



Configuration - BGP Services Avaya Ethernet Routing Switch 5000 Series

Release 6.3
NN47200-511
01.02
October 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

Avaya shall not be responsible for any modifications, additions, or deletions to the original published version of this documentation unless such modifications, additions, or deletions were performed by Avaya. End User agree 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(s) 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 this product. 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://www.avaya.com/support>. Please note that if you acquired the product 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, AND 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(s) and Product(s) provided by Avaya. All content on this site, the documentation(s) and the product(s) 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. 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://www.avaya.com/support/Copyright/>.

Trademarks

The trademarks, logos and service marks ("Marks") displayed in this site, the documentation(s) 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(s) 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 other trademarks are the property of their respective owners.

Downloading documents

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

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://www.avaya.com/support>

Contents

Chapter 1: Purpose of this document.....	5
Purpose of this document.....	5
Chapter 2: BGP Fundamentals.....	7
BGP fundamentals.....	7
Autonomous systems.....	8
Internal BGP routing.....	9
BGP speaker.....	9
Peers.....	9
Supernet advertisements.....	9
Bandwidth and maintenance reduction.....	10
Routing information consolidation.....	10
CIDR and aggregate addresses.....	10
Supernet addressing.....	12
Aggregate routes.....	14
Route reflectors.....	15
BGP communities.....	18
BGP path attributes.....	19
BGP route selection.....	19
BGP updates.....	21
Withdrawn Routes Length.....	21
Withdrawn Routes.....	22
Total Path Attributes Length.....	22
Path Attributes.....	22
Network Layer Reachability Information.....	25
Route policies.....	25
Equal-cost multipath.....	27
MD5 message authentication.....	28
MD5 signature generation.....	28
MD5 signature verification.....	29
BGP and route redistribution.....	29
Circuitless IP.....	30
BGP configuration considerations and limitations.....	30
BGP implementation guidelines.....	31
Configuration guidelines.....	31
BGP neighbor Maximum Prefix configuration.....	32
BGP and OSPF interaction.....	32
Chapter 3: BGP Configuration Using ACLI.....	33
Configuring BGP globally.....	33
Disabling BGP globally.....	38
Viewing BGP configuration.....	39
Viewing global BGP statistics.....	39
Configuring BGP peers or peer groups.....	40
Configuring a BGP peer or peer group password.....	43
Disabling BGP peers or peer groups.....	43

Deleting a BGP peer or peer group password.....	46
Viewing BGP peer information.....	46
Viewing BGP peer group information.....	47
Viewing a summary of BGP configurations.....	48
Configuring aggregate routes.....	48
Viewing BGP aggregate information.....	49
Configuring allowed networks.....	50
Viewing BGP network configurations.....	50
Configuring redistribution to BGP.....	51
Viewing BGP redistributed routes.....	52
Configuring prefix lists.....	52
Configuring route policies.....	53
Configuring AS path lists.....	53
Variable definitions.....	53
Deleting AS path lists.....	54
Viewing AS path information.....	54
Configuring community lists.....	55
Deleting community lists.....	56
Viewing community lists.....	57
Restarting BGP.....	57
Viewing CIDR routes.....	58
Viewing imported routes.....	58
Viewing BGP routes.....	59
Clearing BGP counters.....	60
Chapter 4: BGP Configuration Using EDM.....	61
Configuring BGP globally.....	61
Viewing BGP statistics.....	64
Configuring aggregate routes.....	64
Configuring allowed networks.....	65
Configuring BGP peers.....	66
Configuring peer groups.....	69
Viewing BGP summary route information.....	70
Configuring redistribution to BGP.....	71
Configuring a prefix list.....	72
Configuring route policies.....	74
Configuring an AS path list.....	74
Configuring a community access list.....	75

Chapter 1: Purpose of this document

Purpose of this document

This document describes the conceptual and procedural information to configure Border Gateway Protocol (BGP) services on Avaya Ethernet Routing Switches 5600 switches only.

Following operations are supported by BGP:

- Route Redistribution
- Circuitless IP (CLIP)
- Aggregation
- Route Reflection

Purpose of this document

Chapter 2: BGP Fundamentals

BGP fundamentals

Border Gateway Protocol (BGP) is an inter-domain routing protocol that provides loop-free inter-domain routing between autonomous systems (ASs) or within an AS. BGP systems can exchange network layer reachability information (NLRI) with other BGP systems for the purpose of constructing a graph of AS connectivity. BGP uses this information to prune routing loops and enforce AS-level policy decisions. BGP provides features that allow you to consolidate routing information and to control the flow of BGP updates.

The Ethernet Routing Switch 5600 supports a lightweight version of BGP (BGP Lite). With this new BGP support, the Ethernet Routing Switch 5600 can use only internal BGP (iBGP) to communicate within a single AS. The Ethernet Routing Switch 5600 does not support external BGP (eBGP) for communication between multiple external ASs, or full-table BGP Internet downloads.

A typical BGP Lite configuration is for the Ethernet Routing Switch 5600 to be configured to accept a default route (in the default Virtual Routing and Forwarding (VRF) instance, i.e. VRF 0 only) from a client, such as an internet service provider (ISP), and announce the local routes of the switch to the ISP.

BGP Lite supports the following BGP functions:

- iBGP (eBGP will be supported in a later software release)
- BGP Route Reflector
- BGP Aggregation
- BGP Redistribution
- BGP ECMP

The following BGP Lite feature restrictions exist with the Ethernet Routing Switch 5600:

- BGP Lite is supported only on ERS 5600 standalone and pure 5600 stacks.
- BGP Lite is supported in VRF0 only.
- BGP Lite supports a maximum of 4 BGP peers within a single AS.
- BGP Lite supports a maximum of 4000 BGP routes (4000 is the maximum route table size).

To use BGP, you must have an Advanced License. For more information about licensing, see *Avaya Ethernet Routing Switch 5000 Series Fundamentals, NN467200-104*.

The following sections provide an overview of BGP and includes descriptions of features you can use to optimize your BGP system.

For information about how to use the Avaya command line interface (ACLI) and the Web management interface, Enterprise Device Manager (EDM), see *Avaya Ethernet Routing Switch 5000 Series Fundamentals, NN47200-104*.

Autonomous systems

An autonomous system (AS) is a group of routers and hosts run by a single technical administrator that has a single, clearly defined routing policy. Each AS has its own unique AS number assigned by the appropriate Internet Registry entity.

The following figure shows a sample autonomous system.

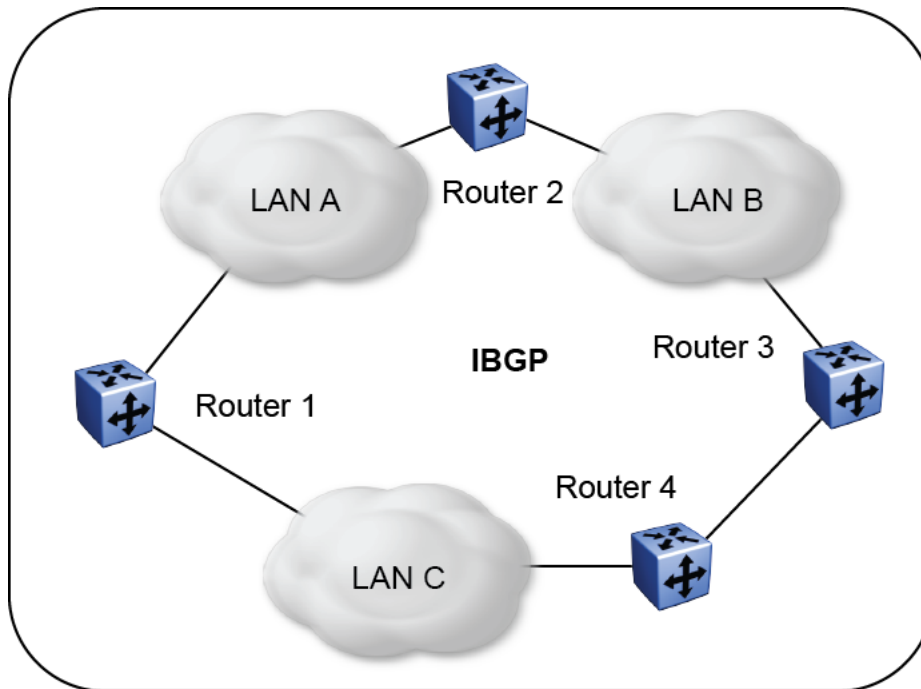


Figure 1: Autonomous system

BGP Lite exchanges information between routers within the same AS. In the preceding figure, routers that are members of the same AS run internal BGP (iBGP) to exchange BGP updates .

Related topics:

[Internal BGP routing](#) on page 9

[BGP speaker](#) on page 9

[Peers](#) on page 9

[Supernet advertisements](#) on page 9

[Bandwidth and maintenance reduction](#) on page 10

Internal BGP routing

With BGP Lite, Avaya supports Internal BGP (iBGP) intra-AS routing. Using iBGP, each router within an AS runs an interior gateway protocol (IGP), such as routing information protocol (RIP), and 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 for exchanging information within the internal AS.

