routers running BGP.
Border Gateway Protocol speaker	An entity within a BGP router that is used to communicate with other BGP speakers by establishing a peer-to-peer session.
Circuitless IP (CLIP)	A virtual interface that does not map to any physical interface. This interface is often called a <i>loopback</i> .

classless interdomain routing (CIDR)

classless interdomain routing (CIDR)	The protocol defined in RFCs 1517 and 1518 for using subnetwork masks, other than the defaults for IP address classes.
cluster	One or more route reflectors and their associated clients that form a relationship where the designated route reflectors provide route reflection for their clients, as well as nonclient peers.
community	A BGP attribute that contains a list of 32-bit values used to identify a route as belonging to a category of routes. All of the routes in the category are treated equally by routing policies.
dampen	Indicates that routes which exhibit instability are not advertised until the routes become stable for a minimum time period.
equal cost multipath (ECMP)	Distributes routing traffic among multiple equal-cost routes.
External BGP (EBGP)	A Border Gateway Protocol (BGP) used by routers that exchange information between two BGP speakers in different Autonomous Systems.
Interior BGP (IBGP)	Routers that use the Border Gateway Protocol (BGP) within an Autonomous System. The router redistributes BGP information to Interior Gateway Protocols (IGPs) that run in the autonomous path.
Interior Gateway Protocol (IGP)	Distributes routing information between routers that belong to a single Autonomous System (AS).
Internet Assigned Numbers Authority (IANA)	The central registry for various assigned numbers, for example, Internet protocol parameters (such as port, protocol, and enterprise numbers), options, codes, and types.
load balancing	The practice of splitting communication into two (or more) routes or servers.
Message Digest 5 (MD5)	A one-way hash function that creates a message digest for digital signatures.
multihomed AS	An AS that has multiple connections to one or more ASs and does not carry transit traffic.
next hop	The next hop to which a packet can be sent to advance the packet to the destination.
prefix	A group of contiguous bits, from 0 to 32 bits in length, that defines a set of addresses.
route reflector	A BGP speaker that advertises routes learned from its route reflector clients to other IBGP neighbors.

route reflector client	A BGP speaker that advertises its learned routes to a route reflector for readvertisement of its routes to the rest of the AS.
routing policy	A form of routing that is influenced by factors other than the default algorithmically best route, such as the shortest or quickest path.
transit AS	An AS that has multiple connections to one or more ASs and is used (with certain policy restrictions) to carry both transit and local traffic.
well-known attribute	A BGP attribute that is required to be known by all BGP implementations.

well-known attribute