



Vanguard Applications Ware
IP and LAN Feature Protocols

Border Gateway Protocol (BGP-4)

Notice

©2008 Vanguard Networks
25 Forbes Blvd
Foxboro, MA 02035
Phone: (508) 964 6200
Fax: (508) 543 0237
All rights reserved
Printed in U.S.A.

Restricted Rights Notification for U.S. Government Users

The software (including firmware) addressed in this manual is provided to the U.S. Government under agreement which grants the government the minimum “restricted rights” in the software, as defined in the Federal Acquisition Regulation (FAR) or the Defense Federal Acquisition Regulation Supplement (DFARS), whichever is applicable.

If the software is procured for use by the Department of Defense, the following legend applies:

Restricted Rights Legend

Use, duplication, or disclosure by the Government
is subject to restrictions as set forth in
subparagraph (c)(1)(ii) of the
Rights in Technical Data and Computer Software
clause at DFARS 252.227-7013.

If the software is procured for use by any U.S. Government entity other than the Department of Defense, the following notice applies:

Notice

Notwithstanding any other lease or license agreement that may pertain to, or accompany the delivery of, this computer software, the rights of the Government regarding its use, reproduction, and disclosure are as set forth in FAR 52.227-19(C).

Unpublished - rights reserved under the copyright laws of the United States.

Notice (continued)

Proprietary Material

Information and software in this document are proprietary to Vanguard Networks (or its Suppliers) and without the express prior permission of an officer of Vanguard Networks, may not be copied, reproduced, disclosed to others, published, or used, in whole or in part, for any purpose other than that for which it is being made available. Use of software described in this document is subject to the terms and conditions of the Vanguard Networks Software License Agreement.

This document is for information purposes only and is subject to change without notice.

Part No. T0100-13, Rev I
Publication Code: TK
First Printing: January 2002

Manual is current for Release 7.3 of Vanguard Applications Ware.

To comment on this manual, please send e-mail to vntechsupport@vanguardnetworks.com

Chapter 1.

Introduction

BGP Functions	1-5
BGP Protocol Entities	1-7
BGP Protocol Operation	1-9
BGP Process	1-11
BGP AS Topology Support	1-14
AS Support	1-15

Chapter 2.

Applications

Dividing a larger AS into Multiple Smaller AS's	2-2
Connecting various sites of an Enterprise	2-3
Inter ISP Communication	2-4
Internetworking in a Branch Network	2-5
Multihoming with BGP	2-6
Redistribution of BGP-4 Routes into RIPv2	2-7
BGP to RIPv2 Redistribution Features	2-9
BGP to RIPv2 Redistribution Functions	2-11
BGP to RIPv2 Route Redistribution Policy.....	2-11
BGP to RIPv2 Route Importing.....	2-12
RIPv2 Route Aggregation.....	2-13
RIPv2 Advertisement.....	2-14
BGP to RIPv2 Redistribution Typical Application	2-15
RIPv2 to OSPF and BGP Interaction Example	2-16
Impacts of RIPv2 Route Redistribution on Other Features	2-18
BGP Aggregation	2-19
BGP Aggregation Features	2-21
BGP Aggregation Functions	2-22
Aggregation Rules and Routes	2-22
BGP Aggregation Examples	2-24
BGP Aggregation Scaling.....	2-24
BGP Aggregation with Multi-Homed AS	2-24
BGP Aggregation with Specific Subnet	2-25
Impacts of Aggregation on Other Features	2-27
BGP Community Attribute	2-28
Application Examples	2-30
Node Boot and Table Boot	2-32
SNMP Network Management	2-33

Chapter 3.

Configuration

Accessing the BGP Parameter Records	3-2
Configuring BGP Routing Parameters	3-4
Configuring OSPF Routing Parameters	3-33

Contents (continued)

Configuring RIPv2 Route Redistribution Parameters	3-39
Configuring BGP Aggregation Parameters	3-48
Configuring BGP Community Attribute	3-57

Chapter 4.

Statistics and Diagnostics

BGP Statistics	4-2
Using the BGP Peer Statistics	4-3
Using the BGP Routing Table Statistics	4-6
Using BGP AS Path Database Display	4-11
Using BGP Aggregation Statistics	4-15
BGP to RIPv2 Redistribution Statistics	4-16
BGP Diagnostics	4-22

Glossary

Overview

The Border Gateway Protocol version-4 (BGP-4) is an Inter-Autonomous System routing protocol which provides loop free inter domain routing between autonomous systems (AS's). AS's are used to provide loop free routing between multiple AS's comprising the internet. BGP4 is an extension of BGP-3 which provides the support for routing information reduction and aggregation based on CIDR architecture.

Vanguard Release 6.2 and prior releases supported redistribution of BGP routing information into OSPF domain. Static routes and default gateways have to be configured in the RIP-only capable routers to route traffic beyond the AS boundary.

Vanguard Release 6.3 and greater includes:

- RIPv2 Route Redistribution
- Aggregation

Vanguard Release 6.4 and greater supports Community Attribute.

Vanguard Release 7.3 introduces the MD5 TCP password to the BGP Peer configuration. This was first available in the 7.2S100 service pack.

More detailed information regarding RIPv2 route redistribution, aggregation, and community attribute can be found in Chapter 2 of this manual.

Limitations

BGP RIPv2 route redistribution and BGP aggregation should be enabled independently even though they are running in the same router simultaneously.

- 1) The BGP RIPv2 route redistribution cannot work if either one of BGP, RIP, and IP modules is not included in the image.
- 2) The BGP aggregation cannot work if BGP module is not included in the image.
- 3) No support of automatic BGP route aggregation.
- 4) BGP RIPv2 Import policy cannot be specified for an IP subnet with subnet prefix length less than 8.
- 5) BGP Aggregation Profiles cannot be specified for an IP subnet with subnet prefix length less than 8.
- 6) BGP-IGP synchronization cannot work properly if the tag value is not set as automatic generation.
- 7) When RIPv2 authentication is used in RIP protocol, the BGP to RIPv2 route redistribution will not set the tag for the imported routes in RIP domain. In this case BGP- IGP synchronization could be impacted if BGP aggregation is enabled and only the aggregate route is redistributed into RIP domain. Proper BGP to RIPv2 route redistribution policies about the aggregate routes and their component routes should be configured to utilize the BGP and IGP synchronization.
- 8) If a BGP route has a next hop that is not a directly connected subnet, this route cannot be imported into Global Routing Table.

- 9) Only well-known community attributes and path backup mechanisms are supported.
- 10) User defined BGP community attributes, BGP extended community and BGP load sharing are not supported.

RIPv2 Domain

Vanguard Release 6.3 and higher introduces the redistribution of BGP routing information into RIP, the RIP capable routers are able to route traffic beyond the AS boundary, without the configuration of static routes and default gateway.

BGP Aggregation

BGP aggregation is used to summarize a list of contiguous IP destinations in the BGP routing table into an aggregate entry. BGP aggregation can reduce the size and slow the growth of the BGP routing table. Vanguard processes aggregated attributes in a received BGP update. The functional specification of BGP aggregation is to support the route aggregation and the sending of the aggregated routes. Vanguard fully supports the BGP aggregation in both directions.

BGP Community Attribute

BGP Community Attribute is a transitive optional attribute, which is mainly used to specify the preference of a path and restrict the advertisement of routes according to administrator's policies.

Objective

This chapter provides the introduction to Border Gateway Protocol version-4 (BGP4), its applications and the specific implementation requirements posed by it. The requirements mentioned in this document assume an incremental phase-wise implementation of BGP on Vanguard Networks routers. Each phase is defined based on the AS topology in which BGP is expected to be used and the level of functionality support and flexibility provided to the user. The requirements in each of the phases are aligned with the main functional area of BGP.

Characteristics

Below is a brief description of some of the important characteristics of BGP4:

- **Path Vector Protocol:** BGP communicates the entire routing path with its neighbors. This includes the list of all AS's a route has traversed. The full path information provides the BGP speaker sufficient information to make a graph of AS connectivity. This loopfree map of AS can be used to suppress the routing loops and eliminate the count-to-infinity problem found in normal distance vector protocols.
- **CIDR enabled:** BGP4 provides the support for necessary mechanisms to achieve classless routing which is one of the major enhancements from BGP3. These mechanisms include support for advertisement of supernatted routes (using IP prefix in advertisement messages) and support for route aggregation and information reduction. This significantly reduces the amount of routing information a BGP speaker has to maintain and exchange with its peers.
- **Authentication Support:** BGP has its own authentication mechanism which can be in addition to the authentication mechanisms provided by the transport protocols (like TCP). It is used for the authentication of protocol messages exchanged between the two BGP speakers. A BGP speaker can optionally negotiate an authentication mechanism to be used at the time of opening the BGP session with its peer. Once it is done further exchange of information (route update or error notification) can be authenticated using the

authentication information negotiated during the session opening time.

- **Hop by Hop Routing Paradigm:** BGP advertises only those routes to its neighbors (peers) which it itself uses. This reflects the "hop-by-hop" routing paradigm used today by internet. Any routing policy conforming to "hop-by-hop" paradigm can be supported by BGP.
- **Incremental Updates:** BGP does not require the periodic refresh of entire routing table. Once BGP speakers has performed their initial exchange of routing information, only incremental updates are sent triggered by local routing table changes which might have been occurred because of change in network topology. This is achieved by maintaining the current version of the entire BGP routing tables of all its peer as long as the BGP connection is alive with its peers.
- **Reliable Transport:** BGP uses TCP as its transport protocol. This eliminates the need for explicit update retransmission, fragmentation, acknowledgment and sequencing to be done by BGP as these functions are inherently supported by any reliable transport protocol including TCP. BGP uses TCP port 179 for establishing its connections with peers. The error notification mechanism used in the BGP assumes that the transport protocol supports the graceful close of the established connection between the BGP peers.
- **Quick Backup:** When there are multiple feasible paths available for the set of destinations associated with the same address prefix, BGP maintains all of them but advertise the one which is selected by its path selection process. This maintenance of multiple feasible paths speeds the adoption of the alternate path in case of primary path failure.
- **Complex AS Topology Support:** BGP supports the complex AS topologies in contrast to centralized AS topology supported by EGP (Exterior Gateway Protocol). It does not pose any topological restrictions on the interconnection of AS's. BGP is applicable in all types of AS topologies namely stub, multihome and transit. This is mainly possible because of the fact that BGP maintains the full graph of AS connectivity and maintains the consistent view of routing within an AS. BGP presents the consistent view of AS to the external world by maintaining the consistent routing information between all the BGP speaker belonging to the AS.
- **Complex Policy Support:** BGP provides the support for complex routing policy to be enforced at AS level. This policy enforcement could be based on the various routing preferences and constraints. Policies are provided to the BGP in the form of local configuration information. BGP allows the enforcement of various complex policy by supporting the following mechanisms.
 - **Complex Path Selection:** BGP supports complex path selection procedures when presented with multiple feasible paths for the set of destination represented by a given address prefix. To decide the preference of the routes, path selection process takes input as information present in the full AS Path as well as the routing constraints provided in local configuration information. Path with highest preference is chosen as the best path.
 - **Controlled exchange of routing information:** BGP provides the mechanisms to control the exchange of routing information (both BGP and Interior Gateway Protocol IGP learned) with its adjacent neighbor AS's. This includes controlled advertisement of routes to BGP neighbors and the

controlled import of routes from its neighbors. This control can be of various granularity based on the information provided in the local policy configuration.

BGP Functions

BGP performs three main functions:

- Exchange of Network Reachability Information
- Policy Enforcement
- Interaction with IGP

Exchange of Network Reachability Information

Exchanging network reachability information with other AS's is the primary function of BGP. The information exchanged serves the two main purpose.

- It contains the sufficient information from which a graph of AS connectivity can be formed which can be used by a BGP speaker to suppress routing loops so that the loopfree routing can be achieved and can be used to enforce routing decision based on the local AS policy.
- A BGP speaker can exchange network reachability information with its multiple peers/neighbors at the same time. In this case BGP speaker maintains reachability information received from all its neighbors and run its decision process to determine the best route for a particular set of destinations. It is this route the BGP speaker uses itself for routing and advertises to its neighbors. This best route selection is done based on the local policy configuration.

Policy Enforcement

This is one of the major functionalities expected from BGP. Policies are set of rules which represents As's routing preferences and constraint put on external traffic. Policies are provide to BGP in form of local configuration and hence BGP can provide support for various complex routing policies. BGP can provide the following types of policies.

- **Path Selection:** This policy provides an AS the capability of prefer a particular path to a destination, in case of multiple paths are available for the same destination. These multiple path can be available from same AS (different BGP speaker) or different AS's. The path selection criteria can be based on various factors like number of AS's is AS_PATH, path origin, presence or absence of a particular AS, local preference. The exact path selection criteria used for path selection process is local policy matter.
- **Advertisement Control:** This policy provides an AS the capability of controlling advertisement of BGP learned route to adjacent AS's. This control can be done with varying dedrees of granularity like Address prefix, originating AS, and neighbor AS. This granularity can be defined by local configuration.
- **Reception Control:** This policy provides an AS the capability of controlling reception of advertisement from other BGP speakers in adjacent AS's. This control can be done with varying degrees of granularity like Address prefix, originating AS, neighbor AS. This granularity can be defined by local configuration.

Interaction with IGP

An AS is providing the transit service to other AS's. It requires a certain degree of interaction and coordination between the BGP and the IGP used by the local AS. This interaction solves the following purpose.

- Transit traffic routing: Transit traffic passes both BGP and non-BGP routes inside the AS. This requires that the exterior routing information be redistributed into IGP so that transit non-BGP routers have information about the external destination and they route the traffic towards the proper exit gateway.
 - Propagation of BGP routing information across the AS: IGP used inside the AS can be used as one of the mechanisms to carry BGP route information across the local AS. In this case BGP routes are carried in IGP packets across the AS. Certain tagged IGP protocols like OSPF/RIP provide a way of propagating the information from an external AS (from which routes have been received). This information can be used for either IGP's local use or for the synchronization of IGP with BGP when IBGP is used between the internal peers of AS to carry BGP route information across AS and when IGP is used for transit traffic routing.
-

BGP Protocol Entities

This section describes the various BGP protocol entities involved in BGP protocol operation.

BGP Speaker

A BGP speaker is a system running BGP. It does not have to be a router.

BGP Neighbors/ Peers

A pair of BGP speakers exchanging the inter-AS routing information. Based on the location of BGP speaker BGP peers can be of two types.

- **Internal Peers:** A pair of BGP speakers in the same AS. Internal BGP speakers need not be directly connected that is they need not be present on the same physical subnetwork. If the AS is providing the transit service to other AS's all internal peers establishes a full-mesh BGP sessions between them for the purpose of exchanging the inter-AS routing information between them. This allows the inter-AS routing information to be carried across the AS. All the internal peers of an AS presents the consistent view of external routing information to the outside world by using consistent set of policies between them. If the two internal peers are not directly connected, they should not advertise the external routing information to other AS's till the internal non-BGP routers has been updated about the external routing information that is till the internal peer has the IGP routes to the set of external destinations to be advertised to external AS's.
- **External Peers:** A pair of BGP speakers in different AS. External BGP speakers need to be directly connected that is they should be present on the same physical subnetwork. External peers establishes a BGP connections between them for the purpose of exchanging the local (originated within the AS) and external routing information (learned from other AS's) between them. External routing information is exchanged only when the local AS is providing the transit service to other AS's. A BGP speaker advertises only those routes to its external neighbors that it uses. For example routes which it has selected for its local use resulting from a combination of local policy and normal path selection process. An AS can have multiple BGP external peers, which are multiple BGP connection to other AS's.

BGP Session

A TCP session between the two BGP neighbors established for the purpose of exchanging the routing information. It is this connection on which BGP peers exchange various protocol messages.

BGP Peer Authentication

BGP Peer Authentication is the newest method to reduce security risks in a BGP network. The Vanguard implementation of BGP peer authentication uses the TCP MD-5 signature as specified in RFC 2385. This algorithm takes a key, the password entered during configuration, and performs an MD-5 hash on the key, and sends the resulting hash to the remote peer. The password itself is never sent over the connection. Both sides of an authenticated BGP peer session must use the same password.

The authentication occurs in the TCP session not on the BGP peer session. It provides added confidence that packets received from the TCP peer actually originated from the authorized TCP peer.

AS Connection

Two AS are said to be connected when both of the following connections exists between them.

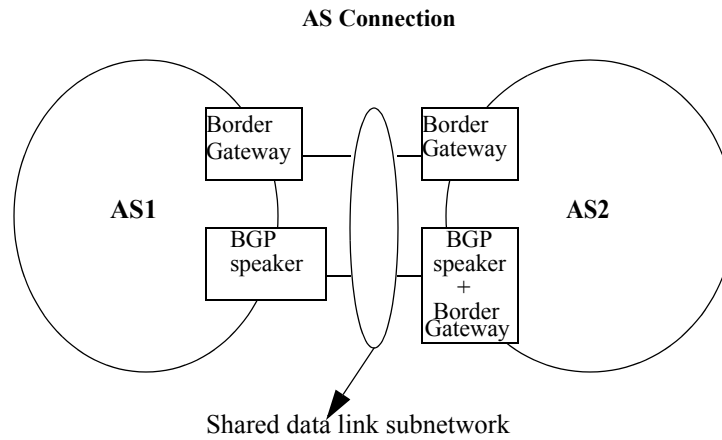


Figure 1-1. AS Connection

Physical Connection

The physical connection is the shared data link subnetwork between the two AS's. On this data link each AS has at least one border gateway belonging to that AS so that traffic between the border gateway can be forwarded without any intra-AS or inter-AS routing help. There can be multiple border gateway at each AS connection.

BGP Connection

The BGP connection is the BGP session between BGP speakers in each of the AS's. BGP speakers that form the BGP connection should share the same data link subnetwork as that of their border gateways so that protocol traffic can be exchanged between BGP speakers without any intra-AS or inter-AS routing help. BGP speaker makes use of BGP connection to exchange the routes to particular set of destination reachable via border gateways present on the same physical network. The BGP speaker which establishes the BGP connection need not be a border router and vice versa. There can be multiple BGP speakers at each AS connections.

BGP Protocol Operation

Following is the brief description of the BGP protocol operation which is used to exchange the routing information between the two AS's. Which are two external BGP speakers. This process of routing information exchange comprises of the following main stages.