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 communicate with other BGP speakers by establishing a peer-to-peer session.

Peers

The transport protocol used with BGP is Transmission Control Protocol (TCP). When any two routers open a TCP connection to each other for the purpose of exchanging routing information, they form a peer-to-peer relationship.

When an IGP is running that allows any two neighbors to logically communicate, 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 route reflectors.

BGP peers exchange complete routing information only after the peer connection is established. 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 are available, BGP compares the path attributes to choose the preferred path. For more information about update messages, see [BGP updates](#) on page 21.

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. For more information about CIDR, see [CIDR and aggregate addresses](#) on page 10.

Bandwidth and maintenance reduction

BGP also provides Route reflectors that reduce the high bandwidth and maintenance costs associated with a large full-mesh topology.

Route reflectors are discussed in the following sections.

Routing information consolidation

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

Related topics:

[CIDR and aggregate addresses](#) on page 10

[Supernet addressing](#) on page 12

[Aggregate routes](#) on page 14

[Route reflectors](#) on page 15

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. Classes D (used for multicast) and E (reserved and currently not used) are not discussed in this book.

For example, network 195.215.0.0, an illegal Class C network number, becomes a legal supernet when it is represented in CIDR notation as 195.215.0.0/16. The /16 is called the prefix length and becomes a way of expressing 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 (see the following table).

Table 1: 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	255.255.0.0	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
/18	255.255.192.0	1111 1111 1111 1111 1100 0000 0000 0000	64 Class C

Prefix	Dotted-decimal	Binary	Network class
/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 any address range.

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

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

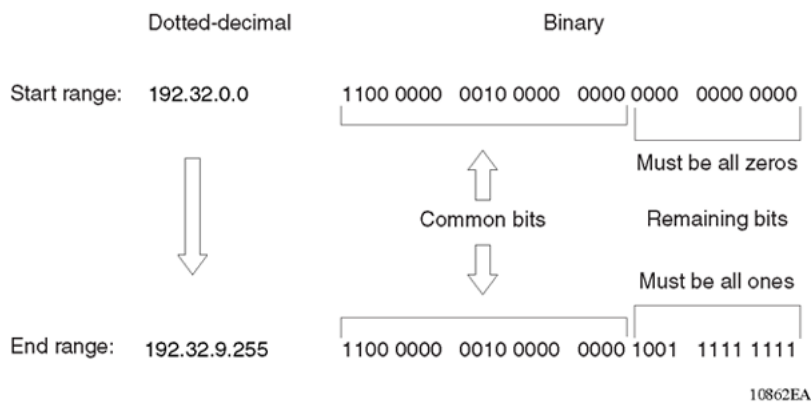
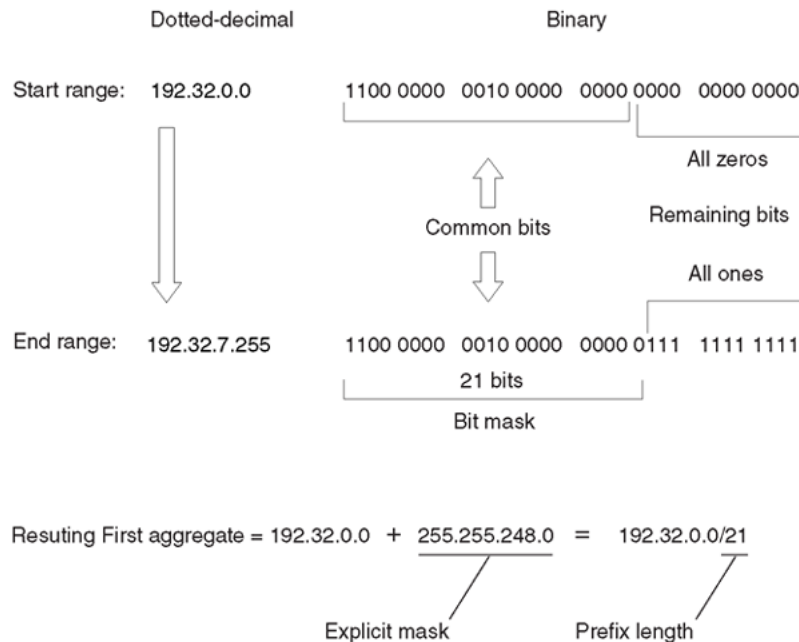


Figure 2: Binary notation conversion

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.

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 start range's common bits.



10863EA

Figure 3: 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 3: First aggregate and prefix length](#) on page 13, 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.

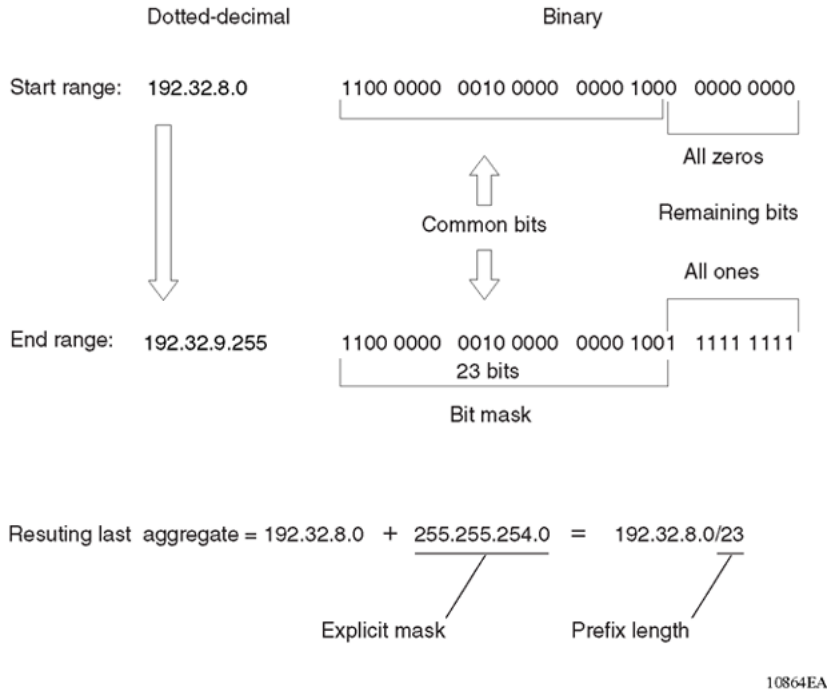


Figure 4: Last aggregate and prefix length

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 is used to advertise 256 separate Class C network numbers 195.215.0.0 through 195.215.255.0.

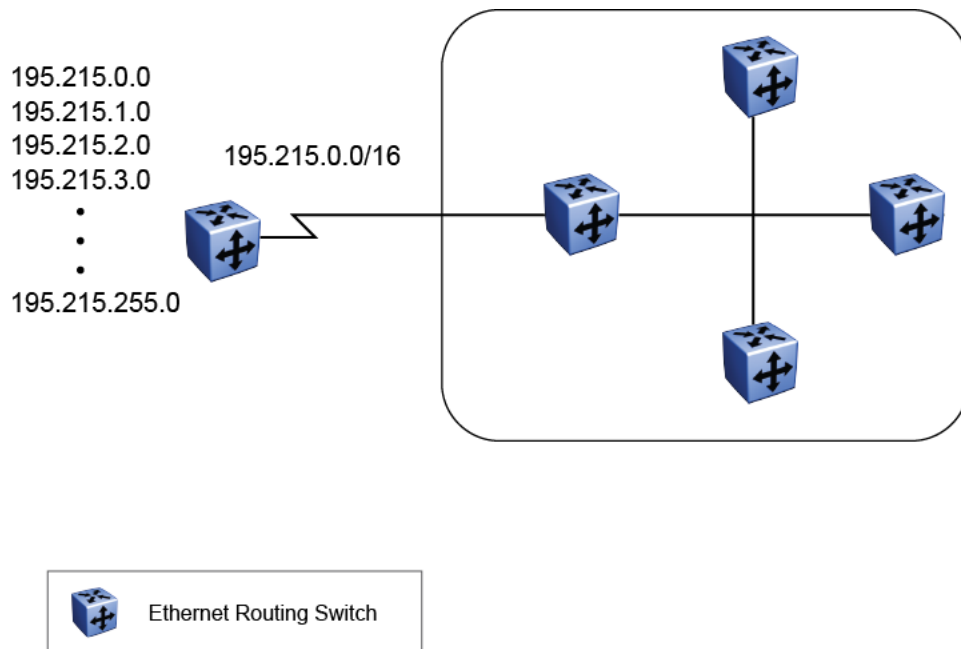


Figure 5: Aggregating routes with CIDR

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 (RR). 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 RR client session with the iBGP route reflector.

In an AS, there can be more than one route reflector cluster. There can also be more than one route reflector in a cluster. When there is more than one reflector in a cluster, special care must be taken to prevent route loops.

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





Figure 6: 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, thus avoiding 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 are not required to be fully meshed. In this case, the assigned route reflector assumes the responsibility for passing iBGP learned routes to a set of iBGP neighbors.


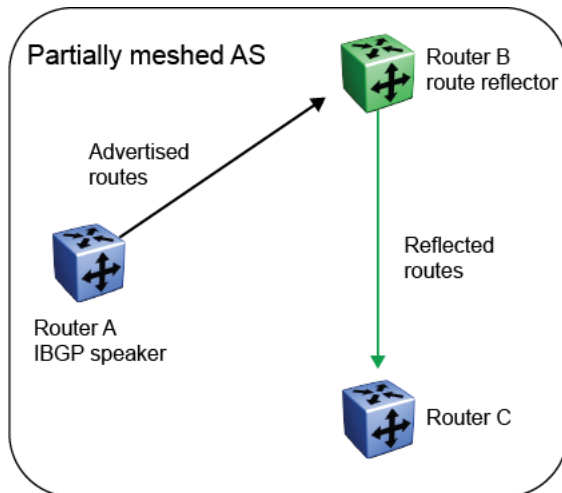


Figure 7: AS with route reflector

When Router B (the route reflector) receives routes advertised from Router A (the iBGP speaker) it advertises them to router C. Conversely, when the route reflector receives routes from internal peers, it advertises those routes to Router A. iBGP sessions are not required between Routers A and C.

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 are not required to be fully meshed, and do not communicate with iBGP speakers outside their cluster. Nonclient peers must be fully meshed with each other.

This concept is shown in the following figure, where Router A is shown as the route reflector in a cluster with client Routers B and C. Routers E and F are fully meshed, nonclient routers.

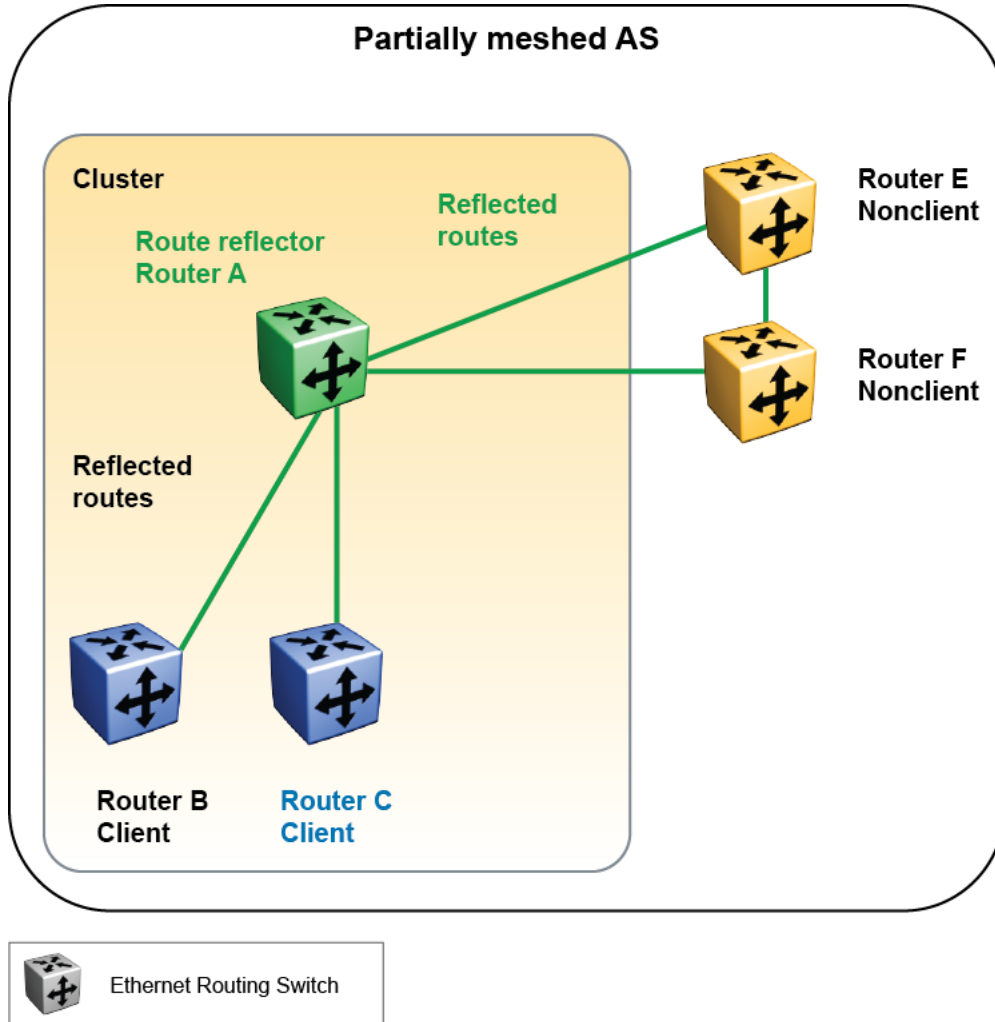


Figure 8: Route reflector with client and nonclient peers

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 a community to control routing policies with respect to destinations. It is common practice to create communities when you have more than one destination and want to share a common attribute.

The following are specific community types:

- Internet—advertise this route to the Internet community
- No Advertise—do not advertise to any BGP peer including iBGP peers

- No Export—do not advertise any destinations outside of a BGP confederation
- No Export SubConfed—do not advertise to external BGP peers, even within the same local AS.

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 specified community value is added to the existing value of the community attribute. Otherwise, the specified community value replaces any community value that was previously set.

BGP path attributes

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

There are four categories of BGP path attributes:

- Well known mandatory attributes must be included in every BGP update message.
- Well known discretionary attributes may or may not be sent in a particular 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 be passed along to other BGP peers.

Path attributes are used by border routers that utilize built-in algorithms or manually configured policies to select paths. 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)
- Multi-Exit Discriminator Attribute (optional non-transitive)
- Local Preference (well-known discretionary)
- Atomic Aggregate (well-known discretionary)
- Aggregator (optional transitive)
- Community (optional transitive)

BGP route selection

One of the most important responsibilities a BGP router performs is determining the best path to a given destination network. This path is then eligible for use in the router's IP forwarding

table. When choosing the best of multiple BGP routes to a given 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 and is 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. The preference is commonly manipulated using route policies to influence path selection.

- prefer locally originated paths

The path that was locally originated using the network, redistribution, or aggregate command is preferred over a path that was learned through a BGP Update. Local paths sourced by network or redistribute commands are preferred over local aggregates sourced by the aggregate address command.

- shortest AS Path

The AS Path parameter specifies the ASs that the network prefix has traversed. The AS Path is commonly used to determine the best path. For example, a router can choose a path based on whether the network passed through a given AS. A route policy can be configured to match the AS and modify the Local Preference. Also, the AS Path can be padded before it is advertised to a peer AS, so that the advertised network path is less likely to be preferred by downstream routers.

- lowest origin type

The origin type can be used to prefer a route. The order of preference is IGP, EGP, or incomplete (INC).

- lowest Multi-Exit Discriminator (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. 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.

Update received with no MED is assigned a MED of 0, unless the global BGP parameter (no-med-path-is-worst) is enabled. When enabled, BGP treats an update that is missing a MED attribute as the worst path. The parameter no-med-path-is-worst is enabled by default.

- lowest IGP metric to the BGP next-hop

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

- lowest Router ID

The lowest Router ID, or CLIP address, is preferred.

BGP updates

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

The following figure shows the format of an update message.




Figure 9: Update message format

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

Related topics:

[Withdrawn Routes Length](#) on page 21

[Withdrawn Routes](#) on page 22

[Total Path Attributes Length](#) on page 22

[Path Attributes](#) on page 22

[Network Layer Reachability Information](#) on page 25

Withdrawn Routes Length

The Withdrawn Routes Length parameter (referred to in RFC 1771 as the Unfeasible Routes Length field) indicates the total length of the Withdrawn Routes field in octets. The Withdrawn Routes Length field is used to calculate the length of the Network Layer Reachability Information field. For example, a value of 0 indicates that no routes are being 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 being withdrawn from service. The following figure shows the format of an IP prefix.



Figure 10: 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 (where: the network mask 255.255.0.0 is represented by the /16 which indicates the number of bits in the Length parameter).

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 is used to calculate the length of the Network Layer Reachability Information parameter. For example, a value of 0 indicates that no Network Layer Reachability Information field is present in this update message.

Path Attributes

The Path Attributes parameter is a variable-length sequence of path attributes that is present in every BGP Update. The path attributes contain BGP attributes that are associated with the prefixes in the Network Layer Reachability Information parameter.

For example, the attribute values allow you to specify the prefixes that can be exchanged in the BGP session, 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. The fields are described in the sections that follow.