Opening and confirming a BGP connection with the neighbor

BGP communication between the two neighbors commences with TCP transport connection being established between them. Once the transport connection is established, each BGP speaker sends an OPEN message to each other to establish the BGP connection between them. OPEN message defines the originating BGP speaker's AS number, its BGP router identifier and hold time for the connection. If no other BGP messages are received for a period of hold time, the originating speaker assumes an error, sends the notification message and closes the connection. BGP speaker optionally can negotiate with its peer the authentication mechanism to be used to authenticate the protocol message exchange with its peer. This is specified as part of connection OPEN message sent to the peer which includes the authentication mechanism code and the authentication data to be used to validate the protocol messages exchanged between the peers. An acceptable OPEN message is acknowledged or confirmed by another OPEN message. The reception of the second OPEN message in response to BGP OPEN message indicates the successful connection open which can now be used by BGP peers to further exchange other routing protocol messages.

Maintaining the BGP Connection

BGP speakers periodically exchanges the KEEPALIVE messages between them to maintain the BGP connection. If no BGP messages are received for a period of hold time, which was negotiated during the connection open time, the originating speaker assumes an error, sends the notification message and closes the connection.

Sending Reachability Information

BGP speakers send UPDATE messages to its peers for exchanging reachability information. BGP speakers uses the UPDATE messages to either advertises a single feasible route or to withdraw multiple unfeasible routes from service. A BGP speaker may simultaneously advertise a single feasible route and withdraw multiple infeasible routes by sending a single UPDATE message to its peer. The UPDATE message sent by BGP speaker to its peer contain the following information:

- **Network Layer Reachability Information (NLRI):** It indicates the set of feasible routes in form of IP prefix routes. There can be variable number of IP prefixes advertised by BGP speaker to its neighbor in a single UPDATE message
- **PATH Attributes:** This represents the additional information about the path of the advertised feasible routes. These path attributes are used by a BGP speaker to indicate the information like Origin of the route, list of AS numbers that the routes have traversed, IP addresses of the next hop border router for the listed feasible routes. A BGP speaker can optionally advertise the additional information about the route like AS number and IP address of the last BGP speaker which performed the aggregation before advertising the routes, or whether a less specific route was chosen without selecting the more specific route by a particular BGP speaker in the path.

- **Withdrawn Routes:** This indicates the list of routes which are withdrawn from the service or they are no longer valid. A BGP speaker removes all the routes indicated by this withdrawn route list, from its routing table. BGP speaker can withdraw multiple IP prefixes using a single UPDATE message.

Notifying Errors

BGP speakers send the notification message to its peer when it detects any error condition. BGP speaker also closes the transport connection after sending the notification message to its peer. Following types of error messages are notified by the BGP speaker.

- **BGP Message Error:** This indicates error in BGP protocol message like BGP message header, OPEN message error, UPDATE message error.
 - **Hold Timer Expired:** This indicates the expiration of BGP connection hold time.
 - **Finite State Machine Error:** This indicates some problem in BGP FSM.
 - **Cease:** This indicates the explicit close of BGP connection by other peer without any presence of error condition.
-

BGP Process

Following is the step-by-step description of BGP process used by BGP speaker to receive, update and advertise the routing information to its peers. At each step BGP process association with the local AS policies is described.

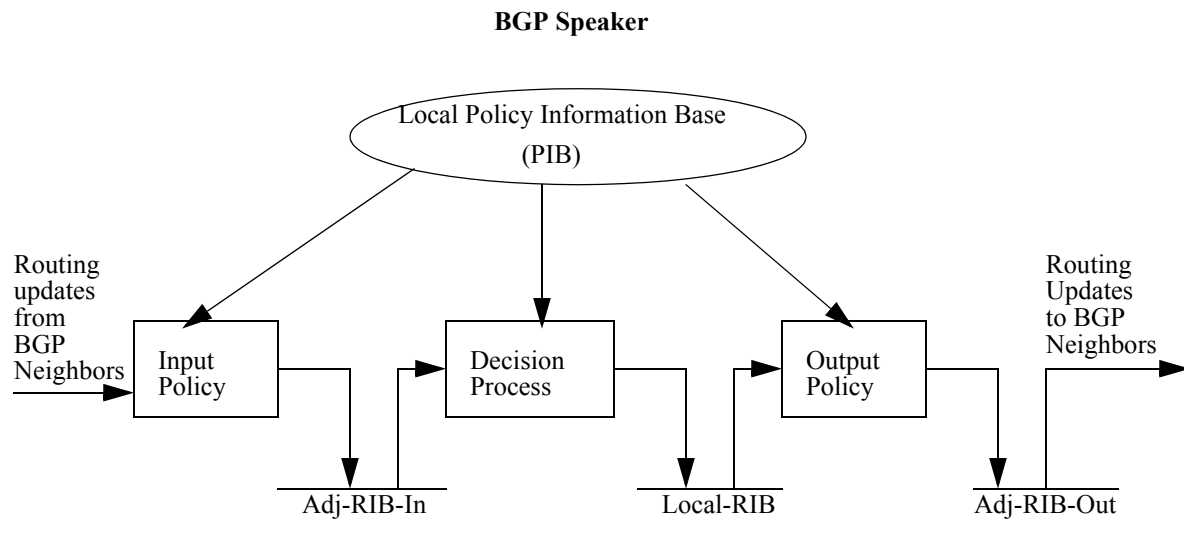


Figure 1-2. BGP Speaker

- 1) **Receiving of Update messages from Neighbor:** Update messages are received from the neighbor and BGP speaker takes the following actions for the routes present in the UPDATE message.
 - If the UPDATE message contains the non-empty WITHDRAWN ROUTES field, then routes specified in this field should be removed from the Adj-RIB-In and the decision process is run to indicate that the previously advertised are no longer available for use.
 - If the UPDATE message contains some feasible routes then those routes processed by the local Input Policy Engine as mentioned in Step2.
- 2) **Input Policy Processing:** This includes filtering the route information received in the update message which is further considered for processing by local BGP speaker. This filtering could be based on IP prefixes present in NLRI or AS numbers present in AS path information of the route. For e.g. BGP speaker may decide not to receive route containing a particular AS in its AS path segment or routes from a particular neighbor AS or routes originated by a particular AS. This filtering information is provided to the local BGP speaker in form of local Input policy configuration. This configuration information represents the local AS's policy for the routes which the local AS is willing to process. The routes

which are allowed (not filtered) are put in the Adj-RIB-In, on the other hand routes which are filtered, are not further processed and dropped by BGP speaker. In case of newly received routes or routes specifying the new path information for the previously received routes, BGP speaker runs the BGP Decision Process as mentioned in Step3.

- 3) **BGP Decision Process:** BGP decision process operates on the routes present in the Adj-RIB-In and selects the routes to be used by local BGP speaker and to be further advertised to its internal and external peers. This route selection process is formalized by calculating degree of preference for each of the newly received or replacement feasible routes. This degree of preference is calculated using a function which takes attributes (like ORIGIN, NLRI, AS_PATH, LOCAL_PREF) of the given route and local path selection policy as an argument and assign a number denoting the degree of preference for that route. Once degree of preference is assigned to each of the candidate routes, routes with highest preference is chosen as the best route for each distinct destination and is put in the local-RIB s. These are routes which are used by the local system to forward the packet and further advertised to BGP speaker's internal peers and considered for further processing to advertise them to BGP speaker's external neighbors. Unfeasible routes are removed from the Local-RIB and also the corresponding unfeasible routes are removed from Adj-RIB-In.
- 4) **Output Policy Processing:** All feasible routes selected by selection process are put in local-RIB and are input to output policy processing. This processing includes optional filtering of routes to be advertised to external neighbors and optional aggregation of feasible routes present in the local-RIB before advertising them to neighbors.
 - **Outbound Filtering:** This means filtering the routes to be advertised to the external neighbors based on IP prefixes present in NLRI or AS numbers present in AS path information of the route or the neighbor AS to which the routes are being advertised. For example BGP speaker may decide not to advertise route containing a particular AS in its AS path segment or routes originated by a particular AS or routes to a particular neighbor AS.
 - **Route Aggregation:** This is the process of combining the characteristics of several different routes in such a way that a single route can be advertised. This included aggregation of both NLRI and other path attribute (for e.g. AS_PATH) of the components routes so that single route with particular NLRI and other path attribute information represents all the components. BGP speaker does this to reduce the amount of route information it has to store and exchange with other peers in the neighbor AS's.

This output policy information is provided to the local BGP speaker in form of local output policy configuration. This configuration information represents the local AS's policy for the routes which the local AS is willing to advertise to its neighbors and optional aggregation of routes to be advertised to its neighbors. Routes resulting after Outbound Filtering and Route Aggregation processing are put in Adj-RIB-Out to get advertised to external neighbors.

- 5) **Sending of Update message to Neighbors:** The Routes put in Local-RIB-Out resulted from above mentioned output policy processing are put in an UPDATE message and advertised to each of the neighbors. The sending of UPDATE messages to BGP neighbors can be optionally done in a controlled fashion. This may require to reduce the link bandwidth needed to advertise the UPDATE messages and the processing power to process the information contained in these update messages. For example, you may be required to control the rate at which a BGP speaker advertises or originates the routes to particular set of destinations to its neighbors.
-

BGP AS Topology Support

BGP provides the supports for complex AS topologies. This is possible because BGP can operate in various types of AS differing in terms of their connectivity to the external world and the type of traffic carried by them.

AS Traffic

BGP categorizes the AS traffic in following two categories.

- **Local Traffic:** traffic that either originates in or terminates in the local AS, that is. the source or the destination IP address of the packet is that of a system inside the AS.
 - **Transit Traffic:** traffic that neither originates in or terminates in the local AS, that is the source and the destination IP address of the packet is that of a system outside the AS.
-

AS Support

Based on the AS connectivity to other AS's and type of traffic carried by them, BGP provides the support for following types of AS's.

Stub AS:

An AS that has only a single inter-AS connection to one other AS. A stub AS only carries local traffic.

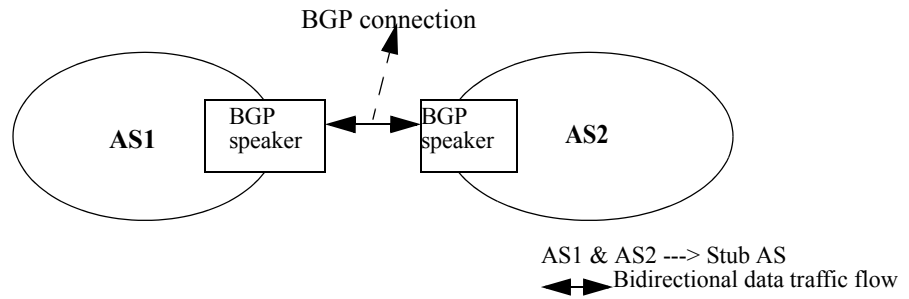


Figure 1-3. STUB AS Topology

Multihome AS

An AS that has connections to more than one other AS, but refuses to carry transit traffic.

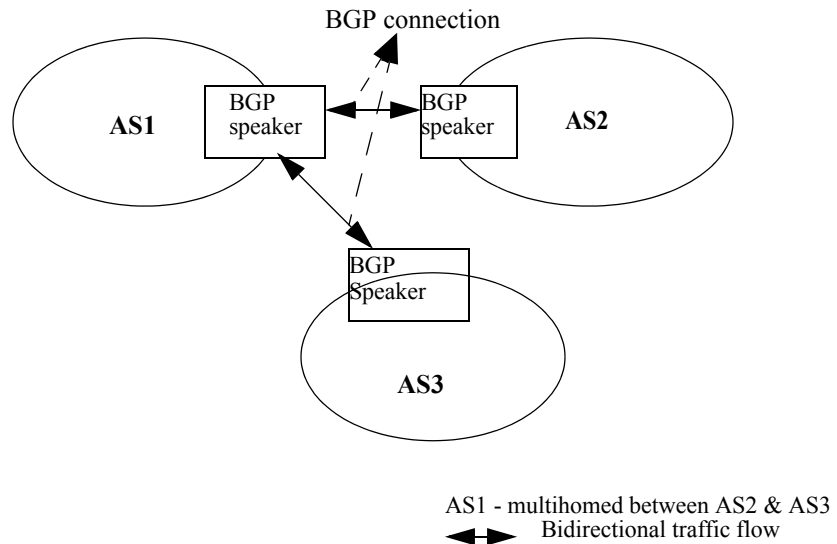


Figure 1-4. Multihome AS Topology

Transit AS

AS that has connections to more than one other AS and carries both local and transit traffic. Such an AS may impose the policy restrictions on what transit traffic will be carried through it.

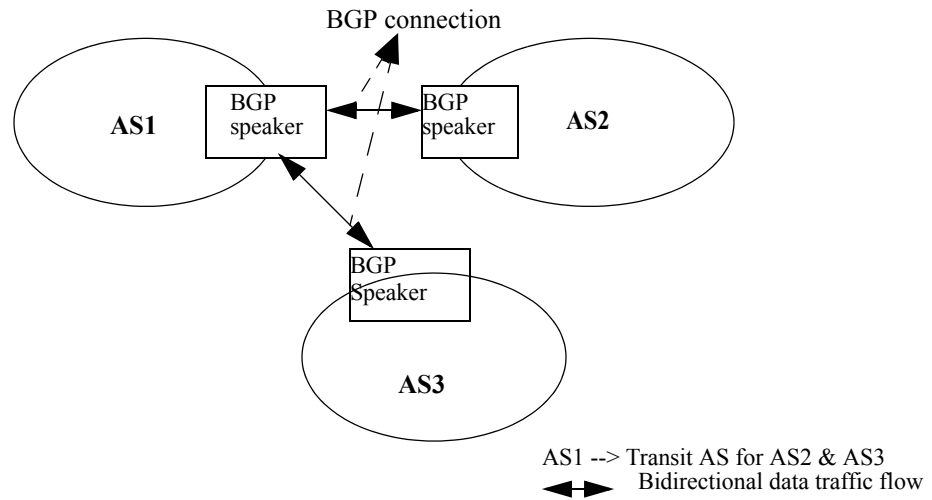


Figure 1-5. Transit AS Topology

Overview

This chapter describes the scenarios and applications in which BGP can be used to exchange routing information between different autonomous systems (AS's) and to enforce the AS level policy governing the inter-AS routing of internet traffic. This chapter also describes RIPv2 Route Redistribution, Aggregation and Community Attribute.

Dividing a larger AS into Multiple Smaller AS's

Dividing a larger AS into multiple smaller AS's

A very large network can have scaling problems in terms of routing traffic update processing and routing table overflow. As the network size increases the management of the network becomes more complex. Enforcing network wide policies like controlling the traffic through the network or enforcing a particular routing path for certain traffic etc. becomes more difficult.

Solution

Scaling problems in terms of routing traffic update processing and routing table overflow can be solved by dividing the very large network into multiple smaller networks. This solves almost all of the scaling problem, as mentioned above, of large networks. Each of these smaller networks can be configured as the AS's and can run their own IGP for example using OSPF for intra-AS routing. BGP can be used between these smaller AS's to provide inter-AS routing of internet traffic. Transportation of BGP routing information can be done either by using IBGP or IGP running in that AS. This allows the routing broadcast to be contained within the smaller AS's or networks. Route Aggregation capability of BGP can be used to advertise the aggregates between the AS's to achieve inter-AS routing. Also BGP policy capabilities can be used to enforce the AS level policies to achieve the organization's administrative goals.

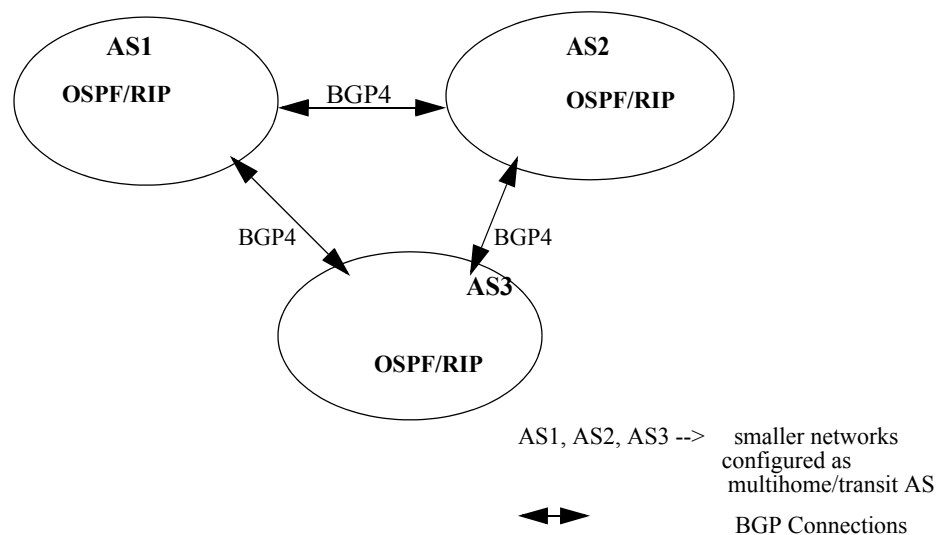


Figure 2-1. Dividing a large AS into multiple, smaller AS's

Connecting various sites of an Enterprise

Connecting Various Sites of an Enterprise

The service provider may want to provide the service of connecting various sites of an enterprise on a common IP backbone. In absence of BGP this can be done by running RIP/OSPF in the different enterprise sites and using EGP/OSPF in the ISP backbone to provide the inter-site routing. This solution does not scale well in terms of amount of routing traffic exchanged between the sites and the backbone network and also the types of policies that can be enforced in the enterprise or ISP backbone network.

Solution

Connecting various sites of an enterprise on a common IP backbone can be solved using BGP for inter-site communication. Each of the enterprise sites can be configured as stub AS's and can use BGP to advertise their routes into ISP backbone. Within each site traffic destined to external networks can be defaulted to border gateway connecting to the ISP backbone. ISP backbone can be configured as the transit AS and can run BGP to provide transit service for the traffic between the enterprise sites. ISP's transit AS can run EBGP connections to each of the enterprise sites and can run IBGP/IGP to carry external routing information across the transit AS.