Figure 11: Path attribute encoding

Related topics:

[Attribute Type](#) on page 23

[Attribute Length](#) on page 25

[Attribute Value](#) on page 25

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 12: Attribute Type fields

Attribute Flags

The Attribute Flags parameter is a bit string that contains four binary values that describe the attribute, and four bits that are unused. 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 set (1) the attribute is optional. When this bit is clear (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 to be present 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. It defines whether the information contained 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. It defines whether the Attribute Length is one octet (when set to 0) or two octets (when set to 1). Extended Length may be used 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. They must be zero (and must be ignored when received).

Attribute Type Code

The Attribute Type Code parameter contains the attribute type code, as defined by the Internet Assigned Numbers Authority (IANA). The Attribute Type Code is used to uniquely identify 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 supported Attribute Type Codes are shown in the following table.

Table 2: 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 path was advertised using an EGP by the last AS in the AS 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 must be traversed to reach the given destinations. Each AS path segment is represented 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).

Attribute	Type code	Description
Multiexit discriminator	4	This attribute is used on external (internal-AS) links to discriminate among multiple exit or entry points to the same neighboring AS.
Local preference	5	Indicates the preference that AS border routers assign to a chosen route when advertising it to iBGP peers
Atomic aggregate	6	Ensures that certain network layer reachability information (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 and is implemented according to the values in the Attribute Flags and the Attribute Type Code parameters.

Network Layer Reachability Information

The Network Layer Reachability Information parameter is a variable length field that contains a list of prefixes. The number of prefixes in the list is limited only by the packet size that can be sent between BGP speakers.

Route policies

BGP route maps also referred to as route policies is an important functionality of BGP. You statically configure these policies in each BGP speaker. These policies are flexible enough to

cause changes in the attributes of a route learned from a BGP peer, thereby affecting the decision process for the route. Both input policies and output policies are supported in BGP.

Input policy

An input policy applies on the ingress direction on receiving a BGP update. Input policies are supported on a peer basis, i.e. an input policy is associated to a specific BGP peer. An input policy can also be associated with multiple BGP peers as well. Each input policy can have a set of MATCH and/or SET criteria.

MATCH criteria is applied to various parameters in a BGP update to decide whether to allow the update for further processing or drop the update. The following attributes will be supported in MATCH criteria:

- AS Path list
- IP Prefix list
- Next Hop
- Local Preference
- Multi-Exit Discriminator (MED)
- Community
- Peer Address

SET criteria are used to modify the BGP updates with the corresponding values configured in the criteria. As SET criteria are applied as the first step in the BGP update processing, it would have impact on the routes selection in the BGP decision process. The updates can be one of the following:

- addition of new attributes (e.g. metric) along with the other path attributes received in the update message
- replacement of an existing path attribute received in the update message (e.g. replacement of next hop attribute in a BGP update)
- appending to an existing path attribute (e.g. appending to AS path attribute in a BGP update)

The following attributes are supported in SET criteria:

- AS Path
- Next Hop
- Local Preference
- MED
- Community

Output policy

An output policy applies on the egress direction before transmitting a BGP update to a peer. Output policies are supported on a peer basis, i.e. an output policy is associated to a specific BGP peer. An output policy can also be associated with multiple BGP peers as well. Each output policy can have a set of MATCH and/or SET criteria.

MATCH criteria is applied to various parameters in a BGP update to decide whether to allow the update for further transmission or block the update from being transmitted to the peer. The following attributes will be supported in MATCH criteria:

- AS Path list
- IP Prefix list
- Next Hop
- Local Preference
- Multi-Exit Discriminator (MED)
- Community
- Peer Address

SET criteria are used to modify the BGP updates with the corresponding values configured in the criteria. As SET criteria are applied to BGP updates before being transmitted, the path attributes in the actual BGP updates being transmitted would have changes based on the policy. The updates can be one of the following:

- addition of new attributes (e.g. metric) along with the other path attributes received in the update message
- replacement of an existing path attribute received in the update message (e.g. replacement of next hop attribute in a BGP update)
- appending to an existing path attribute (e.g. appending to AS path attribute in a BGP update)

The following attributes are supported in SET criteria:

- AS Path
- Next Hop
- Local Preference
- MED
- Community

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 Ethernet Routing Switch 5000 Series Configuration — IP Routing, NN47200-503*.

MD5 message authentication

The switch authenticates BGP messages by using Message Digest 5 (MD5) signatures. When 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 some BGP peers configured with authentication enabled and other BGP peers with authentication disabled.
- The secret keys are always stored encrypted. When you enable BGPv4 TCP MD5 authentication, the router computes an MD5 signature for each TCP packet based on the TCP packet and a per-peer secret key. The router adds this MD5 signature to the TCP packet containing 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 trying to masquerade as the sender, however, cannot generate an authentic signature because it does not know the secret key.

In the CLI 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 secret keys are stored in encrypted form on the switch.

Related topics:

[MD5 signature generation](#) on page 28

[MD5 signature verification](#) on page 29

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 (MSS) by 18 octets, which is the length of the TCP MD5 option. It also 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

As shown in the following table, 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 3: 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
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. When a packet has passed 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 be sent through a BGP network, then configure redistribution of OSPF routes through BGP. This sends OSPF routes to a router that uses BGP.

The Ethernet Routing Switch 5600 supports route redistribution between BGP, RIP, OSPF, and between direct and static routes. For more information about RIP and OSPF route redistribution, see *Avaya Ethernet Routing Switch 5000 Series Configuration — IP Routing, NN47200-503*. For information about configuring route redistribution, see [Configuring redistribution to BGP](#) on page 51 for ACLI commands, or [Configuring redistribution to BGP](#) on page 71 for EDM.

Circuitless IP

Circuitless IP (CLIP) is a virtual (or loopback) interface that is not associated with a physical port. You can use a CLIP interface to provide uninterrupted connectivity to your switch as long as there is an actual path 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).

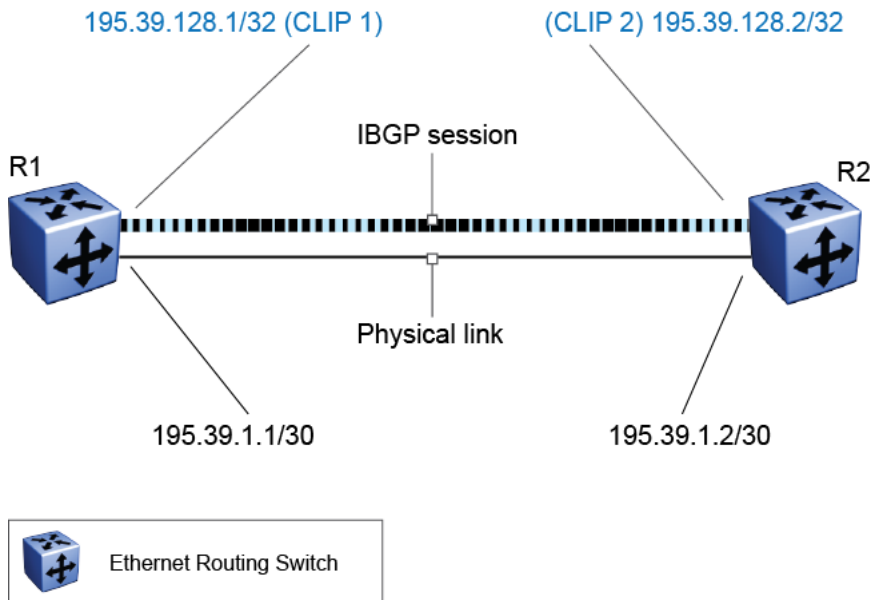


Figure 13: Routers with iBGP connections

For information about configuring CLIP interfaces, see *Avaya Ethernet Routing Switch 5000 Series Configuration — IP Routing, NN47200-503*.

BGP configuration considerations and limitations

Use the information in this section to help you configure BGP on the Ethernet Routing Switch 5600.

Related topics:

- [BGP implementation guidelines](#) on page 31
- [Configuration guidelines](#) on page 31
- [BGP neighbor Maximum Prefix configuration](#) on page 32
- [BGP and OSPF interaction](#) on page 32

BGP implementation guidelines

To successfully configure BGP, follow these guidelines:

- You must enable IP routing in the switch for BGP to work.
- If the BGP router ID is not configured, the switch 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 router's circuitless (virtual) IP interface as the local peer address. In this way, you ensure that BGP is reachable as long as there is an active circuit 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.

Configuration guidelines

On the Ethernet Routing Switch 5600, you must configure the following minimum parameters:

- Router ID
- Local AS Number
- Enable BGP globally
- BGP Neighbor Peer Session: remote IP addresses
- Enable BGP peer

BGP policies can be added to the BGP peer configuration to influence route decisions. BGP policies are applied to the peer through ACLI soft-reconfiguration commands.

After the Ethernet Routing Switch 5600 is configured for BGP, some parameter changes may require the BGP global state or the neighbor admin-state to be disabled or enabled.

BGP policies are dynamically modified. On the global level, the BGP redistribution command has an apply parameter that causes the policy to be applied when it is issued.

BGP neighbor Maximum Prefix configuration

By default, the Maximum Prefix parameter is set to limit 4000 network layer reachability information (NLRI) messages per neighbor. The Maximum Prefix parameter limits the number of routes that the Ethernet Routing Switch 5600 can accept.

The Maximum Prefix parameter prevents large numbers of BGP routes from flooding the network in the event of a misconfiguration. You can configure the Maximum Prefix limit to any value, including 0 (0 means a maximum of 4000 prefixes). When you configure the Maximum Prefix value, consider the maximum number of active routes that your equipment configuration can support.

BGP and OSPF interaction

RFC 1745 defines the interaction between BGP and OSPF when OSPF is the IGP within an autonomous system. A BGP route policy must be configured to allow BGP advertisement of OSPF routes.

Chapter 3: BGP Configuration Using ACLI

Configuring BGP globally

Use the following procedure to configure BGP globally to provide loop-free inter-domain routing within an AS.

Before you begin

- For initial BGP configuration, you must know the AS number.

Procedure

1. Log on to ACLI in Global Configuration mode.
2. At the command prompt, enter the following command:

```
router bgp
```

Next steps

Use the data in the following table to configure BGP as required.

Variable definitions

Use the data in the following table to configure BGP.

Variable	Value
[default] [no] aggregate-address <a.b.c.d/len> [as-set] [summary-only]	Adds an aggregate address in a BGP routing table. <ul style="list-style-type: none">• <a.b.c.d/len>> is an IP address and an integer value (between 0 and 32)• <i>as-set</i> enables autonomous system information. The default value is disable.• <i>summary-only</i> enables the summarization of routes not included in routing updates. This parameter creates the aggregate route and suppresses advertisements of more specific routes to all neighbors. The default value is disable.

Variable	Value
	<ul style="list-style-type: none"> • <i>default</i> configures an aggregate address in a BGP routing table to default values. • <i>no</i> disables or deletes an aggregate address in a BGP routing table.
[default] [no] auto-peer-restart enable	<p>Enables the process that automatically restarts a connection to a BGP neighbor. The default value is enable.</p> <ul style="list-style-type: none"> • <i>default</i> enables the process that automatically restarts a connection to a BGP neighbor. • <i>no</i> disables the process that automatically restarts a connection to a BGP neighbor.
[default] [no] auto-summary	<p>Enables BGP to summarize networks based on class limits (For example, Class A, B, C networks). The default value is enable.</p> <ul style="list-style-type: none"> • <i>default</i> enables BGP to summarize networks based on class limits. • <i>no</i> disables BGP from summarizing networks based on class limits.
[default] [no] bgp aggregation enable	<p>Enables the aggregation feature on this interface. The default value is enable. You cannot change the value when BGP is enabled.</p> <ul style="list-style-type: none"> • <i>default</i> enables the aggregation feature on this interface. • <i>no</i> disables the aggregation feature on this interface.
[default] [no] bgp always-compare-med	<p>Enables the comparison of the multi-exit discriminator (MED) parameter for paths from neighbors in different autonomous systems. A path with a lower MED is preferred over a path with a higher MED. The default value is disable.</p> <ul style="list-style-type: none"> • <i>default</i> disables the comparison of the multi-exit discriminator (MED) parameter for paths from neighbors in different autonomous systems. • <i>no</i> disables the comparison of the multi-exit discriminator (MED) parameter for paths from neighbors in different autonomous systems.
[default] [no] bgp client-to-client reflection	<p>Enables route reflection between two route reflector clients. This option is applicable only if the route reflection value is set to enable. The default value is enable. Route reflection may be enabled even when</p>

Variable	Value
	<p>clients are fully meshed. In this event, route reflection is not required.</p> <ul style="list-style-type: none"> • <i>default</i> enables route reflection between two route reflector clients. • <i>no</i> disables route reflection between two route reflector clients.
[no] bgp cluster-id <A.B.C.D>	<p>Sets a cluster ID. This option is applicable only if the route reflection value is set to enable, and if multiple route reflectors are in a cluster. <A.B.C.D> is the cluster ID of the reflector router.</p> <ul style="list-style-type: none"> • <i>no</i> deletes a cluster ID.
[default] [no] bgp default local-preference <0-2147483647>	<p>Specifies the default value of the local preference attribute. The default value is 100. You cannot change the default value when BGP is enabled.</p> <ul style="list-style-type: none"> • <i>default</i> configures the default local preference attribute value to 100. • <i>no</i> disables the default local preference attribute.
[default] bgp multiple-paths <1-4>	<p>Sets the maximum number of equal-cost-paths that are available to a BGP router by limiting the number of equal-cost-paths that can be stored in the routing table. The default value is 1.</p> <ul style="list-style-type: none"> • <i>default</i> configures the maximum number of equal-cost-paths that are available to a BGP router to 1.
[default] [no] debug-screen <off on>	<p>Displays debug messages on the console. Disable BGP screen logging (off) or enable BGP screen logging (on). The default is off.</p> <ul style="list-style-type: none"> • <i>default</i> configures BGP debug screen logging to default (off). • <i>no</i> disables the logging of BGP debug messages on the console.
[default] default-information originate	<p>Enables the advertisement of a default route to peers, if it is present in the routing table. The default value is disable.</p> <ul style="list-style-type: none"> • <i>default</i> disables the advertisement of a default route to peers, if it is present in the routing table.
[default] [no] default-metric <val>	<p>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 -1.</p>

Variable	Value
	<ul style="list-style-type: none"> • <i>val</i> is an integer value between -1 to 2147483647 • <i>default</i> resets the default metric value to 0 • <i>no</i> resets the default metric value to 0
<p>[default] [no] global-debug mask <value></p>	<p>Displays specified debug information for BGP global configurations. The default value is none.</p> <ul style="list-style-type: none"> • <value> is a list of mask choices separated by commas with no space between choices. <p>Mask choices are: <i>none</i> disables all debug messages. <i>all</i> enables all debug messages. <i>error</i> enables display of debug error messages. <i>packet</i> enables display of debug packet messages. <i>event</i> enables display of debug event messages. <i>trace</i> enables display of debug trace messages. <i>warning</i> enables display of debug warning messages. <i>state</i> enables display of debug state transition messages. <i>init</i> enables display of debug initialization messages. <i>filter</i> enables display of debug messages related to filtering. <i>update</i> enables display of debug messages related to sending and receiving updates.</p> <ul style="list-style-type: none"> • <i>default</i> resets/clears the global debug mask to none. • <i>no</i> resets/clears the global debug mask to none.
<p>[default] [no] ibgp-report-import-rt enable</p>	<p>Configures BGP to advertise imported routes to an interior BGP (iBGP) peer. This command also enables or disables advertisement of nonBGP imported routes to other iBGP neighbors. The default value is enable.</p> <ul style="list-style-type: none"> • <i>default</i> configures the default value (enable) for BGP to advertise imported routes to an interior BGP (iBGP) peer. • <i>no</i> disables BGP from advertising imported routes to an interior BGP (iBGP) peer.
<p>[default] [no] ignore-illegal-rtrid enable</p>	<p>Enables BGP to overlook an illegal router ID. For example, you can enable BGP to accept a connection from a peer that sends an open message using a router ID of 0 (zero). The default value is enable.</p>

Variable	Value
	<ul style="list-style-type: none"> • <i>default</i> enables BGP to overlook an illegal router ID. • <i>no</i> disables BGP from overlooking an illegal router ID.
[default] [no] neighbor <nbr_ipaddr peer-group-name>	Creates and manages BGP peers and peer groups. For more information, see Configuring BGP peers or peer groups on page 40.
neighbor-debug-all mask <value>	<p>Displays specified debug information for BGP neighbors. The default value is none.</p> <ul style="list-style-type: none"> • <value> 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"> <i>none</i> disables all debug messages. <i>all</i> enables all debug messages. <i>error</i> enables display of debug error messages. <i>packet</i> enables display of debug packet messages. <i>event</i> enables display of debug event messages. <i>trace</i> enables display of debug trace messages. <i>warning</i> enables display of debug warning messages. <i>state</i> enables display of debug state transition messages. <i>init</i> enables display of debug initialization messages. <i>filter</i> enables display of debug messages related to filtering. <i>update</i> enables display of debug messages related to sending and receiving updates.
[default] [no] neighbor-debug-all	<ul style="list-style-type: none"> • <i>default</i> sets BGP neighbor debug information to the default mask value (none) • <i>no</i> sets BGP neighbor debug information to the default mask value (none).
[default] [no] network <a.b.c.d/len>	<p>Specifies IGP network prefixes for BGP to advertise for redistribution. This command imports routes into BGP.</p> <ul style="list-style-type: none"> • <a.b.c.d/len> the network address and mask. • <i>default</i> configures BGP to advertise the IGP network prefix with metric 0 for redistribution. • <i>no</i> disables BGP from advertising IGP network prefixes for redistribution.