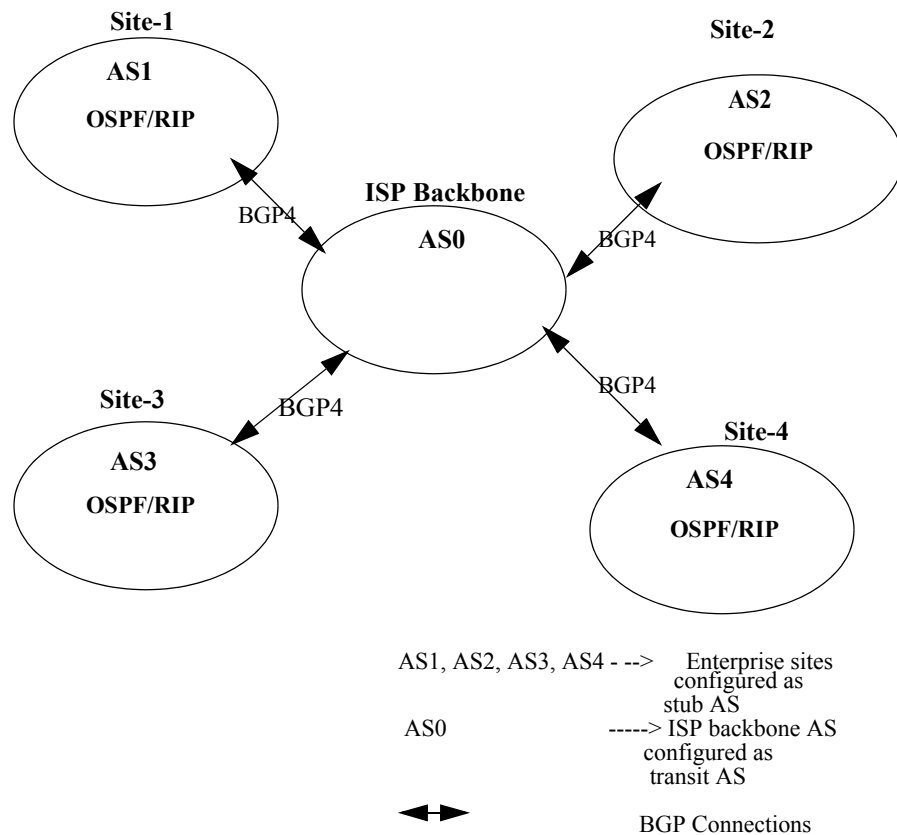


Figure 2-2. Connecting Various Enterprise Sites over a ISP Backbone

Inter ISP Communication

Inter ISP Communication

ISP networks need to be connected to the other ISP networks (probably the next ISP in the hierarchy) to perform routing of internet traffic. Today this is achieved by using either OSPF/RIP/EGP between the ISP AS's. This solution does not scale because of the inherent limitation of OSPF/RIP/EGP in terms of routing traffic overhead and the topological restrictions put on AS connectivity.

Solution

This is the typical application of BGP, BGP provides the complex AS topology support and allows complex AS level policies to be enforced between the ISP's networks. ISP AS's can be interconnected in any topological model and can use BGP to exchange external routing information between them. Also the ISPs can enforce complex policies for the routing information exchange between other ISP or for the complex path selection policies for a particular set of destinations via particular ISPs.

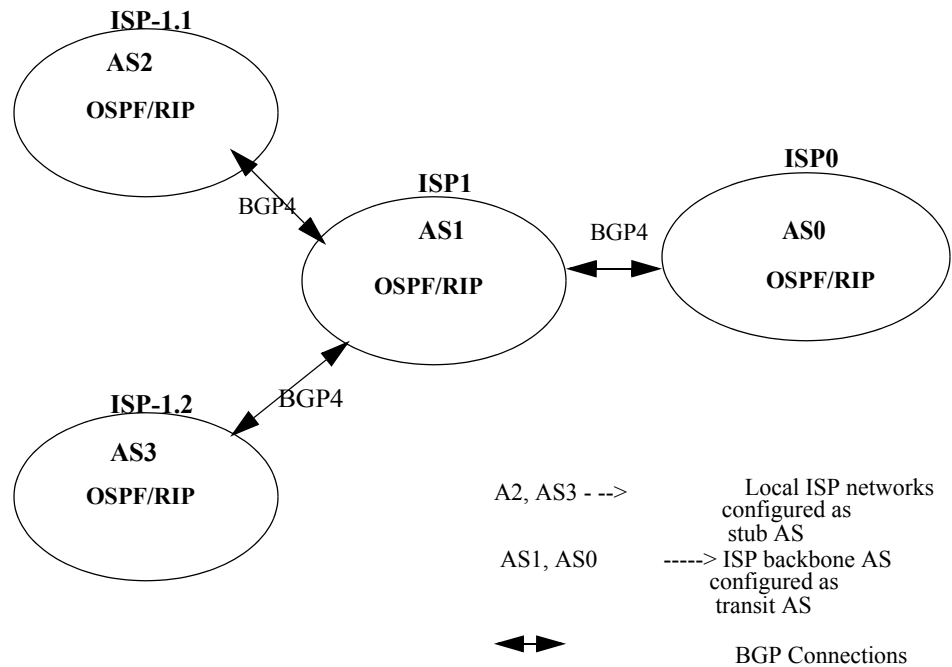


Figure 2-3. Inter ISP Communication

Internetworking in a Branch Network

Branch Network Scenario

In a branch network scenario, each of the branches of an enterprise connects to a central site router which is finally connected to ISP network. Today each of these branches run OSPF/RIP for inter branch communication and central site routers runs OSPF/EGP with ISP border router for routing of traffic in the internet. This model has some scaling problem as number of routes in the branch network, number of branch sites, local policy enforcement for inter branch or branch-Internet communication.

Solution

BGP can be used in branch network scenario to provide both inter branch and branch-Internet routing. Each of the branches can be configured as stub AS and can run OSPF/RIP for intra-site communication and BGP with central site router for inter-branch routing. Central site router can run BGP with ISP router for branch-Internet routing. This require the central site router to connect to multiple branches (AS's).

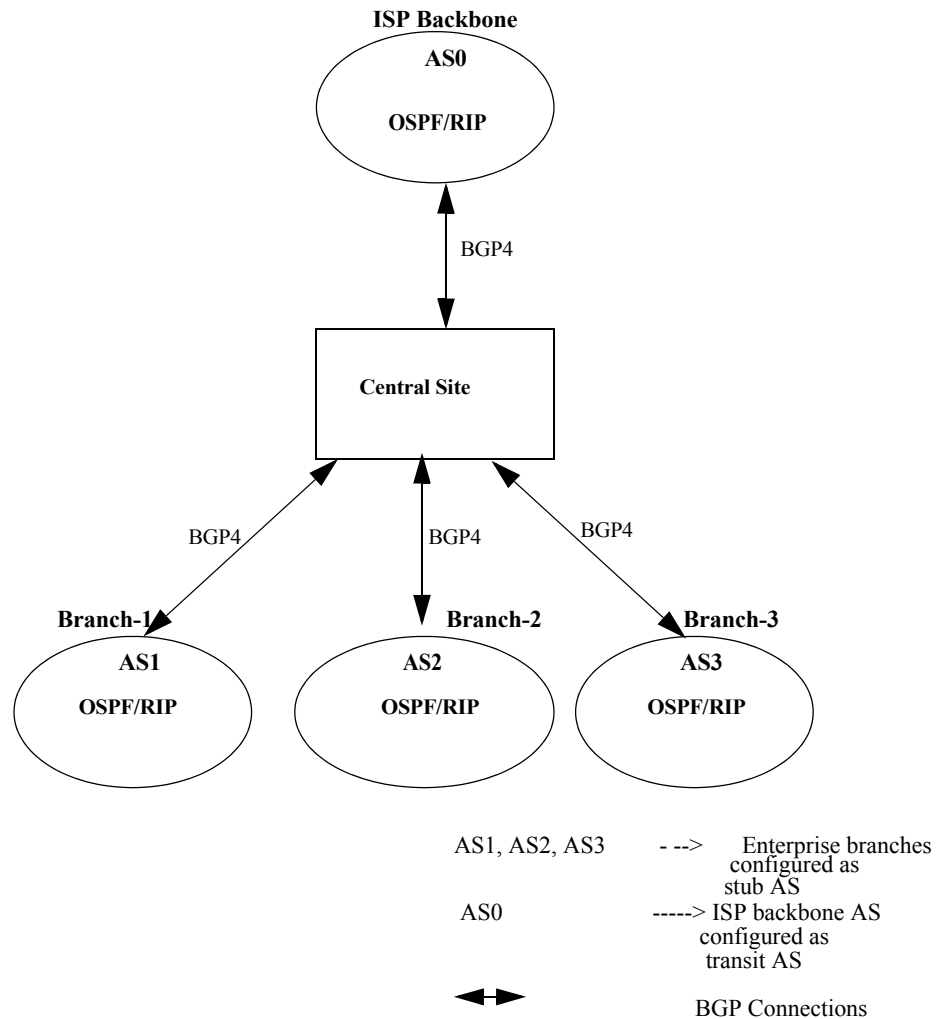


Figure 2-4. Internetworking in Branch Networks

Multihoming with BGP

Multihoming

An enterprise network may want to connect to multiple ISP's for enhancing network connectivity in terms of doing load balancing on multiple paths or creating the multiple backup paths for particular set of external networks. Today this is achieved by either using IGP between enterprise and ISP AS or using static routes for external networks inside the enterprise AS. In case of IGP's, this solution does not scale well in term of convergence time and the inherent limitation of IGP's for maintenance of multiple routing paths and also the limited flexibility in the path selection process. In case of static routes, this solution also this solution does not scale as it poses extra configuration overhead and provide limited flexibility in path selection process.

Solution

BGP can be well suited for above mentioned scenario. Enterprise network can be configured as stub AS and can connect to multiple ISPs. At each of the ISP's connection BGP can be used to exchange routing information with ISP. BGP policies capabilities can be used to enforce complex policies for the controlled exchange of routing information between enterprise and ISP's networks and for enforcing complex path selection policies for set of external destinations.

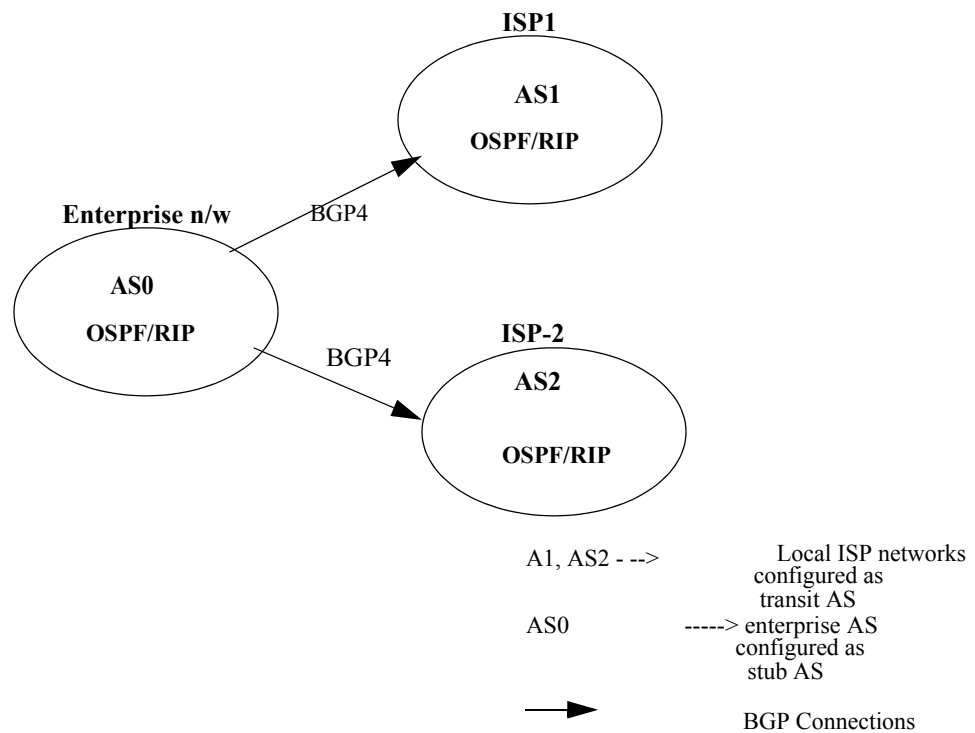


Figure 2-5. Multihoming with BGP

Redistribution of BGP-4 Routes into RIPv2

Introduction

This section describes the redistributing of BGP routes into RIPv2 domain. This feature is available with release 6.3 or greater software.

When both BGP and RIP run in a router, the router works as an Autonomous System Border Router (ASBR). The router uses the RIP as its IGP to interact with other routers within the AS and BGP as its EGP to interact with other BGP speakers in other ASs. RIP is responsible to collect the routing information for the specific AS, build and maintain IGP routing tables, and prepare routing information for BGP if this router is an ASBR. BGP is responsible to import routing information from IGP, send the routing information to other BGP speakers, receive and process other BGP speakers' updates, and build and maintain the BGP routing table.

Redistribution of BGP routes into RIPv2 is used to import the qualified route entries from the BGP routing table to RIPv2 domain so that the routers (that use RIPv2 as the only routing protocol within a specific AS) can obtain the BGP routing information that is currently only available in ASBR.

A set of policies defined for BGP to RIPv2 redistribution is similar to BGP to OSPF redistribution. If a new BGP route can satisfy any one of the policies set, the route is input into RIPv2 domain and a new RIPv2 route entry is created based on this BGP route. The path attributes in the BGP route are reflected in the protocol fields of the new RIP route. RIPv2 checks the new route against the route aggregation rules in its domain if user enables the RIPv2 route aggregation in the Vanguard router. RIPv2 advertises the new route(s) by sending an unsolicited response packet in either a periodic or triggered update. An alarm is generated to inform the users of the result. The user can get the statistics and debug information through the CTP. The network administrator is able to diagnose the functionality using the alarms and statistics.

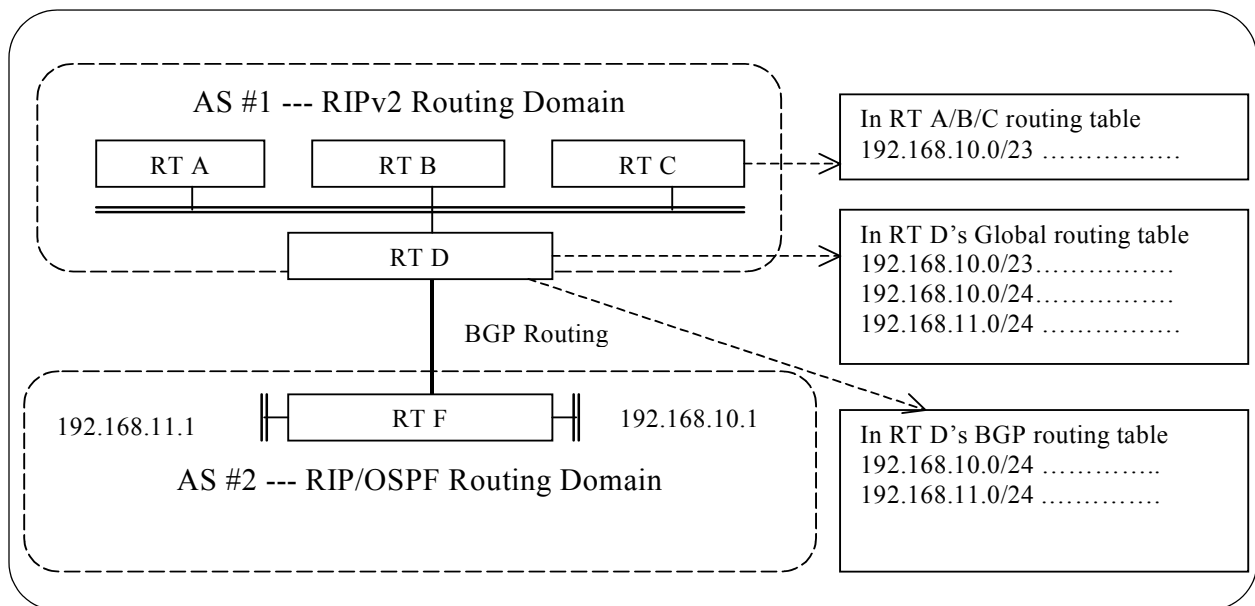


Figure 2-6. Redistributing BGP Routing Information into a RIP Domain

Redistribution Example

An example of the redistribution process is shown in Figure 2-6. AS #1 and AS #2 are two IGP domains which are using RT D and RT F as their ASBR respectively. RT F and RT D are using both BGP and IGP. RT F's IGP could be RIP or OSPF or a combination. RT D's IGP is RIPv2. RT A, RT B and RT C are all RIPv2 only routers in AS #1.

BGP route aggregation is not implemented in Vanguard routers, therefore RT F sends a BGP update packet to RT D which includes the 192.168.10.0/24 and 192.168.11.0/24. When RT D receives RT F's BGP update packet it puts 192.168.10.0/24 and 192.168.11.0/24 into its BGP routing table and lets RIPv2 take these two routes into GRT. If route aggregation is enabled in RT D, then another route 192.168.10.0/23 is formed as an aggregate route. RIPv2 advertises these route entries to other routers in the AS, either as 192.168.10.0/24 and 192.168.11.0/24 or 192.168.10.0/23 depending on the configuration in the corresponding interface.

BGP to RIPv2 Redistribution Features

Features

Following is the brief summary of various features supported by Vanguard BGP to RIPv2 Redistribution.

BGP to RIPv2 Import Policy Support: The policy provides a mechanism to control the route information which is imported from BGP into local RIPv2 routing domain. This control can differ in terms of route attributes like the Network Layer Reachability Information (NLRI) and AS path attribute.

- **Complex BGP route attributes selection:** Vanguard BGP to RIPv2 Redistribution supports complex route selection based on the NLRI attributes. The NLRI attributes can be selected in different match styles, exact or range. The NLRI attributes can be applied separately or together.
- **RIP Metric:** RIP metric with which BGP routes are imported. The metric is configured by the user according to preference.
- **RIP Tag:** IGP tag with which BGP routes are imported in RIP. Both manual and automatic generation of IGP tag is supported. Automatic RIP tag generation is used to carry the BGP route information (like next hop AS, originating AS) across the local AS via RIP.
- **Permit/Deny Policy:** This policy provides the administrator the capability of selectively importing /filtering the routes from BGP into the RIP domain.
- **Specific BGP route import:** This policy provides the network administrator the capability of importing/filtering the specific routes from BGP into RIP domain if corresponding aggregate route exists in the BGP routing table.
- **Default BGP to RIPv2 import policy:** The default policy includes most of common policies that the user needs. In case the network administrator does not provide any explicit route importing policy, the default route importing policy is used.