Variable	Value
[default] [no] no-med-path-is-worst enable	Enables BGP to treat an update without a multi-exit discriminator (MED) attribute as the worst path. The default value is disable. <ul style="list-style-type: none"> • <i>default</i> sets no-med-path-is-worst command to default (disable) • <i>no</i> disables no-med-path-is-worst command
[default] [no] redistribute <direct ospf rip static>	Configures a redistribute entry to announce routes of a certain source protocol type into BGP. <ul style="list-style-type: none"> • <i>default</i> sets redistribute command to default. • <i>no</i> disables/deletes the BGP redistribute command.
no neighbor <nbr_ipaddr peer-group-name>	Disables BGP peers and peer groups. For more information, see Disabling BGP peers or peer groups on page 43.
[default] [no] quick-start enable	Enables the quick-start flag for exponential backoff. <ul style="list-style-type: none"> • <i>default</i> disables the quick-start flag. • <i>no</i> disables the quick-start flag.
[default] [no] route-reflector enable	Enables the reflection of routes from iBGP neighbors. The default value is enable. <ul style="list-style-type: none"> • <i>default</i> enables the reflection of routes from iBGP neighbors. • <i>no</i> disables the reflection of routes from iBGP neighbors.
[default] [no] router-id <A.B.C.D>	Specifies the BGP router ID in IP address format. <ul style="list-style-type: none"> • <A.B.C.D> is the BGP router IP address. • <i>default</i> resets the BGP router ID to default. • <i>no</i> resets the BGP router ID to default.
[default] [no] traps enable	Enables BGP traps. <ul style="list-style-type: none"> • <i>default</i> disables BGP traps. • <i>no</i> disables BGP traps.

Disabling BGP globally

Use the following procedure to disable BGP globally to discontinue providing loop-free inter-domain routing within an AS.

Procedure

1. Log on to ACLI in Global Configuration mode.
 2. At the command prompt, enter the following command:
`no router bgp enable`
-

Viewing BGP configuration

Use this procedure to view information about the BGP configuration.

Procedure

1. Log on to ACLI in Privileged EXEC mode.
 2. At the command prompt, enter the following command:
`show ip bgp conf`
-

Viewing global BGP statistics

Use this procedure to view global BGP statistics.

Procedure

1. Log on to ACLI in Privileged EXEC mode.
 2. At the command prompt, enter the following command:
`show ip bgp stats`
-

Configuring BGP peers or peer groups

Use the following procedure to configure BGP peers or peer groups to assemble neighbors with the same update policies into peer groups and peer associations. You can also use this procedure to configure default values for BGP neighbor parameters.

Before you begin

- If required, route policies exist.

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. Log on to ACLI in BGP Router Configuration mode.
 2. At the command prompt, enter the following command:

```
[default] neighbor <nbr_ipaddr|peer-group-name>
```
-

Variable definitions

The following table defines optional parameters that you enter with the **neighbor** `<nbr_ipaddr|peer-group-name>` command.

Variable	Value
default-originate	Enables the switch to send a default route advertisement to the specified neighbor or peer group. A default route does not have to be in the routing table. The default value is disable. Do not use this command if default-information originate is globally enabled.
enable	Enables the BGP neighbor.

Variable	Value
in-route-map <route policyname>	<p>Applies a route policy rule to all incoming routes that are learned from the local BGP router's peers, or peer groups. The local BGP router is the BGP router that allows or disallows routes and sets attributes in incoming updates.</p> <p><route policyname> is an alphanumeric string length (1 to 64 characters) that indicates the name of the route map or policy.</p>
max-prefix <0-4000>	<p>Sets a limit on the number of routes that can be accepted from a neighbor. The default value is 4000 routes and the range is 0 to 4000 (0 means a maximum of 4000 prefixes).</p>
MD5-authentication enable	<p>Enables TCP MD5 authentication between two peers. The default value is disable.</p>
neighbor-debug mask <value>	<p>Displays specified debug information for a BGP peer. The default value is none.</p> <p><value> 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"> <i>none</i> disables all debug messages. <i>all</i> enables all debug messages. <i>error</i> enables display of debug error messages. <i>packet</i> enables display of debug packet messages. <i>event</i> enables display of debug event messages. <i>trace</i> enables display of debug trace messages. <i>warning</i> enables display of debug warning messages. <i>state</i> enables display of debug state transition messages. <i>init</i> enables display of debug initialization messages. <i>filter</i> enables display of debug messages related to filtering. <i>update</i> enables display of debug messages related to sending and receiving updates.
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 is generating the iBGP update. The default value is disable.</p> <p>The next-hop parameter can only be configured when the neighbor is disabled.</p>
out-route-map <route policyname>	<p>Applies a route policy rule to all outgoing routes that are sent to the local BGP router's peers, or peer groups. The local BGP router is the BGP router that allows or disallows routes and sets attributes in outgoing updates.</p>

Variable	Value
	<route policyname> is an alphanumeric string length (1 to 64 characters) that indicates the name of the route map or policy.
peer-group <grpname>	Adds a BGP peer to the specified subscriber group. You must create the specified subscriber group before you issue this command.
remote-as <0-65535>	Configures the remote AS number of a BGP peer or a peer-group. You cannot configure this option when the admin-state is enable.
retry-interval <1-65535>	Sets the time interval (in seconds) for the ConnectRetry Timer. The default value is 120 seconds.
route-reflector-client	Configures the specified neighbor or group of neighbors as its route reflector client. The default value is disable. All neighbors that are configured become members of the client group and the remaining iBGP peers become members of the nonclient group for the local route reflector.
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	When enabled, the router relearns routes from the specified neighbor or group of neighbors without resetting the connection when the policy changes in the inbound direction. The default value is disable.
timers <KeepAliveTime><HoldTime>	Configures timers (in seconds) for the BGP speaker for this peer. <ul style="list-style-type: none"> • <KeepAliveTime> integer value between 0 to 21845. • <HoldTime> integer value 0 or between 3 to 65535.
update-source <A.B.C.D>	Specifies the source IP address when BGP packets are sent to this peer or peer group. Change does not take effect until neighbor restarts. <A.B.C.D> is the specified source IP address.
weight <0-65535>	Specifies the weight of a BGP peer or peer groups, or the priority of updates that can be received from that BGP peer. The default value is 0. If you have particular neighbors that you want to prefer for most of your traffic, you can assign a higher weight to all routes learned from that neighbor.

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.

Procedure

1. Log on to ACLI in BGP Router Configuration mode.
 2. At the command prompt, enter the following command:

```
neighbor password <nbr_ipaddr|peer-group-name> <passwd>
```
 3. To enable MD5 authentication, enter the following command:

```
neighbor <nbr_ipaddr|peer-group-name> md5-authentication enable
```
-

Variable definitions

The following table defines parameters for the `neighbor password <nbr_ipaddr|peer-group-name>` command.

Variable	Value
<passwd>	Specifies an alphanumeric string length from 1 to 80 characters.

Disabling BGP peers or peer groups

Use the following procedure to disable BGP neighbor peer groups, to delete neighbors from peer groups and peer associations or to discontinue using specific BGP neighbor parameters.

Procedure

1. Log on to ACLI in BGP Router Configuration mode.
 2. At the command prompt, enter the following command:

```
no neighbor <nbr_ipaddr|peer-group-name>
```
-

Variable definitions

The following table defines optional parameters that you enter after the `no neighbor <nbr_ipaddr | peer-group-name>` command.

Variable	Value
default-originate	Prevents the switch from sending a default route advertisement to the specified neighbor. A default route does not have to be in the routing table. The default value is disable.
enable	Disables the BGP neighbor.
in-route-map	Discontinues the application of a route policy rule to all incoming routes that are learned from, or sent to, the local BGP router's peers, or peer groups. The local BGP router is the BGP router that allows or disallows routes and sets attributes in incoming or outgoing updates.
MD5-authentication enable	Disables TCP MD5 authentication between two peers. The default value is disable.
neighbor-debug-mask <WORD 1-100>	<p>Discontinues displaying specified debug information for a BGP peer.</p> <p><WORD 1-100> is a list of mask choices separated by commas with no space between choices. For example: <i>{<mask>,<mask>,<mask>...}</i>.</p> <p>Mask choices are:</p> <ul style="list-style-type: none"> <i>none</i>: disables all debug messages. <i>all</i>: enables all debug messages. <i>error</i>: enables display of debug error messages. <i>packet</i>: enables display of debug packet messages. <i>event</i>: enables display of debug event messages. <i>trace</i>: enables display of debug trace messages. <i>warning</i>: enables display of debug warning messages. <i>state</i>: enables display of debug state transition messages. <i>init</i>: enables display of debug initialization messages.