BGP to RIPv2 Route Importing Support: Once the redistribution of BGP routes into RIPv2 is allowed, Vanguard router imports the routes in the local AS. By importing the BGP routes into RIPv2 domain the route importing policies are applied.

- **Importing the BGP routes in batch:** This normally happens when BGP routes are imported into RIPv2 in the beginning.
- **Importing/withdrawing the BGP routes in incremental way:** This normally happens after the initial importing to keep the RIPv2 routing table consistent with BGP routing table.
- **Importing only BGP aggregate route:** If BGP aggregate route is in the BGP routing table, the network administrator can choose to import the specific route or not.

RIPv2 Route Aggregation: The redistributed routes take part in the RIPv2 route aggregation process to reduce the size of RIP routing table. The user should have the CTP interface to control whether a redistributed BGP route takes part in the RIPv2 route aggregation process.

- **Aggregation of redistributed BGP routes into RIPv2 routes:** The redistributed BGP routes are combined into an aggregate route if possible. The aggregate route bears the mapped BGP route attributes if the user allows the Vanguard to do so. If a redistributed BGP route cannot be aggregated, then the route is stored in the RIP routing table as is.

- De-aggregation of redistributed BGP routes from RIPv2 aggregate route: The BGP route can be withdrawn anytime in the BGP routing table and the withdraw is reflected in the RIPv2 aggregate route. The aggregate route is changed to either more specific aggregate route(s) or no aggregate route.
- Route aggregation user interface: In the CTP menu of route import policy of BGP to RIPv2, one parameter is provided to let the network administrator to choose if the imported routes is eligible to be aggregated. If eligible, the route is considered in the RIPv2 aggregation process.

RIPv2 Route Advertisement: The redistributed BGP routes are advertised in the local AS. RIPv2 must set the tag of BGP routes and aggregate routes which include the BGP routes properly. BGP routes cannot be advertised in RIPv1 protocol.

- Periodical updates: Redistributed BGP routes are included in RIPv2 periodical update packet in the local AS.
 - Triggered updates: For incrementally changed BGP routes, RIPv2 sends the triggered update packet in the local AS.
 - RIPv2 route request and response: Imported BGP routes are included in the RIPv2 response packet no matter whether the request packet is for the whole routing table or some specific routes.
 - RIPv2 route tag: The route tag of RIPv2 route in the update packet should be set to the value of stored route in the routing table. This value is controllable by the network administrator. The RIPv2 route tag should not be lost in the chain of RIPv2 update process.
 - Default Route: The user can select the option of sending a default route with specific routes if BGP routes exist.
-

BGP to RIPv2 Redistribution Functions

BGP to RIPv2 route redistribution major functions are:

- BGP to RIPv2 Route Redistribution Policy
- BGP to RIPv2 Route Importing
- RIPv2 Aggregation
- RIPv2 Advertisement

BGP to RIPv2 Route Redistribution Policy

Route Redistribution Policy

BGP route redistribution policy is used to set the criteria that is used by the route importing process. If a BGP route satisfies the criteria, the route is imported to RIPv2 routing domain. A set of single route redistribution policy forms the router's whole policy of BGP to RIPv2 route redistribution. Each single policy is an element of the router's whole policy set. Stating that BGP route satisfies the route importing criteria indicates that the BGP route meets the requirement of at least one single policy in the router's whole policy set. If a BGP route satisfies one single policy, then the route satisfies the criteria. This route can be imported into RIPv2 routing domain if the action type of the policy is Permit.

There are two different levels of policy in the router:

- Non-default policy
- Default policy

Both policies are elements of the router's whole policy set. The default policy exists in the node and works once BGP to RIPv2 route redistribution is enabled. Although you can change the parameters of default policy, it is not required to configure explicitly. The non-default policy is not stored in the router unless you configure it explicitly. If the network administrator wants to enforce more policies in traffic engineering, he must configure non-default policies to implement it. If both non-default policy and default policy exist in a router, non-default policy always takes precedence.

When the user configures the policy, several levels should be provided to the user. The user can choose what kind of routes are selected, based on the IP address and the AS number(s) in the AS path attributes of the route. The user should be able to choose to import the route type, such as, whether a specific route is selected when an aggregate route is there. The user should be able to choose the action on the selected route, either permit the route to enter RIPv2 domain or deny.

If a route is permitted to enter the RIPv2 domain, the user should have ability to set the metric and tag fields of RIPv2 route. The user can choose to set them to a fixed value to control the traffic flow or let the router calculate the value. For the tag field the user even could disable the value setting.

The above mentioned requirements must be reflected in the configured menu in CTP. Users may configure policy through CTP.

■Note

The router should support as many as 1,024 non-default policy entries.

BGP to RIPv2 Route Importing

Route Importing

BGP to RIPv2 route importing is a key component of BGP to the RIPv2 route redistribution feature. Route importing is responsible for searching the BGP routing table, selecting the qualified route entries, translating the qualified route entries from BGP format, putting them into the global routing table, and setting the proper flags for these imported entries. From the global routing table, RIPv2 can advertise the imported BGP routes in the local AS.

There are two route importing methods:

- Batch input
- Incremental input

Batch and incremental input methods are supported for the BGP to RIPv2 route redistribution feature. In batch input method all the route entries in BRT are checked. Based on the checking result, the qualified route entries are input in RIPv2 routing table in block-by-block manner until all the qualified routes are done. In each block, many routes are included (each block can contain as many as 24 qualified routes). The reason why batch input method is introduced is that the number of RIPv2 triggered updates can be drastically decreased compared to the incremental input method. In incremental input method only the modified (changed/added/deleted) routes are put into the RIPv2 routing table in one-to-one way until all the qualified routes entries are done. If the BGP routing table is very large, it's not a good idea to go through the whole routing table when only one entry is modified. Normally in a stable network, a route entry is not modified frequently. By utilizing the incremental input method, the router could avoid a lot of resource consumption.

Not every BGP route is qualified to enter the RIPv2 routing domain. The BGP routes must satisfy a few predefined conditions. First, the BGP routes that were input from the local AS are not eligible to be re-input back to IGP again. That indicates the originating AS of a BGP route is equal to the configured local AS number then this route is not considered in route input process.

To qualify to be imported to RIPv2 domain, a BGP route must satisfy the following conditions:

- Not imported from the local AS
- Meet at least one single policy in the router's whole policy set
- Is an aggregate BGP route or is a specific route and specific route is allowed to be selected
- The action value of the matched policy is Permit

If a BGP route passes all the checks, then the action is used to determine if this route should be permitted to enter RIPv2 routing domain or should be denied. If a route is allowed to enter RIPv2 domain, a flag is set to indicate that a BGP route is input to RIPv2.

To add the imported BGP route through RIPv2 in the GRT, several RIPv2 fields have to be filled according to the attributes of BGP routes:

- The IP address field (the destination IP address) is copied directly from the destination IP address field of BGP route.
- The subnet mask is copied directly from the mask field of the BGP route.
- The next hop is copied from the BGP route.
- The metric is set according to the user's configuration:
 - If the configured value is between 1 and 15, this value is copied to the metric field.
- The route tag field is set according to user's configuration:
 - If the configuration is Disable or Manual but has invalid value, set the value as 0,
 - If the configuration is Manual and has a valid value, set the value as the configured one,
 - If the configuration is AdvertisingAS/OriginatingAS_Automatic, set the value as the advertising/originating AS number.

If the network administrator put a value between 1-15, then this value is copied to the metric field. If automatic is chosen then the number of the originating AS or of the advertising AS of the BGP route are copied. The next hop field should be set as (1) IP address of the peer's router interface from which the best path of the BGP route is relayed, if the peer is IBGP peer, or (2) IP address of the router interface from which the best path of the BGP route is received for all other cases. When RIPv2 receives these imported BGP routes, RIPv2 should not check from which interface this route came. RIPv2 should set the corresponding flag to indicate that this route was imported from BGP.

When a BGP route is modified (deleted or changed) and this route is input to RIPv2 (by checking the proper flag), this modified route is input to RIPv2 to make corresponding modifications in local AS. If the BGP route is deleted, the corresponding imported route should be deleted as well by setting the metric as 16. If the route has changed, the metric and/or tag of imported route in RIPv2 must be changed accordingly.

RIPv2 Route Aggregation

RIPv2 Route Aggregation

RIPv2 route aggregation is used to combine many specific routes into an aggregate route entry. RIPv2 route aggregation is an existing feature in Vanguard routers. The introduction of BGP to RIPv2 route redistribution in Release 6.3 and greater, add new requirements to the RIPv2 route aggregation algorithm.

If the user set the imported routes as not eligible to be aggregated in BGP to RIPv2 route redistribution policy configuration, the imported routes do not take part in the route aggregation process. This check should be done when route aggregation is trying to form the aggregate route.

If the imported routes can take part in the aggregation process, the aggregate route carries the route tag of the specific route. This is required because of BGP-IGP route synchronization. The other BGP speaker in the local AS could not see the IGP route if aggregate route is used and BGP-IGP does not recognize the aggregate route. By carrying the route tag in the aggregate route, the other BGP speaker could determine if an aggregate IGP route can be considered as the IGP route or not.

RIPv2 Advertisement

RIPv2 Advertisement

The RIPv2 Advertisement is used to send RIPv2 request and response packets. For the imported BGP routes the RIPv2 advertisement treats them as normal routes except that the tag value is copied into the packet and then the value is installed into the routing table.

BGP to RIPv2 Redistribution Typical Application

BGP/RIPv2 Interaction

A typical application for BGP to RIPv2 is shown in Figure 2-7. If RIPv2 is the only IGP in an AS and BGP is the EGP of the AS, then either default route or static routes have to be configured in the routers within the AS so that traffic can be correctly routed beyond the AS boundary. The scaling problem makes the routing table in each router within AS overflow and make the management task a big problem if no RIPv2 aggregation is applied to imported BGP routes. Enforcing network wide policies like controlling the traffic through the network or enforcing a particular routing path becomes more difficult.

The problem can be solved by enabling the BGP to RIPv2 route redistribution and configuring appropriate route import policies. The BGP routes are checked against the configured policies and imported to the local AS so that there is no need to configure the default route and/or static routes. The RIPv2 automatically combines the imported BGP routes into an aggregate route (if possible). The scaling problem in the AS could be solved. By configuring the different policies, the network wide management policies could be enforced to control the traffic, making network management more favorable.

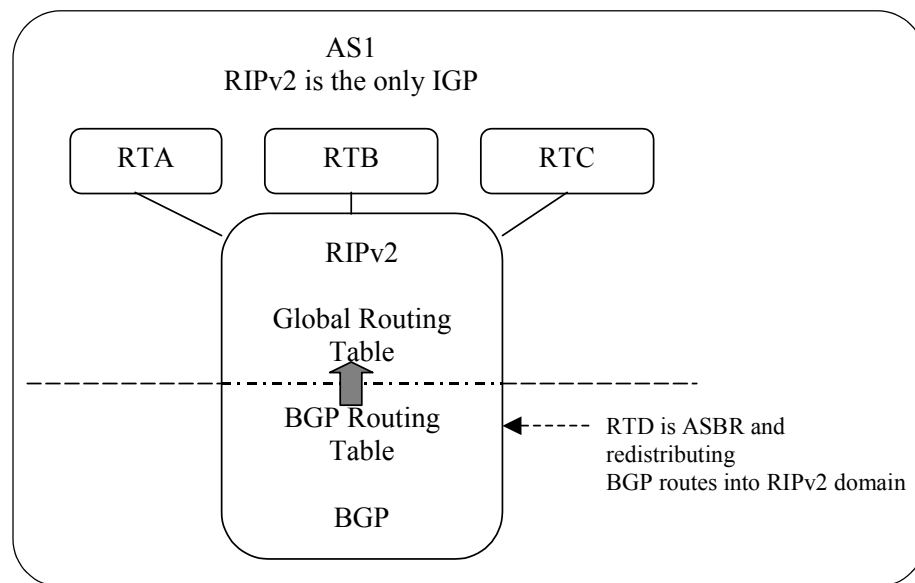


Figure 2-7. BGP Routes Imported into RIPv2 Domain

RIPv2 to OSPF and BGP Interaction Example

OSPF and RIPv2

It is common that both OSPF and RIPv2 are the IGP of an AS and this AS uses BGP as its EGP. Both OSPF and RIPv2 are capable of importing BGP routes into IGP, this may cause a problem. The routes imported by RIPv2 could be imported by OSPF again, and these routes could be re-imported back to RIPv2 domain. These routes could be redistributed back to BGP.

There are two solutions available to the network administrator:

- 1) The network administrator disables the BGP to OSPF route redistribution and solely depends on the BGP to RIPv2 route redistribution. BGP to OSPF route redistribution does not provide more substance than BGP to RIPv2 (the preferred route redistribution method should be BGP to RIPv2). If only one of two route redistribution conducts is enabled, the imported routes are obtained by the other IGP.
- 2) The network administrator configures the BGP to RIPv2 and BGP to OSPF policies properly, especially the metric and tag of the imported route. Whenever there is a disagreement between the BGP to RIPv2 and BGP to OSPF routes, the Vanguard chooses the more preferable route as the current path, based on the metric and tag, and make the less preferable route as the secondary path. The disagreement is solved and only the more preferable route is used in the local AS. Using the tag information of the routes, the Vanguard can check it against the ASBRs of the local to prevent these routes from re-distributing to the BGP routing domain.

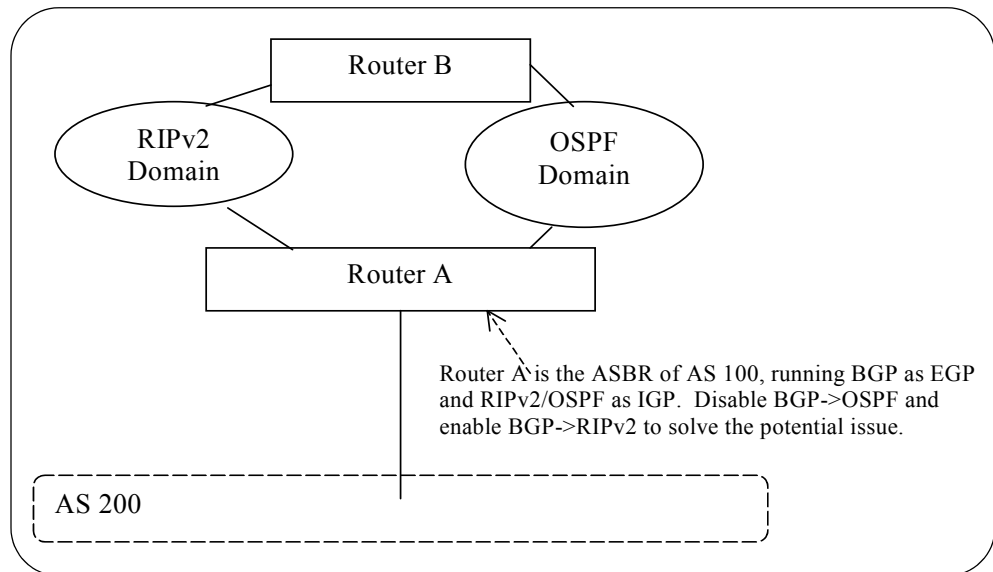


Figure 2-8. One Route Importing Method - Both OSPF and RIPv2 Co-exist in the Local AS

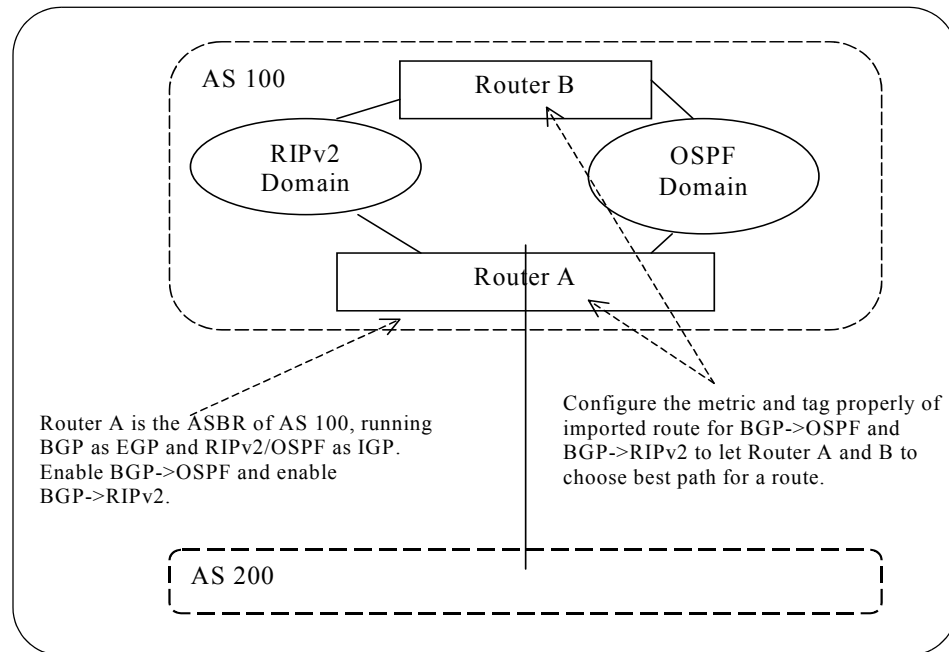


Figure 2-9. Configure Router A properly when both OSPF and RIPv2 Co-exist in local AS

Impacts of RIPv2 Route Redistribution on Other Features

Route Overwritten When using BGP to RIPv2 route redistribution sometimes one route created by IGP could have an identical destination as the BGP route. In this case, if the BGP route could be imported to IGP, and the IGP route could be overwritten. Some of the IGP routes are replaceable, while others are not. In case of possible overwriting, the following rules are applicable.

- Static and direct routes cannot be overwritten by BGP routes.
- BGP dynamic routes can be prioritized over the non-default dynamic routes, created by IGP, if the user enables the "BGP to RIP Nondefault Route Override" in **Configure->Router->Configure IP->IP Parameters**.
- The default BGP route can only be prioritized over the Default IGP route if the user enables the "BGP to RIP Default Route Override" in **Configure->Router->Configure IP->IP Parameters**.

BGP - IGP Synchronization

The BGP-IGP synchronization is a feature which is used to synchronize the routes advertisement in a local AS. A BGP route can be advertised in two ways, through IBGP and through IGP (RIPv2 or OSPF). The IBGP speaker does not send a route to its EBGp peer until it receives a route with the same destination network, which is advertised by IGP.

RIPv2 aggregation may cause the IBGP speaker to not find a route with the exact same destination. The IBGP speaker cannot send the route to its EBGp peers forever. The BGP-IGP synchronization was modified. If the IBGP speaker finds a route which is aggregate route, the destination is a super-network of a BGP route and the tag information agrees with the BGP route, the IBGP speaker sends the BGP route to its EBGp peers.

BGP Aggregation

Introduction

This section describes BGP aggregation. This feature is available with release 6.3 or greater software.

Using CIDR as its cornerstone, BGP aggregation aggregates the routes in its routing table and send out the aggregated routes to its peers. According to RFC 2519, two major components are suggested by the standard:

- 1) The first component is that route aggregation is from the originating AS. In its outbound route announcement, each AS aggregates the BGP routes originated by itself, by dedicated AS, and by private ASs. Although the proxy aggregation, which refers to route aggregation done by an AS other than the originating AS, is also supported, it is not recommended.
- 2) The second component is that BGP community "no-export" is set toward upstream providers. In its route announcements toward its upstream providers, an AS tags the BGP community "no-export" to routes it originates that do not need to be propagated beyond its upstream providers. Vanguard BGP implementation currently does not support community attributes, this BGP community "no-export" toward upstream providers is not supported. What the local BGP speaker sends to its peers is the aggregate route entries and explicit configured specific route entries. The BGP route aggregation is only based on the manually configured profiles. Automatic aggregation is not supported.

With the user configured route aggregation profiles, the BGP speaker checks every new route which is originated by its own AS. If more than one route can be aggregated into one, then an aggregated route is formed. This aggregated route is reflected in its update packet; the more specific routes are not included in the update packet unless explicitly configured as specific routes by user. That indicates the Vanguard BGP sends out a NLRI entry with `ATOMIC_AGGREGATE` and `AGGREGATOR` in its path attributes.

Multi-homed AS aggregation and AS path attributes loss are major subjects in BGP aggregation. Multi-homed AS is the AS which connects to two or more ISPs. If not processed properly, aggregation with multi-homed AS could generate a black hole in the routing table. For multi-homed AS, the user needs to configure the corresponding specific routes, and specific routes are advertised to upper layer's ISP.

Aggregation could cause some AS attribute loss if proxy aggregation is enabled in a BGP speaker. The originating AS information could be lost for the aggregated route. For proxy aggregation, if the aggregate route is formed from several routes with different AS path attributes, monitor the configuration to keep the attributes.

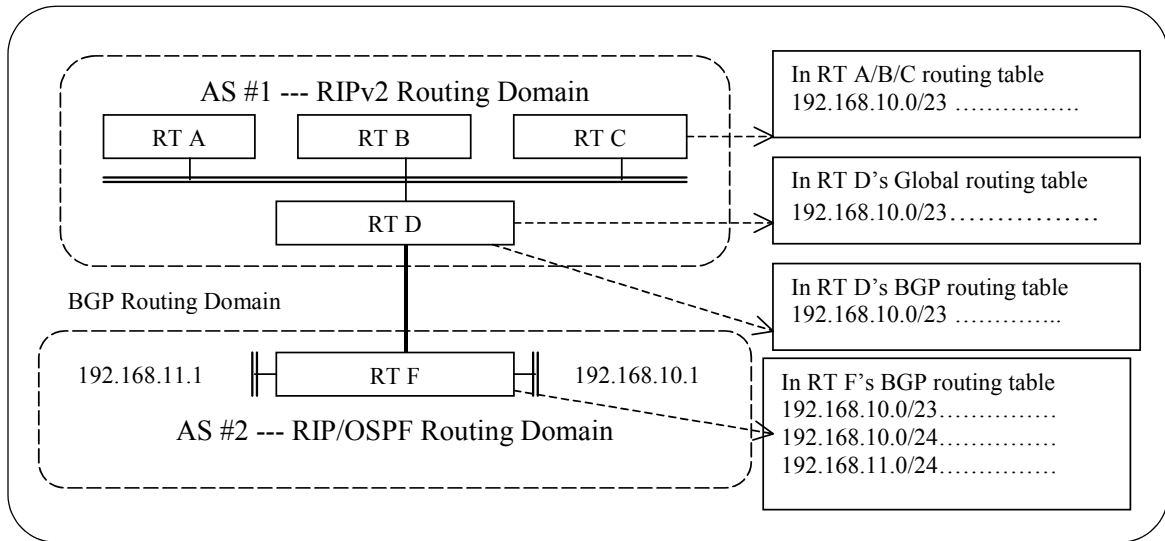


Figure 2-10. BGP Aggregation Example

The example in Figure 2-10 shows how BGP aggregation works. RT F is configured to aggregate 192.168.10.0/24 and 192.168.11.0/24 to 192.168.10.0/23. When 192.168.10.0/24 and 192.168.11.0/24 are added to RT F BGP routing table, a new route is created (according to configuration, 192.168.10.0/23). 192.168.10.0/23 is an aggregate route and this route is sent to RT D in AS #1, while 192.168.10.0/24 and 192.168.11.0/24 is not sent to RT D.

BGP Aggregation Features

Features

BGP aggregation supported in the Vanguard Router:

Aggregation Profiles Support: The BGP aggregation profiles allow the network administrator to control what kinds of routes are aggregated. The network administrator has several levels of controlling whether specific routes are sent out and proxy aggregation is enabled.

- **Configurable aggregate network:** The network administrator is allowed to configure the aggregate network. The aggregate network, which is a range of IP addresses, is used to match to more specific networks. If there is more than one specific network, which falls in the range of the aggregate network, the aggregate network is created and advertised to the other BGP peers.
- **Configurable more specific component networks:** The network administrator is allowed to configure a set of more specific component networks based on which aggregate network is created. The network administrator can choose a set of more specific component networks or exclude a set of more specific networks for an aggregate network.
- **Configurable proxy aggregation:** Although proxy aggregation is not recommended in the Vanguard router, it is configurable. By configuring this parameter the network administrator has the right to control if a Vanguard router aggregates a route that is not originated by its own AS. Vanguard BGP routers can act as the aggregation proxy of a third party product.
- **Configurable attribute of aggregate route:** The network administrator is allowed to change the attribute value of the aggregate route to implement traffic engineering strategy.
- **Configurable specific routes:** The network administrator is allowed to configure which specific routes are sent out to other AS and which routes are suppressed. By configuring these parameters the users can handle complex network layout.
- **Configurable suppressed route:** The network administrator is allowed to configure which specific routes are suppressed.

Aggregation and De-aggregation Support: When a new BGP route is added to BGP routing table, the route is compared to the configured profiles to determine if an aggregate route should be formed or not. If an aggregate route is formed then a new aggregate route entry is created and added to routing table. The attribute of a newly formed route is filled accordingly. If a route is withdrawn from the routing table, then a de-aggregation algorithm is applied. If one specific route is corresponding to an aggregate route, then the aggregate route entry is deleted from the routing table.

Full ATOMIC_AGGREGATE and AGGREGATOR Support: Currently, the Vanguard BGP implementation only supports receiving of ATOMIC_AGGREGATE and AGGREGATOR attribute. With the introduction of BGP aggregation, ATOMIC_AGGREGATE and AGGREGATOR attribute is fully supported; that is, Vanguard send ATOMIC_AGGREGATE and AGGREGATOR attributes if needed.

AS path attributes preserved: By forming an aggregate route, the originating AS or the AS information before the proxy aggregation could be lost. By using more AS attribute information in the routing table and updates packet, Vanguard can preserve the AS path attribute successfully.

BGP Aggregation Functions

BGP aggregation functions:

- BGP Aggregation Rules
- BGP Route Aggregation
- BGP Route Update

■ **Note**

BGP-4 supports CIDR. This reduces the number of routes in the routing table.

Aggregation Rules and Routes

Rules and Routes

Aggregation applies to routes which exist in the BGP routing tables and can be applied if more than one specific route of the aggregate exists in the BGP routing table. Aggregation adds a new level of complexity when deciding what routes to advertise to peers and how to forward packets because of a loss of information associated with route summarization. The following general rules must be adhered to for aggregation to work properly.

- Longest match - forwarding is done based on the longest match (most specific route) found in the routing table.
- Aggregation can only be performed if (at least) more than one specific route exists in the routing table.
- The IP addresses must be assigned on hierarchical or topological lines for aggregation to have its optimal benefits.
- Destinations which are multi-homed relative to a routing domain can be explicitly announced into that routing domain.
- In order to prevent routing loops, routers that aggregate multiple routes must discard packets which match the aggregate route but do not match any of the explicit routes which make up the aggregate.

The following situations describe the update options which must be supported.

- 1) **Aggregate Only, Suppress More Specific:** An aggregate is advertised and all of its specific routes are suppressed. This is usually done when the more specific routes do not offer any extra benefits such as making a better forwarding decision when forwarding traffic. An example of this is a single homed network
- 2) **Aggregate Plus More Specific Routes:** Situations exist in which it is beneficial to advertise the aggregate and its more specific routes. This usually occurs when a customer is multi-homed to a single provider. The provider would use the more specific routes to make a better decision when sending traffic towards the customer. At the same time the provider can propagate the aggregate only towards the NAP to minimize the number of routes leaked to the Internet. The use of this method allows the provider to balance the load to the customer. This is accomplished by sending different metrics for different routes on each of the links. In a backup situation with a primary/secondary topology, the forwarding of specific routes on the primary causes all traffic to be sent on this link unless the link fails and the routes are withdrawn from the routing domain.

3) Aggregation With a Subnet of More Specific Routes: In some situations a subset of more specific routes need to be advertised in addition to the aggregate. This is used to direct certain traffic to a specific AS such as in the case of multi-homed, geographically dispersed networks. This type of configuration allows the administrator to direct traffic to routes closer to the user in very large ASs. In some situations it is required that the attributes of an aggregate be changed.

ORIGIN Attribute If at least one route among routes that are aggregated has ORIGIN with the value INCOMPLETE, then the aggregated route must have the ORIGIN attribute with the value INCOMPLETE. Otherwise, if at least one route among routes that are aggregated has ORIGIN with the value EGP, then the aggregated route must have the origin attribute with the value EGP. In all other cases, the value of the ORIGIN attribute of the aggregated route is INTERNAL (IGP).

AS_PATH Attribute If routes to be aggregated have identical AS_PATH attributes, then the aggregated route has the same AS_PATH attribute as each individual route. When proxy aggregation is being performed (routes being aggregated from different ASs) then the AS_PATH attribute can be

- The AS number of the router doing the aggregation, if the type is AS_SEQUENCE.
- An unordered set of ASs that the aggregate has formed from with the aggregating routers, with its own AS number in the last position

The detailed information that existed in the specific prefixes are lost when summarized in the form of aggregates. The purpose of AS_SET is to try to save the attributes carried in the specific routes.

Without AS_SET the aggregate formed is considered to have originated from the AS in which it was formed and all the specific attributes are lost. With the use of the AS_SET, type the specific attribute information is retained .

**ATOMIC_
AGGREGATE
Attribute**

The ATOMIC_AGGREGATE attribute is used to indicate a loss of AS_PATH information if an aggregate is formed. The sources of the aggregate can have different attributes. If a system propagates an aggregate that causes a loss of AS_PATH information then it is required to attach the ATOMIC_AGGREGATE attribute to that route.

If a BGP speaker receives an UPDATE with the ATOMIC_AGGREGATE attribute set, it must not remove the attribute from the route when propagating it to other speakers.

**AGGREGATOR
Attribute**

The AGGREGATOR attribute specifies the autonomous system and the router that has generated the aggregate.

BGP Aggregation Examples

The following figures show some aggregation examples.

BGP Aggregation Scaling

BGP Aggregation

Problem: A very large network can have a scaling problem in terms of routing traffic update processing and routing table overflow. The network size makes the management of network more complex, and it becomes more difficult to keep the router running efficiently, especially increased traffic throughput and reasonable resource consumption of the whole network.

Solution: As stated in the CIDR and route aggregation RFCs, the size of the routing table causes problems in this application scenario. By utilizing BGP aggregation, the size of the routing table could be decreased significantly and the resource consumption can be restricted to reasonable level. This solves the scaling problem.

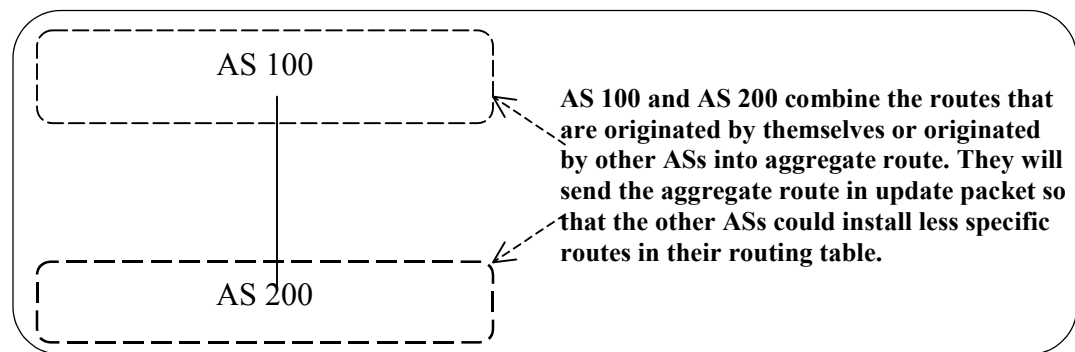


Figure 2-11. BGP Route Aggregation

BGP Aggregation with Multi-Homed AS

Multi-Homed AS

Problem: An enterprise network can connect to multiple ISP's for enhancing its network connectivity in terms of load balancing on multiple paths or creating multiple backup paths for a particular set of external networks. If proxy aggregation is enabled in the ISP sites, route aggregation could generate some routes that include black holes. Destination Based Routing (DBR) uses the longest match algorithm to find the next hop, traffic could be forwarded to the wrong network and discarded.

Solution: In this case the network administrator can configure the multi-homed AS combining with specific subnet routes. The specific routes associated with the multi-homed AS in the upper layer ISPs would be found using the longest match algorithm and traffic could be forwarded properly.

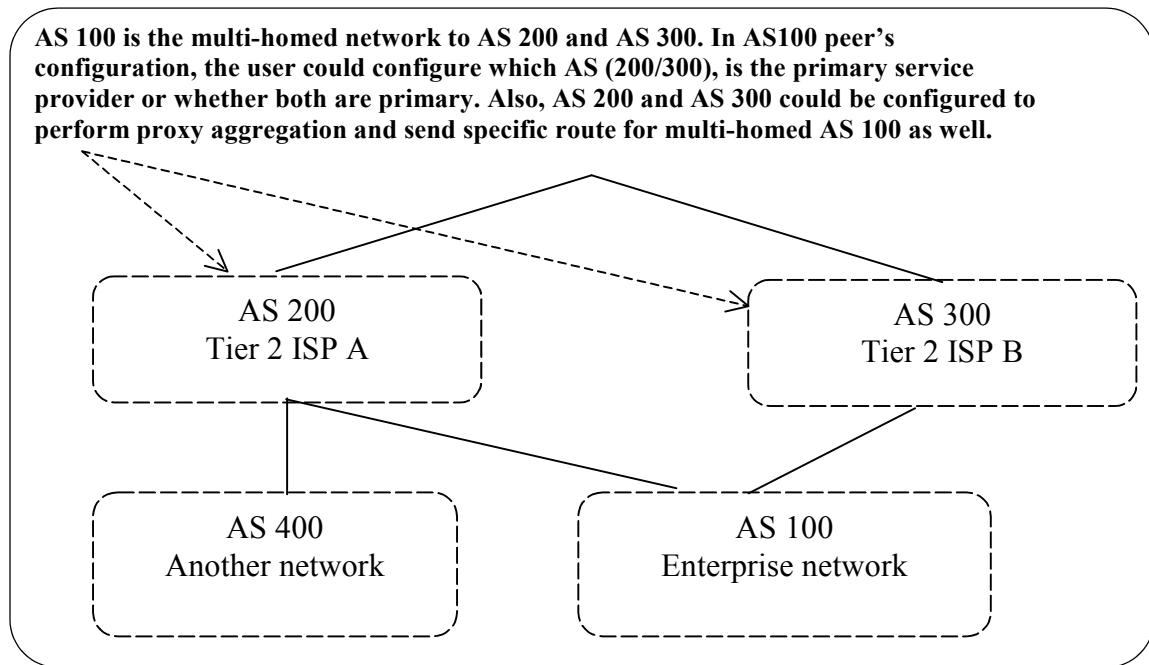


Figure 2-12. BGP Aggregation with Multi-Homed AS

■ **Note**

When AS 1000 is set to do proxy aggregation, BGP routes sent by AS 200 and AS 300 could be combined into an aggregate route. Make sure the aggregate route keeps the AS information, 200/300, in the path attributes.

BGP Aggregation with Specific Subnet

Specific Subnet

Problem: If the user's network migrates from one ISP to another ISP, the IP addresses of the user's network still use the allocated IP addresses from the previous ISP. If the user's network was setup even before the CIDR and aggregation concept was introduced, the IP addresses of user's network might be not in a continuous range. As stated in the RFC of CIDR and route aggregation, re-allocation of IP addresses for the user's network is recommended. However, the users may not want to change the IP addresses of their network due to various reason. Thus these situations give BGP aggregation big trouble.

Solution: Vanguard BGP aggregation provides the network administrator the capability to configure specific routes to handle this problem. The IP addresses of a user's network could be sorted to different ranges; some of them could be more generic and some of them could be more specific. BGP aggregation combines the generic addresses ranges into more generic aggregate routes, and makes more specific address ranges into more specific aggregate routes. Both the more generic and specific aggregate routes are sent in the BGP update packet. With the more specific routes supported, the above mentioned problems is solved.

Another benefit of using the specific routes is that the network administrator can implement the traffic engineering policies by sending and suppressing specific routes in different AS exits.