Variable	Value
	<p><i>filter</i>: enables display of debug messages related to filtering.</p> <p><i>update</i>: enables display of debug messages related to sending and receiving updates.</p>
next-hop-self	<p>Discontinues specifying that the next-hop attribute in an iBGP update is the address of the local router or the router that is generating the iBGP update. The default value is disable.</p> <p>The next-hop parameter can only be configured when the neighbor is disabled.</p>
out-route-map	<p>Discontinues applying a route policy rule to all outgoing routes that are learned from, or sent to, the local BGP router's peers, or peer groups. The local BGP router is the BGP router that allows or disallows routes and sets attributes in incoming or outgoing updates.</p>
peer-group	<p>Deletes a BGP neighbor peer group.</p>
remote-as	<p>Deletes the BGP peer or a peer-group remote AS.</p>
route-reflector-client	<p>Disables the specified neighbor or neighbors route reflector clients. The default value is disable.</p>
send-community	<p>Disables the switch from sending the update message community attribute to the specified peer. The default value is disable.</p>
soft-reconfiguration-in enable	<p>When disabled, the router does not relearn routes from the specified neighbor or group of neighbors without resetting the connection when the policy changes in the inbound direction. The default value is disable.</p>
update-source	<p>Disables specifying the source IP address when BGP packets are sent to this peer or peer group. You cannot configure this parameter when the admin-state is enable.</p>
weight	<p>Disables specifying the weight of a BGP peer or peer groups, or the priority of updates that can be received from that BGP peer.</p>

Deleting a BGP peer or peer group password

Use this procedure to delete a BGP peer or peer group password to discontinue using TCP MD5 authentication between two peers.

Procedure

1. Log on to ACLI in BGP Router Configuration mode.
 2. At the command prompt, enter the following command:


```
no neighbor <nbr_ipaddr|peer-group-name> md5-authentication
enable
```

OR

```
default neighbor <nbr_ipaddr|peer-group-name> md5-
authentication enable
```
-

Variable definitions

The following table defines parameters for the `[no] [default] neighbor <nbr_ipaddr|peer-group-name>` command.

Variable	Value
<nbr_ipaddr>	Specifies the IP address.
<peer-group-name>	Specifies an alphanumeric string length from 1 to 80 characters.
md5-authentication enable	Disables md5 authentication.

Viewing BGP peer information

Use this procedure to display information about BGP peer configuration.

Procedure

1. Log on to ACLI in Privileged EXEC mode.
2. To display information about BGP peers, enter the following command:


```
show ip bgp neighbors [{A.B.C.D}]
```

- To display information about routes advertised to the neighbor, enter the following command:

```
show ip bgp neighbors {A.B.C.D} advertised-routes [<prefix/len>] [longer-prefixes]
```

- To display information about routes accepted from the neighbor, enter the following command:

```
show ip bgp neighbors {A.B.C.D} routes [<prefix/len>] [community <enable|disable>] [longer-prefixes]
```

- To display statistics for BGP peer, enter the following command:

```
show ip bgp neighbors {A.B.C.D} stats
```

Variable definitions

The following table defines optional parameters that you can enter with the `show ip bgp neighbors <A.B.C.D>` command.

Variable	Value
{A.B.C.D}	Specifies the IP address.
community <enable disable>	Enables or disables the display of community attributes.
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).
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/0 to A.B.C.D/32).

Viewing BGP peer group information

Use this procedure to display information about BGP peer groups.

Procedure

- Log on to ACLI in Privileged EXEC mode.
- At the command prompt, enter the following command:

```
show ip bgp peer-group [WORD]
```

Variable definitions

The following table defines optional parameters that you can enter with the `show ip bgp peer-group` command.

Variable	Value
<WORD>	Specifies the peer group name, an alphanumeric character string ranging from 1 to 64 characters.

Viewing a summary of BGP configurations

Use this procedure to display summarized information about BGP.

Procedure

1. Log on to ACLI in Privileged EXEC mode.
 2. At the command prompt, enter the following command:
`show ip bgp summary`
-

Configuring aggregate routes

Use this procedure to configure aggregate routes that allow the router to advertise a single route (aggregate route) that represents all destinations.

Before you begin

- Disable BGP before you enable aggregation.
- You need the appropriate aggregate address and mask.

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. Log on to ACLI in BGP Router Configuration mode.
2. At the command prompt, enter the following command:


```
bgp aggregation enable
```

3. Add an aggregate route to the routing table:

```
aggregate-address <prefix/len> [as-set][summary-only]
```

4. Exit to Global Configuration mode:

```
exit
```

5. Enable BGP:

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

Variable definitions

The following table defines parameters for the **aggregate-address** command.

Variable	Value
as-set	Enables autonomous system information. The default value is disable.
<prefix/len>	Specifies an IP address and mask in the form of A.B.C.D/<0-32>.
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.

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.

Viewing BGP aggregate information

About this task

Use this procedure to display information about current aggregate addresses.

Procedure

1. Log on to ACLI in Privileged EXEC mode.
 2. At the command prompt, enter the following command:

```
show ip bgp aggregates [A.B.C.D/<0-32>]
```
-

Configuring allowed networks

Use this procedure to configure BGP allowed networks to determine the network IP addresses that BGP advertises.

The allowed addresses determine the BGP networks that originate from the Ethernet Routing Switch 5000.

Procedure

1. Log on to ACLI in BGP Router Configuration mode.
 2. At the command prompt, enter the following command:

```
network <prefix/len> [metric <0-65535>]
```
-

Variable definitions

The following table defines parameters for the **network** command.

Variable	Value
<prefix/len>	Specifies the network address and mask in the form A.B.C.D/<0-32>
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 peers. The range is 0-65535.

Viewing BGP network configurations

Use this procedure to display information about BGP network configurations.

Procedure

1. Log on to ACLI in Privileged EXEC mode.
 2. At the command prompt, enter the following command:

```
show ip bgp networks [A.B.C.D/<0-32>]
```
-

Configuring redistribution to BGP

Use this procedure to configure a redistribute entry to announce routes of a certain source protocol type into the BGP domain.

Before you begin

- If required, a route policy exists.

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. Log on to ACLI in BGP Router Configuration mode.
 2. To create a redistribution instance, enter the following command:

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

```
redistribute <direct|ospf|rip|static> route-policy <route-policy-name>
```
 4. If required, configure the route metric:

```
redistribute <direct|ospf|rip|static> metric <0-65535>
```
 5. Enable the instance:

```
redistribute <direct|ospf|rip|static> enable
```
 6. Exit BGP Router Configuration mode:

```
exit
```
 7. Apply the redistribution instance configuration:

```
ip bgp apply redistribute <direct|ospf|rip|static>
```
-

Variable definitions

The following table defines parameters for the `redistribute` and `ip bgp apply redistribute` commands.

Variable	Value
<static ospf 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-policy <route-policy-name>	Configures the route policy to apply to redistributed routes, an alphanumeric string up to 64 characters.

Viewing BGP redistributed routes

Use this procedure to display information about BGP redistributed routes.

Procedure

1. Log on to ACLI in Privileged EXEC mode.
 2. At the command prompt, enter the following command:

```
show ip bgp redistribute
```
-

Configuring prefix lists

You can configure prefix lists to allow or deny specific BGP route updates. A prefix list policy specifies route prefixes to match. When there is a match, the route is used.

For information about configuring prefix lists, see Avaya Ethernet Routing Switch 5000 Series Configuration — IP Routing NN47200–503.

Configuring route policies

For information about configuring route policies, see Avaya Ethernet Routing Switch 5000 Series Configuration — IP Routing NN47200–503.

Configuring AS path lists

Use this procedure to 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. Log on to ACLI in Global Configuration mode.
2. At the command prompt, enter the following command:


```
ip as-list <as-list-id> memberid <member-id> <permit|deny>
as-path <as-path-string>
```

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

Variable definitions

The following table defines parameters for the `ip as-list` command.

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

Deleting AS path lists

Use this procedure to delete an AS path list

Procedure

1. Log on to ACLI in Global Configuration mode.
 2. At the command prompt, enter the following command:

```
no ip as-list <as-list-id>
```
-

Variable definitions

The following table defines parameters for the `no ip as-list` command.

Variable	Value
<as-list-id>	Specifies an integer value between 1 and 1024 that represents the AS-path list ID you want to remove.

Viewing AS path information

Use this procedure to display the configured information of AS path lists.

Procedure

1. Log on to ACLI in Privileged EXEC mode.
 2. At the command prompt, enter the following command:

```
show ip as-list [<as-list-id>]
```
-

Variable definitions

The following table defines parameters for the `show ip as-list` command.

Variable	Value
<as-list-id>	Specifies an integer value between 1 and 1024 that represents the AS-path list ID you want to display. If not specified, all the configured as-list entries will be displayed.

Configuring community lists

Use this procedure to configure community lists to specify permitted routes by using their BGP community.

About this task

A community list acts as a filter that matches communities or AS numbers.

Procedure

1. Log on to ACLI in Global Configuration or BGP Router Configuration mode.
2. At the command prompt, enter the following command:


```
ip community-list <list-id> memberid <member-id> <permit|deny> community-string <community-string>
```

Variable definitions

The following table defines parameters that you enter with the `ip community-list` command.