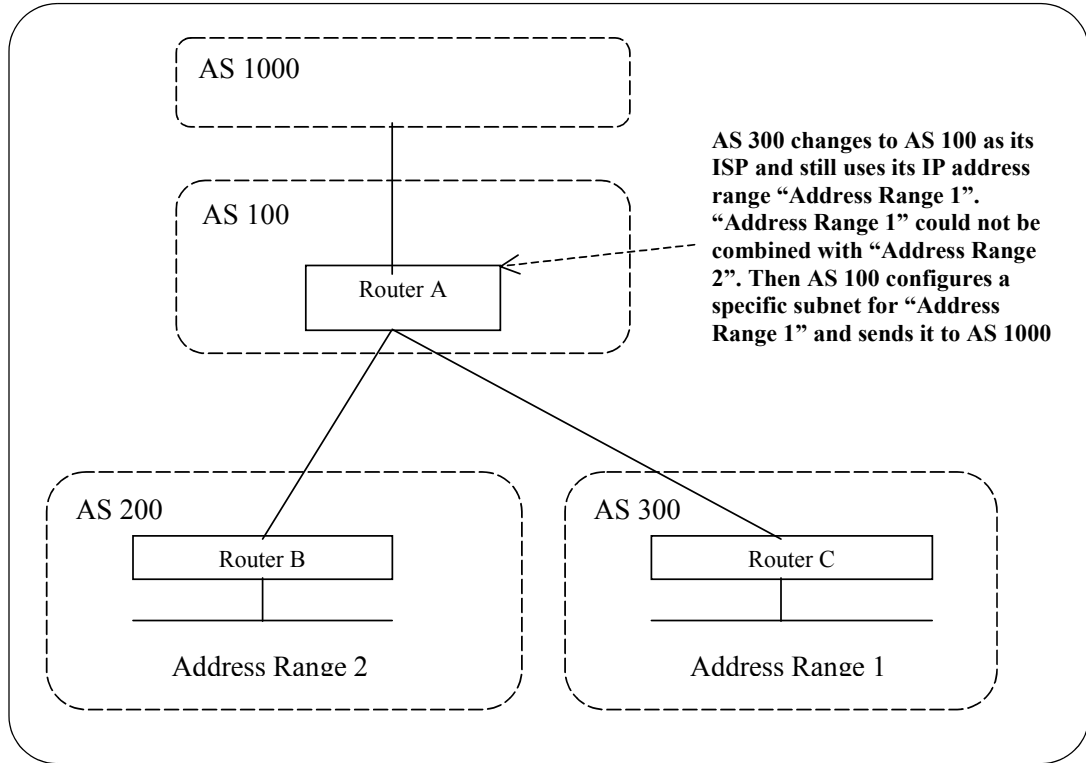


Figure 2-13. BGP Aggregation with Specific Subnets

Impacts of Aggregation on Other Features

BGP Route Table and Packet Forwarding

BGP aggregation creates some new aggregate routes in a BGP speaker which perform the BGP aggregation and BGP proxy aggregation. The aggregate routes, when installed in the routing table and advertised in the network, are used as normal routes.

BGP Route Update

The creation of the BGP aggregate routes also has implications to the BGP update process. If less specific routes are not allowed to advertise to BGP peers, creation of an aggregate implies that the less specific routes should be withdrawn from the peers' routing table. In addition to the normal BGP update packet, the implied withdrawn routes for the BGP peers should be sent by more BGP update packet. Second in the normal BGP update process, the new configured aggregation policies have to be considered in addition to normal output policies. For example: a specific route, which can be advertised according to the output policies, may not be sent in the normal BGP update packet if a specific route is not allowed to do so in the BGP aggregation configuration.

<i>Specific Routes and Aggregate Routes</i>	<i>Can be Advertised (BGP Output Policies)</i>	<i>Cannot be Advertised (BGP Output Policies)</i>
Specific Routes can be advertised (Aggregation Profile)	Send specific routes	Do not send specific routes
Specific Routes cannot be advertised (Aggregation Profile)	Do not send specific routes	Do not send specific routes
Aggregate routes	Send aggregate routes	Do not send aggregate routes

BGP Community Attribute

Introduction

BGP Community Attribute is a transitive optional attribute, which is mainly used to specify the preference of a path and restrict the advertisement of routes according to administrator's policies.

The BGP Community Attribute feature can be summarized as follows.

- Well-known values and user-defined values are supported.
- Community can be configured in BGP policies and aggregation profiles
- Community can be used to control the advertisement of routes.
- Community can be used to define a path's preference.
- Community values can be carried out in the aggregate component route.
- Configuration and statistics display are supported.

Community Attribute

BGP community attribute includes the support of well-known community and user defined values. The well-known values, NO_EXPORT, NO_ADVERTISE, and NO_EXPORT_SUBCONFED, are the basic requirement for our implementation. Other values are defined according to the mapping to users' routing policies. The well-known values are mainly used to control advertisement of routes, while user defined ones are mainly used for traffic engineering.

Community attribute values are organized as a set of four octets, with first two octets as the ASN and last two octets as one of the community values. The well-known values do not follow this rule because they have their predefined values. The path attribute for a route entry, which includes the community attribute, makes the route entry belong to all the communities defined in the attribute.

Since the community attribute values reflect the user's routing policies, community feature supports user's configuration through community profile. The user's policy configuration includes the community profile parameter. The community profile defined in the policy menu is applicable to all the inbound/outbound, AS/NET/AS+NET, and scope policy. The default value for the community profile is NONE, which means that no community is configured. In the same way BGP Aggregation includes the well-known community values, which define advertisement scope of the aggregate route and component routes. BGP community values are displayed in the BGP statistics display.

Community attribute values control the advertisement of routes. When BGP advertisement packet is received, the community values are checked in addition to the existing policies. When BGP routes are advertised to the speaker's peers, the community values in the outbound policies are checked in addition to existing policies. Proper actions, such as no advertisement to some (all) peers and adding community values, are carried out. For the well-known community values, even if the outbound policies are not configured; the attribute values in the path attribute are used to determine the advertisement of route entries.

A received community attribute is explained according to user's configuration. The attribute value is mapped to appropriate route aspects. One of the most common practices is to map the attribute value to the weight (cost, or preference) of route entry. The community value effects the relationship between a subscriber and its service provider.

Community attribute effects the BGP aggregation. In addition to the advertisement of aggregate route and component routes, community attribute values must be fully added to the path attribute of the aggregate route if the ATOMIC_AGGREGATE is not configured for the aggregate route.

Application Examples

Route Advertisement

Route Advertisement is used when aggregation is configured in a user's node and the user would like some component to be advertised to its peer so that a more efficient route could be used for a specific destination. However, if the user does not want the component route to be advertised beyond that peer then the NO_ADVERTISE attribute could be configured for these specific routes. When the peer receives these specific routes, it will not advertise the specific route further.

Figure 2-14 demonstrates the application scenario and its solution.

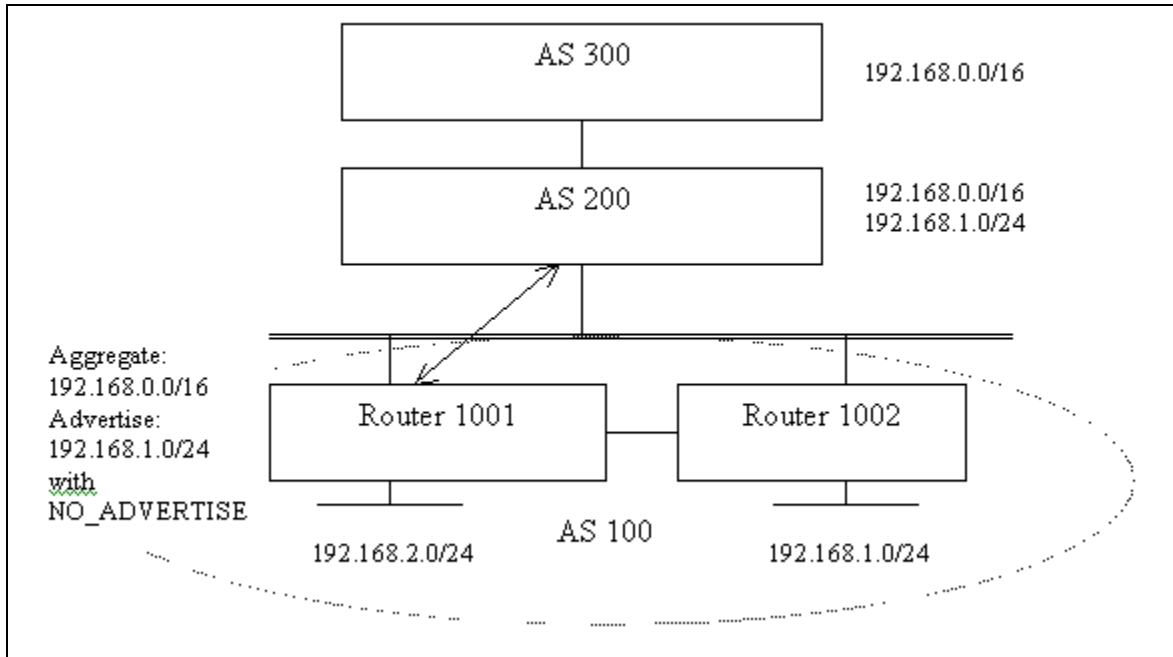


Figure 2-14. Route Advertisement Controlled through Community Attribute

Route Preference

Route Preference is used when two users subscribe service from two different service providers, and they also make a mutual agreement to backup each other. In a normal situation, they use their primary service provider to access the whole network. In case the primary service provider can't provide the service, the backup mechanism will be used. Obviously the primary path takes priority over the backup path. With community attributes' help, the application problem could be easily solved. Each user tags the routes from the other user with lower attribute value and its own route with higher attribute value, suppose lower attribute value represents lower preference. The service provider uses the high attribute value tagged route in normal case and lower attribute value in abnormal case.

Figure 2-15 shows the application scenario:

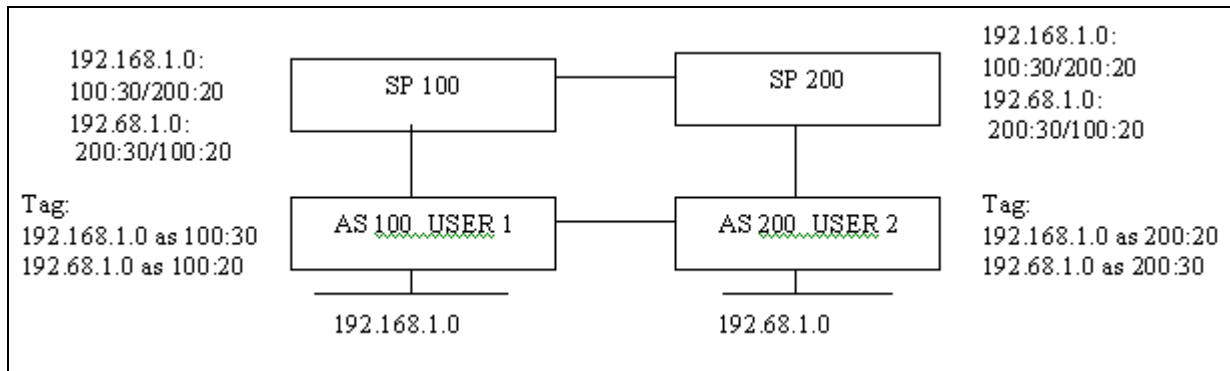


Figure 2-15. Route Path Preference through Community Attribute

Node Boot and Table Boot

A Node Boot is not recommended in the Vanguard Router, the BGP to RIPv2 Route Redistribution and BGP aggregation do not require a Node Boot even when configuration is changed.

The parameter's configuration changes require a table boot. The table boot checks the validity of configuration, copies the parameter from SRAM to DRAM, and performs associated action based on the changes.

RIPv2 Route Redistribution Boot

To boot BGP to RIPv2 Route Redistribution:

Boot->Boot Router->Boot IP->"Boot BGP->RIP Import Policy Table"

The above actions boot the BGP to RIPv2 Route Redistribution.

Aggregation Boot

To boot BGP Aggregation:

Boot->Boot Router->Boot BGP->BGP Aggregation

SNMP Network Management

SNMP

SNMP access is based on the objects defined in the control (.cont) and management information base (.mib) files. The Agent Action Functions (AAF) are used to return appropriate values based on the mib definitions.

Currently BGP4 is implemented under a public management MIB defined from an RFC (rfc1657a.cont and rfc1657a.mib). Existing CMEM configuration parameters are occasionally implemented under this RFC (about nine parameters out of a possible fifty). Since an RFC-based MIB cannot be changed, the old and new BGP4 configuration and statistical parameters are added to privately defined MIB files.

RFC1657

The current BGP MIB files, rfc1657a.cont and rfc1657a.mib, adhere to RFC1657 and include objects for the BGP Peer Table, the BGP Received Path Attribute Table and the BGP-4 Received Path Attribute Table. The RFC1657 Object Identifier (OID) tree is:

Rfc1657OID Tree Location	.iso.org.dod.internet.mgmt.mib-2.bgp
--------------------------	--------------------------------------

BGP4 Configuration

All current BGP4 CMEM configuration parameters (53 parameters under BGP Global Parameters, Peer Parameters and BGP Policies) are added as new Objects and Object Tables to the new BGP4 MIB files (bgp4_opt.cont and bgp4_opt.mib). This is to avoid increasing the Router MIB files by almost fifty percent. The new MIBs are created as a BGP Group under the Router Group. The SNMP MIBs for BGP are included in an image when the BGP module is included. The BGP4 Configuration Group Object Identifier is:

BGP4 Configuration Group OID Tree Location	.iso.org.dod.internet.private.enterprises.codex.cdxProductSpecific.cdx6500.cdx6500Configuration.cdx6500CfgProtocolGroup.cdx6500PCTRouterGroup.cdx6500PCTBgpGroup
--	--

BGP to RIPv2 Redistribution and BGP4 Aggregation Configuration

SNMP support for IP and RIP configuration parameters are privately defined in the router_opt.cont and router_opt.mib files. New BGP to RIPv2 configuration parameters, dealing with IP implementation are located as an extension to the existing IP Parameter configuration parameters. The Router Configuration Group Object Identifier is:

Router Group OID Tree Location	.iso.org.dod.internet.private.enterprises.codex.cdxProductSpecific.cdx6500.cdx6500Configuration.cdx6500CfgProtocolGroup.cdx6500PCTRouterGroup
--------------------------------	---

With the implementation of BGP to RIPv2, new Objects and Object Tables are added for all of the new BGP4-RIPv2 CMEM parameters, as described in this functional specification. This includes any new BGP Global Parameters, the BGP Redistribution Policy and the BGP Aggregation Tables (BGP Aggregate Profiles, IP Map Profiles and BGP Aggregate Attribute Profiles). These objects are added to the BGP Configuration Group OID, cdx6500PCTBgpGroup, as listed above.

Statistics

Statistics that are available through SNMP are shown under General Statistics. The BGP Statistics Group Object Identifier is:

BGP4 Statistics Group OID Tree Location	.iso.org.dod.internet.private.enterprises.codex.cdx- ProductSpecific.cdx6500.cdx6500Statistics. cdx6500StatProtocolGroup.cdx6500PSTRouterGroup .cdx6500PSTBgpGroup
--	---

Overview

Introduction

The BGP configuration consists of five tables. Users need to configure the following tables for basic BGP configuration:

Configure BGP:

- BGP Global Parameters
- BGP Peer Table
- BGP Import/Export Policy Table

Configure OSPF:

- AS Boundary Routing Parameters
- BGP->OSPF Import Policy Table

Configure RIPv2 Redistribution:

- BGP->RIPv2 Redistribution
- BGP->RIPv2 Redistribution Policy

Configure Aggregation:

- Aggregate Profiles
- IP Map Profile
- Aggregate Attribute Profiles

Configure Community Attributes

- Community Attribute Profile

Basic Routing Configuration

Before you can configure BGP-4 parameters you must configure the following:

- LAN and WAN ports
 - Route Selection Table
 - LAN Connection Table
 - Mnemonic Table
 - IP Router Interfaces
 - IP Router Parameters
-

Accessing the BGP Parameter Records

BGP Configuration Menu The configuration path for BGP tables is as follows. From the Control Terminal Port Main menu, select **Configure->Router->Configure BGP**. See Figure 3-1.

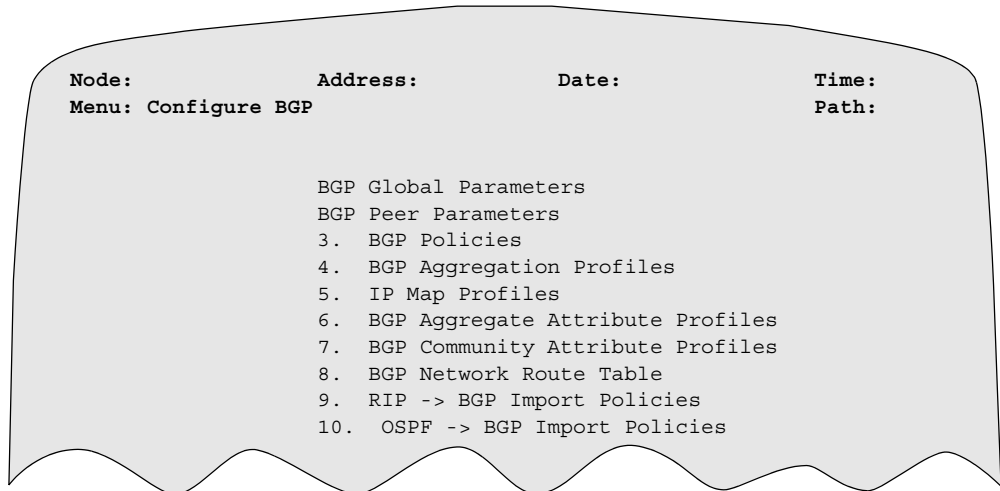


Figure 3-1. Configure BGP Menu

Configuration Tasks and Records This table correlates the common BGP configuration tasks with the configuration records:

To...	Complete This Record...
Enable or disable BGP in the node.	BGP
Specify the maximum number of peers that this router can have simultaneously.	Maximum Peers
Specify the AS number for local BGP speaker.	AS Number
Specify the BGP identifier for local BGP speaker.	BGP Identifier
Specify the size of BGP Routing Table.	BGP Routing Table Size
Specify the default import policy for different types of IGP routes.	Default IGP Import Policy
Specify the default import policy for Incoming BGP routes.	Default BGP Inbound Policy
Specify the default export policy for Outbound BGP routes.	Default BGP Outbound Policy
Control the support of atomic aggregate attribute by local BGP Speaker.	Atomic Aggregate
Control the display of advanced BGP parameters.	Advanced Parameter Display

To.... (continued)	Complete This Record...
Allow the synchronization of BGP with the IGP table before advertising the routes to the external peers.	BGP - IGP Synchronization
Allow configuration of the maximum segment size to be used for the active TCP connections of local BGP speaker with its peers.	Default TCP Segment Size
Controls whether a non-default BGP route can override the corresponding non-default IGP route	Override Non-Default IGP Route
Controls whether a default BGP route can override the default IGP route	Override Default IGP Route

Configuring BGP Routing Parameters

Introduction

This section describes how to configure the parameters to enable BGP routing.

What You See in This Record

Figure 3-2 shows the BGP Global Parameters.

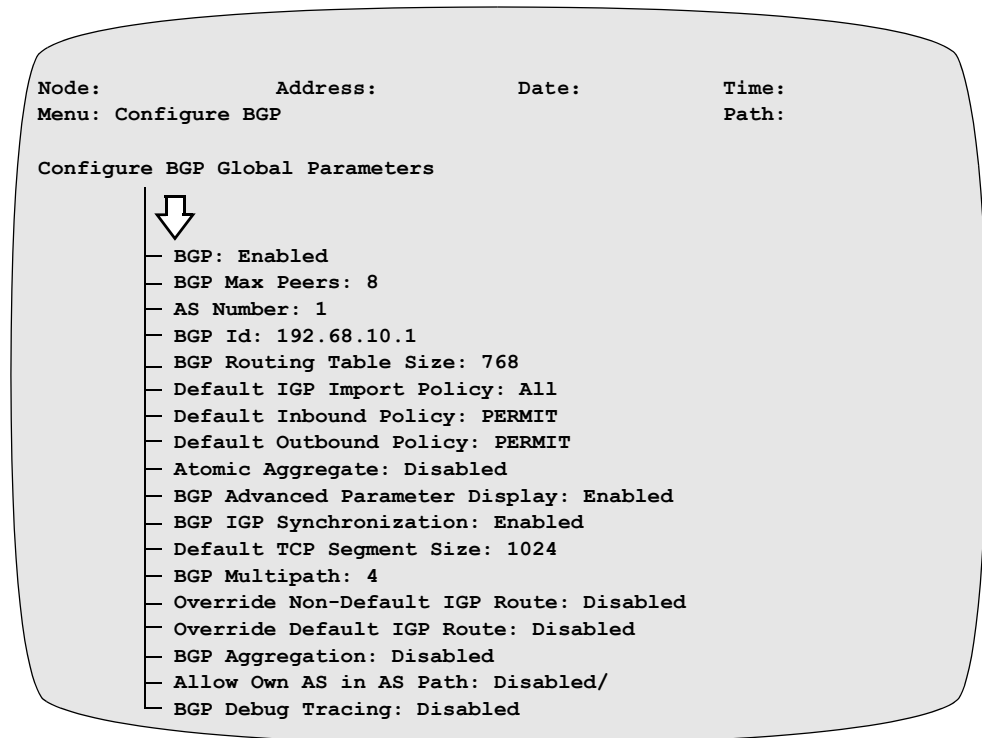


Figure 3-2. BGP Global Parameters

BGP Global Parameters

This tables describe the parameters that make up the BGP Routing Parameters record. Unless otherwise indicated, you must Boot BGP Parameters for changes to these parameters to take effect.

BGP

Range:	ENABLE/DISABLE
Default:	DISABLE
Description:	This parameter enables or disables BGP feature in the node.

Maximum Peers

Range:	1-128 (Vanguard 7300 Series) 1-16 (All others)
Default:	64 (Vanguard 7300 Series) 8 (All others)
Description:	This parameter specifies the maximum number of peering that local BGP speakers can have simultaneously. This controls the maximum number of entries in BGP peer table.

AS Numbers

Range:	1-65535
Default:	Blank
Description:	This parameter specifies the AS number for local BGP speaker.

BGP Identifier

Range:	Any valid IP address
Default:	Blank
Description:	This parameter specifies the BGP identifier for local BGP speaker. It should be one of the IP addresses belonging to the router.

BGP Routing Table Size

Range:	1024-10000
Default:	1024
Description:	This parameter specifies the maximum number of routes which can be stored in BGP Routing Table.

Default IGP Import Policy

Range:	DIR, STATIC, RIP, OSPF_INTERNAL, OSPF_EXT_TYPE1, OSPF_EXT_TYPE2, OSPF_ANY, ALL, NONE
Default:	Blank
Description:	<p>This parameter specifies the default import policy for different types of IGP routes. For all the IGP routes which do not match any filter in IGP import policies, this parameter defines the default import policy. This parameter controls the import of following types of IGP routes.</p> <p>Any combination of these route types can also be configured. For ex: RIP+STATIC, RIP+OSPF_INTERNAL+DIR.</p> <p>A value of ALL means that all types of IGP routes are imported. A value of NONE means that none of the IGP routes are imported.</p>

Default BGP Inbound Policy

Range:	PERMIT, DENY
Default:	DENY
Description:	<p>This parameter specifies the default import policy for Inbound BGP routes. For all the Inbound routes which do not match any policy entry in BGP policy table, this parameter defines the default import policy.</p> <p>A value of PERMIT indicates that all incoming BGP routes which do not match any policy entry in BGP Policy Table are imported into the BGP Routing Table with the following default values for the specified parameters.</p> <p>Path Weight = 0</p> <p>A value of DENY indicates that incoming BGP routes which do not have any matching entry In BGP Policy Table are not imported.</p>

Default BGP Outbound Policy

Range:	PERMIT, DENY
Default:	DENY
Description:	<p>This parameter specifies the default export policy for Outgoing BGP routes. For all the outgoing routes which do not match any policy entry in BGP policy table, this parameter defines the default export policy.</p> <p>A value of PERMIT indicates that all outgoing BGP routes which do not match any policy entry in BGP Policy Table is advertised with the following default values for the specified parameters.</p> <p>MED = 0</p> <p>Number of Extra AS Prepends = 0</p> <p>Local Preference = 0</p> <p>Explicit Advertisement = NOT_ALLOWED</p> <p>A value of DENY indicates that outgoing BGP routes which do not have any matching entry In BGP Policy Table are not advertised.</p>

Atomic Aggregate

Range:	ENABLE, DISABLE
Default:	DISABLE

Atomic Aggregate (continued)

Description:	<p>This parameter controls the support of atomic aggregate attribute by local BGP Speaker.</p> <p>When enabled, the local BGP speaker installs and advertises a less specific route when presented with routes (with overlapping networks). The less specific routes are advertised with Atomic Aggregate to the external peers.</p> <p>When disabled, both less and more specific routes are installed in the local BGP speaker's routing table.</p>
--------------	---

Advance Parameter Display

Range:	ENABLE, DISABLE
Default:	DISABLE
Description:	<p>This parameter controls the display of advanced BGP parameters in all the tables including this parameter record. The advanced parameters appear only when this parameter is enabled otherwise a default value is assumed for those parameters.</p>

BGP- IGP Synchronization

Range:	ENABLE, DISABLE
Default:	ENABLE
Description:	<p>This parameter enables or disables this feature in the node. When set to enabled the local BGP speaker waits for synchronization with its IGP table before advertising the routes learned from internal peers to external peers. When set to disabled it will not wait for synchronization and starts advertising the routes as soon as it learns from other peers.</p>

Default TCP Segment Size

Range:	512-65000
Default:	1024
Description:	<p>This parameter specifies the maximum segment size to be used for the active TCP connections of local BGP speaker with its peers. It can be used to reduce the CPU processing related to fragmentation and reassembly of TCP packets. This parameter is applicable for those peers only which do not have TCP segment size configured in BGP peer table.</p>

BGP Multipath

Range:	1-4
Default:	4

BGP Multipath

Description:	This parameter specifies the maximum number of paths for BGP load balancing.
Boot Type	A change to this parameter requires a BGP Global parameter Boot to take effect.

Allow Own AS in AS Path

Range:	ENABLE, DISABLE
Default:	DISABLE
Description:	This parameter specifies whether to allow the local AS number to be the originating AS number in the AS Path.
Boot Type	A change to this parameter requires a BGP Global parameter Boot to take effect.

BGP Debug Tracing

Range:	ENABLE, DISABLE
Default:	DISABLE
Description:	This parameter enables or disables BGP debug tracing.
Boot Type	A change to this parameter requires a BGP Global parameter Boot to take effect.

Override Non-Default IGP Route

Range:	ENABLE, DISABLE
Default:	DISABLE
Description:	This parameter controls whether a non-default BGP route can override the corresponding non-default IGP route (from any IGP source like Static, RIP, OSPF, etc.) when there are both IGP and BGP routes available for a particular destination network. This overriding rule determines the route to be used by the local border router for forwarding packets in above mentioned scenario.

Override Default IGP Route

Range:	ENABLE, DISABLE
Default:	DISABLE

Override Default IGP Route *(continued)*

Description:	This parameter controls whether a default BGP route can override the default IGP route (from any IGP source like Static, RIP, OSPF, etc.) when there are both IGP and BGP default routes available. This overriding rule determines the default route to be used by the local border router for forwarding packets when both non-default IGP and BGP routes are not available for a particular destination network.
--------------	---

BGP Aggregation

Range:	Enable, Disable
Default:	Disable
Description:	This parameter specifies whether BGP aggregation can run in the node. ■ Note If there are no BGP aggregation profiles configured, this parameter is set to Disable (no matter what value is configured).

■ Note

When BGP Aggregation is set to Enabled, the BGP Proxy Aggregation parameter appears.

BGP Proxy Aggregation

Range:	Enable, Disable
Default:	Disable
Description:	This parameter specifies whether BGP proxy aggregation can run in the BGP speaker or not. The BGP proxy aggregation is used to enable the BGP speaker to aggregate routes, even if these routes are not originated by this local AS. ■ Note If BGP Aggregation is set to Disable this parameter will not display.

BGP Peer Parameters

This table defines the parameters that are applicable to a particular peer only. They are used to form a BGP connection with a BGP peer. In this table the user can also define certain path attributes for the routes sent or received from this peer. The following table gives the overall description of each parameter in BGP Peer Table.

Entry

Range:	1 - MAX PEER
Default:	Increasing number
Description:	This parameter uniquely identifies a peer.

Peer Control

Range:	ENABLE, DISABLE
Default:	ENABLE
Description:	This parameter controls the BGP peering with this peer. When disabled, peering is not formed.

Peer AS Number

Range:	1-65535
Default:	Blank
Description:	This parameter specifies the AS number for this peer.

Peer IP Address List

Range:	A list of valid IP addresses
Default:	Blank
Description:	This is a comma separated list containing maximum of eight IP addresses belonging to this peer which are used to form a TCP connection with the peer. Forming a connection, an address is tried randomly from this list and if not successful then other addresses are tried one by one. Example: 10.10.10.10, 1.1.1.1, 130.1.1.1

Hold Down Time

Range:	0, 10-65535
Default:	90

Hold Down Time *(continued)*

Description:	This parameter specifies the amount of time (in seconds) that the BGP speaker waits before timing out the BGP connection. A value of 0 indicates that no keep alive is sent to this peer. This parameter should be configured in accordance with the Keep alive timer configured for this peer.
--------------	---

Keep Alive Time

Range:	3-60000
Default:	30
Description:	This parameter specifies the time interval (in seconds) between two successive keep alive packets to this peer. This value should be less than the value of the hold down timer. If configured to a value equal to or greater than the hold down timer, a value of 1/3 the hold down timer is used. ■ Note A lower value of this parameter generates more traffic.

Connection Retry Time

Range:	15-65535
Default:	60
Description:	This parameter specifies the amount of time (in seconds) that local BGP speaker should wait before making a second attempt to form a TCP connection with this peer. This parameter is applicable when local speaker goes to the ACTIVE or CONNECT state due to some error in the first attempt.

Connection Restart Time

Range:	0, 30-65535
Default:	120
Description:	This parameter specifies the amount of time (in seconds) that local BGP speaker should wait before it restarts the TCP connection process with this peer. This parameter is applicable when the local speaker goes to IDLE state due to some error. A value of 0 indicates that the local speaker will not restart the connection process itself. In this case the user has to restart the connection process by booting the peer table.

Authentication Type

Range:	SIMPLE_PASSWORD, NONE
Default:	NONE

Authentication Type *(continued)*

Description:	This parameter selects the authentication scheme used to authenticate the peer. Both the ends of the connection should be configured with the same authentication scheme. A value of NONE indicates that no authentication scheme is used for the connection with this peer.
--------------	--

Password

Range:	Character string consisting maximum of 16 characters
Default:	Blank
Description:	Used as the password data for peer authentication as well as all the communications with this peer. Both the ends of the connection should be configured with the same password. This parameter is applicable only when simple password authentication scheme is in use.

MD5 Password

Range:	1-25 alphanumeric characters, use the space character to blank field
Default:	Blank
Description:	Enter password for TCP MD5 Signature Option. A blank password means the feature is turned off.

#MinAS Origination Interval

Range:	0, 10-180
Default:	15
Description:	This parameter specifies the interval (in seconds) between two successive advertisement of locally originated routes to this peer. This is used to reduce the processing overhead of the remote BGP speaker resulting from learning new routes and running the decision process to incorporate them in its routing database.

#MinAS Advertisement Interval

Range:	10 -180,0
Default:	30
Description:	This parameter specifies the interval (in seconds) between two successive advertisement of external BGP routes to this peer. This parameter is applicable to external peers only. This is used to reduce the processing overhead of the remote BGP speaker resulting from learning new routes and running the decision process to incorporate them in its routing database.

Indirect BGP Peering

Range:	ALLOWED, NOT_ALLOWED
Default:	NOT_ALLOWED
Description:	This parameter controls the establishment of BGP connection with this peer even if it is not present on the same IP subnet as the local BGP speaker. If configured ALLOWED then peering is formed only when the local BGP speaker has an IGP or static route to the subnet on which this peer is present. If configured NOT_ALLOWED then peering is formed only when the local BGP speaker and the peer are present on a directly attached sub-network. This parameter is applicable to external peers only.

Note

Indirect peering should not be used when establishing a connection to an external peer. This could create adverse effects when importing BGP prefixes into an IGP protocol.

#TCP Segment Size

Range:	512 - 65535
Default:	1024
Description:	This parameter specifies the maximum segment size to be used for the active TCP connection with this peer. A value of zero indicates that Default TCP Segment Size is used as configured in BGP Global Parameters.

Peer Weight

Range:	0,1- 65535
Default:	0
Description:	This parameter specifies the weight associated with this peer. This is used as one of the factors to calculate the degree of preference for all the routes learned from this peer. A higher value of this parameter means that higher preference is given to the routes learnt from this peer. This parameter is applicable to external peers only. A value of 0 indicates that no special preference is given to the routes learned from this peer.

MED

Range:	0, 1-0xFFFFFFFF
Default:	0

MED

Description:	This parameter specifies the value of the Multi Exit Discriminator (MED) attribute with which routes are advertised to this peer. The routes with lower MED's are preferred over the routes with higher MED's. A value of 0 indicates that a MED's attribute is omitted from the updates being advertised to this peer. This parameter is applicable to external peers only.
--------------	--

Number of Extra AS Prepends

Range:	0, 1-10
Default:	0
Description:	This parameter specifies the additional number of times the local AS number is prepended in the AS path attribute of the routes generated for this peer. This can be used to influence the route selection process at remote BGP speaker for the route being advertised to this peer. This parameter should be set to a non-zero value only when the peer also support this feature. A value of zero indicates that this feature is not used. No extra AS prepends are done for the routes being advertised to this peer.

Local Preference

Range:	0,1-xFFFFFFFF
Default:	0
Description:	This parameter specifies the value of the LOCAL PREF attribute for the routes learned from this peer and being advertised to all internal peers. A value of 0 indicates that the locally calculated DoP is advertised as the local preference.

#BGP NON SELF Nexthop Advertisement

Range:	ALLOWED, NOT_ALLOWED
Default:	NOT_ALLOWED
Description:	This parameter controls the advertisement of a router other than the local BGP speaker as a nexthop of a BGP route being advertised to this peer. If configured as ALLOWED (and a BGP route is reachable through a router which is on the same subnet as the local and remote BGP speaker) then that router is advertised as the nexthop field of the route. Otherwise the address of the local BGP speaker is advertised as the nexthop of all the routes. If configured as NOT_ALLOWED, the address of the local BGP speaker is advertised as the nexthop of all the routes. This parameter is applicable to external peers only.

BGP Policies

The policies in this table allow the user to control the amount of BGP route information that is exchanged by local BGP speaker with its peers and to have a specific policy for the routes specified in this table other than the common policy specified for all the routes to or from a specific peer, in the BGP Peer Table.

Entry Number

Range:	1 - 2048 (Vanguard 7300 Series) 1-768 (All others)
Default:	1
Description:	This parameter uniquely identifies a BGP-OSPF import policy entry in this table.

IP Address

Range:	A valid IP Address
Default:	0.0.0.0
Help Message:	This parameter specifies the IP address of the subnet for which the import/export policy is specified by this entry. A value of 0.0.0.0 for this parameter indicates that this entry matches any subnet.

IP Address Mask

Range:	A valid IP Address Mask
Default:	0.0.0.0
Description:	This parameter specifies the IP address mask of the subnet for which the import/export policy is specified by this entry. This parameter along with IP Address defines a subnet on which the policy specified by this entry is applied. ■ Note 0.0.0.0 is not a valid value for this parameter.

Match Type

Range:	EXACT, RANGE
Default:	RANGE
Description:	This parameter specifies whether the subnet specified by this entry matches a single subnet or all constituent routes. When configured EXACT, the subnet specified by this entry corresponds to a single route. When configured as RANGE, the subnet specified by this entry corresponds to a range of routes within it. ■ Note This parameter is displayed only when the IP Address configured with a non-zero value.

Originating AS

Range:	0, 1 - 65535
Default:	0
Description:	This parameter specifies the AS which originated the BGP route specified by this entry. A value of 0 for this parameter indicates that this entry matches any Originating AS (that is the value of this parameter is ignored while matching this policy entry).