Variable	Value
<member-id>	Specifies an integer value from 0–65535 that represents the member ID in the community list.
<list-id>	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.

Variable	Value
<i><community-string></i>	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).

Deleting community lists

Use this procedure to delete community lists to discontinue specifying permitted routes using the route BGP community.

Procedure

1. Log on to ACLI in Global Configuration or BGP Router Configuration mode.
 2. At the command prompt, enter the following command:

```
no ip community-list <id> [memberid <member-id> | community-string <community-string>]
```
-

Variable definitions

The following table defines parameters that you enter with the `no ip community-list` command.

Variable	Value
<i><member-id></i>	Specifies an integer value from 0–65535 that represents the member ID in the community list.
<i><id></i>	Specifies an integer value from 1–1024 that represents the community list ID.
<i><community-string></i>	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)

Variable	Value
	Well known communities include: internet, no-export, no-advertise, local-as (known as NO_EXPORT_SUBCONFED).

Viewing community lists

Use this procedure to display the configured information of community lists.

Procedure

1. Log on to the Privileged EXEC mode in ACLI
 2. At the command prompt, enter the following command:

```
show ip community-list <id>
```
-

Variable definitions

The following table defines parameters for the `show ip community` command.

Variable	Value
<id>	Specifies an integer value from 1–1024 that represents the community list ID you want to display. If not specified, all the configured community list entries will be displayed.

Restarting BGP

Use this procedure to restart BGP to force the restart of global BGP, BGP for a specific peer, or BGP for a specific peer group.

Procedure

1. Log on to the Privileged Exec, Global Configuration, or BGP Router Configuration mode in ACLI.
2. At the command prompt, enter the following command:

```
ip bgp restart-bgp [neighbor <nbr_ipaddr | peer-group-name>
[soft-reconfiguration <in|out>]]
```

Variable definitions

The following table defines parameters for the `ip bgp restart-bgp` command.

Variable	Value
neighbor <nbr_ipaddr peer-group-name>	Restarts BGP for a specific BGP peer or peer group.
soft-reconfiguration <in out>	Allows the router to relearn routes from the specified neighbor or group of neighbors without resetting the connection when the policy changes. If you do not specify in or out, both inbound and outbound soft configurations are triggered.

Viewing CIDR routes

Use this procedure to display information about classless interdomain routing (CIDR) routes.

Procedure

1. Log on to the Privileged EXEC mode in ACLI
 2. At the command prompt, enter the following command:
`show ip bgp cidr-only [A.B.C.D/<0-32>]`
-

Viewing imported routes

Use this procedure to display information about BGP imported routes.

Procedure

1. Log on to the Privileged EXEC mode in ACLI
2. At the command prompt, enter the following command:

```
show ip bgp imported-routes [A.B.C.D/<0-32>] [longer-
prefixes]
```

Variable definitions

The following table defines optional parameters that you can enter with the `show ip bgp imported-routes` command.

Variable	Value
A.B.C.D/<0-32>	Shows paths with this prefix. The prefix is the IP address and exact mask length (must be an integer value from 0-32).
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.D/32).

Viewing BGP routes

Use this procedure to display information about BGP routes.

Procedure

1. Log on to the Privileged EXEC mode in ACLI
 2. At the command prompt, enter the following command:

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

Variable definitions

The following table defines optional parameters that you can enter with the `show ip bgp route` command.

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).

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.D/32).
community <enable disable>	Enables or disables the display of community attributes.
ip {A.B.C.D}	Specifies an IP address.

Clearing BGP counters

Use this procedure to clear all BGP counters or counters specific to a BGP peer.

Procedure

1. Log on to the Privileged EXEC mode in ACLI
 2. To clear all BGP counters, enter the following command:
`ip bgp stats-clear-counters`
 3. To clear all counters specific to a peer, enter the following command:
`ip bgp stats-clear-counters neighbor <A.B.C.D | peer-group-name>`
-

Chapter 4: BGP Configuration Using EDM

Configuring BGP globally

Use this procedure to configure general BGP parameters to define how BGP operates on the system, and to enable BGP so that BGP runs on the router.

About this task

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 route reflectors as required.
 10. Click **Apply**.
-

Generals field descriptions

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

Name	Description
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.
LocalAs	Configures the local AS number in the range of 0 – 65535. You cannot change the LocalAS if AdminStatus is enable.
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 neighbors in different ASs. The system prefers a path with a lower MED over a path with a higher MED. 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 NoMedPathsWorst. When enabled, BGP treats an update that is missing a MED attribute as the worst path. This would be checked only for updates with the same latest AS number in the AS Path sequence (the left most AS number should be the same in the paths for this comparison to occur). The default is enable.
DebugMask	Displays the specified debug information for BGP global configurations. The default value is none. Other options are

Name	Description
	<ul style="list-style-type: none"> • <i>none</i> : disables all debug messages. • <i>event</i>: enables the display of debug event messages. • <i>state</i>: enables display of debug state transition messages. • <i>update</i>: enables display of debug messages related to updates transmission and reception. • <i>error</i>: enables the display of debug error messages. • <i>trace</i>: enables the display of debug trace messages. • <i>init</i>: enables the display of debug initialization messages. • <i>all</i>: enables all debug messages. • <i>packet</i>: enables the display of debug packet messages. • <i>warning</i>: enables the display of debug warning messages. • <i>filter</i>: 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.
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 1–4.
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.
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.
TrapEnable	Enables or disables BGP traps.
ReflectionEnable	Enables or disables the reflection of routes from IBGP neighbors. The default is enable.

Name	Description
ReflectorClusterId	Configures a reflector cluster ID IP address. This variable applies only if you enable ReflectionEnable, 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 ReflectionEnable is enable. The default is enable.

Viewing BGP statistics

Use this procedure to display BGP statistics.

Procedure

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

Configuring aggregate routes

Use this procedure to configure aggregate routes that allow the router to advertise a single route (aggregate route) that represents all destinations.

Before you begin

- Enable aggregate routes globally.
- You need the appropriate aggregate address and mask.
- If required, policies exist.

About this task

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 **PrefixLen**.
 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.
PrefixLen	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.

Configuring allowed networks

Use this procedure to configure BGP allowed networks to determine the network IP addresses that BGP advertises.

The allowed addresses determine the BGP networks that originate from the Ethernet Routing Switch 5000.

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
NetworkAfAddr	Specifies the network prefix that BGP advertises.
NetworkAfPrefixLen	Specifies the network subnet mask.
NetworkAfMetric	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 peers. The range is 0–65535.

Configuring BGP peers

Use this procedure to configure BGP peers to connect two routers to each other for the purpose of exchanging routing information.

About this task

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 **true**.
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).
RemoteAs	Configures a remote AS number for the peer or peer-group in the range 0–65535.
Enable	Controls whether the peer connection is enabled or disabled. The default is disabled.
RoutePolicyIn	Specifies the policy (by name) that applies to all routes learned from this peer.
RoutePolicyOut	Specifies the policy (by name) that applies to all outgoing route updates.
UpdateSourceInterface	Specifies the source IP address to use when the switch sends BGP packets to this peer or peer group.
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 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

Name	Description
	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.
DefaultOriginate	This parameter enables or disables sending the default route information to the specified neighbor or peer. The default value is false.
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 4000 routes and the range is 0–4000. A value of 0 means no limit exists.
DebugMask	Displays the specified debug information for the BGP peer. The default value is none. <ul style="list-style-type: none"> • <i>none</i>: disables all debug messages. • <i>event</i>: enables the display of debug event messages. • <i>state</i>: enables display of debug state transition messages. • <i>update</i>: enables display of debug messages related to updates transmission and reception. • <i>error</i>: enables the display of debug error messages. • <i>trace</i>: enables the display of debug trace messages. • <i>init</i>: enables the display of debug initialization messages. • <i>all</i>: enables all debug messages. • <i>packet</i>: enables the display of debug packet messages. • <i>warning</i>: enables the display of debug warning messages. • <i>filter</i>: enables the display of debug messages related to filtering.
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.
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

Name	Description
	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.

Configuring peer groups

Use this procedure to 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.

Name	Description
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.
KeepAlive	Specifies the time interval, in seconds, between sent BGP keep alive messages to remote peers. The range is 0–21845 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 0 or 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 400 routes and the range is 0–4000.
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.
SendCommunity	Enables or disables sending the community attribute of the update message to the specified peer group. The default value is disable.

Viewing BGP summary route information

Use the following procedure to display current BGP route information.

Procedure

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

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 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 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
RouteSource	Specifies the source protocol for the route redistribution entry.
Enable	Enables (or disables) a BGP redistribute entry for a specified source type.
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

Use this procedure to configure a prefix list to allow or deny specific route updates.

About this task

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 route policies

For information about configuring route policies, see Avaya Ethernet Routing Switch 5000 Series Configuration — IP Routing NN47200–503.

Configuring an AS path list

Use this procedure to 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 1–1024.
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

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 1–1024.
MemberId	Specifies the community 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).
Community	Specifies the community access list community string.