Advertising AS

Range:	0, 1 - 65535
Default:	0
Description:	This parameter specifies the AS which has advertised (Adjacent AS) the route specified by this entry, to the local BGP Speaker. A value of 0 for this parameter indicates that this entry matches any Advertising AS. The value of this parameter is ignored while matching this policy entry.

Any AS List

Range:	A comma separated list of AS numbers containing a maximum of 8 elements.
Default:	0
Description:	This parameter specifies a list of AS numbers any of which, if present in the AS_PATH attribute of the BGP Route, the policy specified by this entry is applicable. An AS number present in this list should not be same as the AS number configured in "Originating AS" or "Advertising AS", if they are configured to be non-zero values. For example 100, 500. A value of 0 for this parameter indicates that this entry matches any AS in the AS_PATH attribute of the route specified by this entry. The value of this parameter is ignored while matching this policy entry.

Direction

Range:	INBOUND, OUTBOUND
Default:	INBOUND
Description:	This parameter determines the direction in which the policy specified by this entry is an applicable policy (specified as an Inbound or Outbound policy). When configured INBOUND, the policy specified by this entry is applied to incoming BGP Routes from the specified BGP Peers. When configured OUTBOUND, the policy specified by this entry is applied to outgoing BGP Routes to the specified BGP Peers.

Peer Scope

Range:	ALL, Internal, External, Specific
Default:	ALL
Description:	<p>This parameter determines the peers to which the policy specified by this entry is applicable. A policy can be applied to All, Internal, External or a specific set of peers. When configured ALL, the policy specified by this entry is applicable to all peers. When configured External, the policy specified by this entry is applicable to external peers only. When configured Internal, the policy specified by this entry is applicable to internal peers only. When configured Specific, the policy specified by this entry is applicable to a specific set of Peers as specified in "Peer List".</p> <p>■ Note Policy configured in this table for a route, takes precedence over the corresponding Peer specific policy.</p>

Peer List

Range:	A comma separated list of Peer ID ranges containing maximum of 16 elements.
Default:	Blank
Description:	<p>This parameter specifies a list of peers to which the policy specified by this entry is applicable. The specific set of Peers, as identified by their Peer ID configured in "BGP Peer Table", can be configured as follows. Example: 1, 5 - 7, 10 This parameter is displayed only when the Peer Scope is set to SPECIFIC_PEERS.</p>

Filtering Action

Range:	PERMIT , DENY
Default:	DENY
Description:	<p>This parameter specifies the filtering action for the route(s) specified by this policy entry. When configured Permit, the route(s) specified by this entry are permitted from/to the Peers specified by this entry. When configured Deny, then the route specified by this entry is filtered from or to the Peers specified by this entry.</p>

Path Weight

Range:	0, 1 - 65535
Default:	0

Path Weight

Description:	<p>This parameter specifies the weight associated by local BGP Speaker to the route specified by this policy entry. This is used as one of the factors to calculate the degree of preference for the route(s) specified by this entry. A higher value of this parameter means that higher preference is given to the route(s) specified by this entry.</p> <p>■ Note For a matching entry, path weight is added to an existing path weight defined in peer parameters.</p>
--------------	---

MED

Range:	0, 1 - 0xFFFFFFFF
Default:	0
Description:	<p>This parameter specifies the value of Multi Exit Discriminator (MED) attribute with which routes specified by this entry are advertised to specified peers. The configured value is put in the MED Attribute of the update message containing the route specified by this entry. The routes with lower MED are preferred over the routes with higher MED.</p> <p>When set to 0, MED attribute for the route specified by this entry is set as per the peer specific policies configured in the "Peer Table".</p> <p>■ Note</p> <ol style="list-style-type: none"> 1. This parameter is applicable only for Outbound Permit policy entries, "having external peer(s) in its scope. 2. This parameter is not applicable, if the net parameters (IP Address, IP Address Mask) of the policy entry are configured with non-zero values.

Number of Extra AS Prepends

Range:	0, 1 - 10
Default:	0
Description:	<p>This parameter specifies the additional number of times the local AS number is prepended in the AS path attribute of the route(s) specified by this entry when advertised to the specified peers.</p> <p>A route with shorter AS_PATH length is preferred over a route with longer one.</p> <p>When set to 0, extra AS prepends for the route specified by this entry is done as per the peer specific policies configured in the "Peer Table". This parameter is applicable only for Outbound Permit policy entries, having external peer(s) in its scope. This parameter is not applicable, if the net parameters (IP Address, IP Address Mask) of the policy entry are configured with non-zero values.</p>

Local Preference

Range:	0, 1 - 0xFFFFFFFF
Default:	0
Description:	This parameter specifies the value of local preference attribute for the route(s) specified by this entry, when received from the specified external peers and being advertised to internal peers. Route with higher local preference is preferred over routes with lower local preference. When set to 0, local preference for the route specified by this entry is set as per the peer specific policies configured in the "Peer Table". This parameter is applicable only for Inbound Permit policy entries, having external peer(s) in its scope. This parameter is not applicable, if the net parameters (IP Address, IP Address Mask) of the policy entry are configured with non-zero values.

BGP Community Attribute Policy Record

Three new community attribute profile parameters are added to BGP policy record. The community profiles define the list of profile ranges, which define the applicable profiles. The well-known attributes are also included in the list. Community operation defines the operations of how to use the community profiles, MATCH, APPEND, DELETE. MATCH operation is meaningful for inbound policies and defines a new filter. APPEND and DELETE operations are meaningful for outbound policies. To replace an old attribute value, DELETE operation should be configured first, the APPEND could be added. The parameters are:

Match Community Profiles

Range:	A list of community profile ranges. NONE, Well-known community attribute, NO_EXPORT, NO_ADVERTISE, NO_EXPORT_SUBCONFED, NO_PEER
Default:	NONE
Description:	This parameter specifies the community profiles for an inbound policy. If any one of specified profiles is matched, the condition of community attribute is satisfied. A list of ranges plus the well-known community value can be input for this parameter. Default value is NONE. Example: 1-3, 32, NO_EXPORT

Delete Community Profiles

Range:	A list of community profile ranges. NONE, Well-known community attribute, NO_EXPORT, NO_ADVERTISE, NO_EXPORT_SUBCONFED, NO_PEER
Default:	NONE

Delete Community Profiles

Description:	This parameter specifies the community profiles for an outbound policy. If any one of specified profiles is matched, the condition of community attribute is satisfied and the matched community values are deleted from the attribute value set. A list of ranges plus the well-known community value can be input for this parameter. Default value is NONE. Example: 1-3, 32, NO_EXPORT
--------------	---

Append Community Profiles

Range:	A list of community profile ranges. NONE, Well-known community attribute, NO_EXPORT, NO_ADVERTISE, NO_EXPORT_SUBCONFED, NO_PEER
Default:	NONE
Description:	This parameter specifies the community profiles for an outbound policy. A list of ranges plus the well-known community value can be input for this parameter. Default value is NONE. Example: 1-3, 32, NO_EXPORT

BGP Aggregation

Profiles BGP route aggregation as defined in RFC 2519. BGP routes can be summarized into CIDR blocks to reduce the number of routes in the network. The aggregation profile specifies the IP address and mask to which routes will be summarized. It also references the list of IP routes that are included or excluded from the summarization. A BGP aggregate attribute indicated will provide the BGP attributes for the summarized route.

Entry Number

Range	1-64
Default	1
Description	Entry number used to reference this table record.
Boot Type	A change to this parameter requires a BGP Aggregation Profiles Boot to take effect.

Aggregate Address

Range	IP Address in dot notation
Default	0.0.0.0
Description	This parameter specifies the IP address of an aggregate route which will be created based on the component routes in the BGP Routing Table.
Boot Type	A change to this parameter requires a BGP Aggregation Profiles Boot to take effect.

Aggregate Address Mask

Range	IP Address in dot notation
Default	255.0.0.0
Description	This parameter specifies the IP address mask of the aggregate route which will be created based on the component routes in the BGP routing table.
Boot Type	A change to this parameter requires a BGP Aggregation Profiles Boot to take effect.

Include Network Profiles

Range	A list of entry number ranges of IP Map Profiles records, or None
Default	None
Description	This parameter specifies the entry numbers of IP Map Profiles record. The IP address ranges in the designated IP Map Profile records are used to specify which route is eligible to be a component route. The aggregate route will be created solely based on the component routes which are falling into the IP address ranges in the corresponding IP Map Profiles records. If None is input, all prefixes are eligible. Example: 1,2-5,6
Boot Type	A change to this parameter requires a BGP Aggregation Profiles Boot to take effect.

Exclude Network Profiles

Range	A list of entry number ranges of IP Map Profiles records, or None
Default	None
Description	This parameter specifies the entry numbers of IP Map Profiles record. The IP address ranges in the designated IP Map Profile records are used to specify which route is not eligible to be a component route. The aggregate route will be created solely based on the component routes which are not falling into the IP address ranges in the corresponding IP Map Profile records. If None is input, no prefixes will be excluded. Example: 1,2-5,6
Boot Type	A change to this parameter requires a BGP Aggregation Profiles Boot to take effect.

Peer List

Range	A list of ranges of BGP peer numbers, or All
Default	ALL
Description	This parameter specifies a list of BGP peers to which specific routes can be advertised along with an aggregate route. To be successfully advertised to a BGP peer, a specific route must match a BGP outbound policy and the action type of the policy must be PERMIT. If All is input, then all BGP peers are included. Example: 1,2,3-5
Boot Type	A change to this parameter requires a BGP Aggregation Profiles Boot to take effect.

Specific Map Profiles

Range	A list of entry number ranges of IP Map Profiles records, or None
Default	None
Description	This parameter specifies the entry numbers of IP Map Profile records. The IP address ranges in the designated IP Map Profile records are used to specify which route is eligible to be advertised with the aggregate route. If None is input, no component prefixes will be advertised. Example: 1,2,3-5
Boot Type	A change to this parameter requires a BGP Aggregation Profiles Boot to take effect.

Append Community Profiles

Range	NO_EXPORT, NO_ADVERTISE, NO_EXPORT_SUBCONFED, NO_PEER, NONE
Default	None
Description	This parameter specifies the entry numbers of IP Map Profile records. The IP address ranges in the designated IP Map Profile records are used to specify which route is not eligible to be advertised with the aggregate route. If None is input, no component prefixes will be advertised depending on Specific Map Profiles applied. Example: 1,2,3-5
Boot Type	A change to this parameter requires a BGP Aggregation Profiles Boot to take effect.

Suppress Map Profiles

Range	A list of entry number ranges of IP Map Profiles records, or None
Default	None
Description	This parameter specifies the entry numbers of IP Map Profile records. The IP address ranges in the designated IP Map Profile records are used to specify which route is not eligible to be advertised with the aggregate route. If None is input, no component prefixes will be advertised depending on Specific Map Profiles applied. Example: 1,2,3-5
Boot Type	A change to this parameter requires a BGP Aggregation Profiles Boot to take effect.

Path Setting

Range	AUTOMATIC,ORIGINATOR
Default	AUTOMATIC
Description	This parameter specifies the attribute setting method in the aggregate route. Automatic: The attribute in the aggregate route will be calculated based on the component routes automatically. No attribute in the component routes will be lost. Originator: The attribute in the aggregate route will be populated as a self-originated route.
Boot Type	A change to this parameter requires a BGP Aggregation Profiles Boot to take effect.

Attribute Profile

Range	1-32
Default	None
Description	This parameter specifies the entry number of BGP Aggregate Attribute Profile record. The content of the designated record will be applied to the aggregate route. If None is input, there will be no Attribute Profile.
Boot Type	A change to this parameter requires a BGP Aggregation Profiles Boot to take effect.

IP Map Profiles

The IP Map profiles are a list of routes and masks that are used in the BGP Aggregation Profile.

Entry Number

Range	1-256
Default	1
Description	Entry number used to reference this table record.
Boot Type	A change to this parameter requires a BGP IP MAP Profiles Boot to take effect.

IP Networks

Range	0.0.0.0/8
Default	0.0.0.0/8
Description	This parameter specifies a list of networks ranges. Each element of the list specifies one network range in IP Address/Mask Length format (e.g. X.X.X.X/yy). As many as 16 network elements can be specified in this parameter. Example: 150.6.1.0/24-150.6.16.0/24, 150.6.24.0/24
Boot Type	A change to this parameter requires a BGP IP MAP Profiles Boot to take effect.

BGP Aggregate Attribute Profiles

The BGP Aggregate Attribute Profile is the list of BGP attributes that are assigned to the Aggregate Route specified in the BGP Aggregation Profile.

Entry Number

Range	1-32
Default	1
Description	Entry number used to reference this table record.
Boot Type	Boot BGP Aggregate Attribute Profiles

MED

Range	0x00000000-0xFFFFFFFF
Default	Automatic
Description	This parameter specifies the MED attribute of the corresponding aggregate route. If Automatic is input, the internally calculated MED value will be used.
Boot Type	Boot BGP Aggregate Attribute Profiles

Local Preference

Range	0x00000000-0xFFFFFFFF
Default	Automatic
Description	This parameter specifies the Local Preference attribute of the corresponding aggregate route. If Automatic in input, the internally calculated Local_Pref will be used.
Boot Type	Boot BGP Aggregate Attribute Profiles

Origin

Range	IGP, EGP, Incomplete, Automatic
Default	Automatic
Description	This parameter specifies the Origin attribute of the corresponding aggregate route. If Automatic in input, the internally calculated Origin will be used.
Boot Type	Boot BGP Aggregate Attribute Profiles

Next Hop

Range	IP Address in dot notation
Default	0.0.0.0
Description	This parameter specifies the Next Hop attribute of the corresponding aggregate route. A value of 0.0.0.0 (default) will use the calculated next hop address.
Boot Type	Boot BGP Aggregate Attribute Profiles

AS Path Prepend

Range	1-65535
Default	None
Description	This parameter specifies the AS number(s) which will be prepended to the calculated AS_Path attribute in the corresponding aggregate route. If None is input, no AS number will be prepended in the AS_Path attribute. Example: 200,300,150
Boot Type	Boot BGP Aggregate Attribute Profiles

BGP Community Attribute Profiles

BGP Community implementation as described in RFC 1997. Communities provide a mechanism to handle a variety of routes with a single policy. This is used in conjunction with BGP Policies.

Entry Number

Range	1-128
Default	1
Description	Entry number used to reference this table record.
Boot Type	A change to this parameter requires a BGP Community Attribute Profiles Boot to take effect.

AS Number

Range	1-65534, *, MY_AS
Default	MY_AS
Description	This parameter specifies the AS number of a community attribute value. If default value, MY_AS, is configured, the AS number configured in the BGP Global Parameters will be used. The wildcard, *, matching any AS number with the same community value, is only applicable for DELETE community profiles in Policy configuration.
Boot Type	A change to this parameter requires a BGP Community Attribute Profiles Boot to take effect.

Community Value

Range	1-65535, *
Default	1
Description	This parameter specifies the value for the community attribute. The previous parameter, AS Number and this parameter, are combined together to define one element in the set of community value. The wildcard, *, matching any community value with the same AS number, is only applicable for DELETE community profiles in Policy configuration.
Boot Type	A change to this parameter requires a BGP Community Attribute Profiles Boot to take effect.

BGP Network Route Table

BGP Network Route Table provides a way for BGP to initiate the advertisement of a route without redistributing it from an IGP such as RIP or OSPF. This feature can be loosely described as a BGP static route. The Network Route Table provides the capability, for example, of advertising a route that is reachable by the router only by its default gateway. The IGP synchronization flag is used to determine if a route must be in the IGP routing table before BGP can advertise it. This is new to release 7.1.

Entry Number

Range	1-256
Default	1
Description	Entry number used to reference this table record.
Boot Type	Boot BGP Network Route Table

IP Network/Subnet

Range	IP Address in dot notation
Default	0.0.0.0
Description	The IP address of a destination network or subnet.
Boot Type	Boot BGP Network Route Table

IP Address Mask

Range	255.0.0.0 to 255.255.255.255 A valid IP Address mask in dotted notation in the format X.X.X.X where X can take a maximum value of 255.
Default	255.255.255.0
Description	The mask associated with the IP Network/Subnet address. For example: if the destination is a subnet of a class B network and the third byte of the IP address is used as the subnet portion, the address mask would be set to 255.255.255.0.
Boot Type	Boot BGP Network Route Table

IGP Synchronization

Range	Enabled, Disabled
Default	Disabled
Description	If set to enabled, the route is imported into BGP only if presenting the IGP table. If set to disabled, the route is unconditionally imported into BGP.
Boot Type	Boot BGP Network Route Table

RIP -> BGP Import Policies

RIP->BGP Route filtering is necessary to control the RIP routes being introduced into BGP. When redistribution into BGP is necessary, it is not usually best practice to advertise all RIP routes into BGP. In particular, cases where BGP is redistributing routes into the IGP table, those routes must not be redistributed back into BGP. This would cause routing loops. The resultant action of the policy is order dependent. This is new to release 7.1.

With the release of 7.2R00A, a new parameter, "Peer weight" has been added to allow more control and flexibility over the routing decisions made by BGP regarding RIP imported routes.

Entry Number

Range	0-256
Default	1
Description	Entry number used to reference this table record.
Boot Type	Boot BGP RIP to BGP Import Policies

IP Address

Range	IP Address in dot notation
Default	0.0.0.0
Description	This parameter specifies the IP address of the subnet for which the import policy is specified by this entry. A value of 0.0.0.0 for this parameter indicates that this entry matches any subnet. A value of 0.0.0.0 for this parameter and 255.255.255.255 for the IP Address Mask parameter indicates that this policy entry matches the default route.
Boot Type	Boot BGP RIP to BGP Import Policies

IP Address Mask

Range	255.0.0.0 to 255.255.255.255 A valid IP Address mask in dotted notation in the format X.X.X.X where X can take a maximum value of 255.
Default	0.0.0.0

IP Address Mask *(continued)*

Description	<p>This parameter specifies the IP address mask of the subnet for which the import policy is specified by this entry.</p> <p>This parameter along with the IP address defines a network/subnet on which the policy specified by this entry will be applied.</p> <p>A value of 0.0.0.0 for both this parameter and the IP Address parameter indicates that this entry matches any subnet.</p> <p>A value of 255.255.255.255 for this parameter and 0.0.0.0 for the IP Address parameter indicates that this policy entry matches the default route.</p>
Boot Type	Boot BGP RIP to BGP Import Policies

Match Type

Range	Exact, Range
Default	Range
Description	<p>This parameter specifies whether the subnet specified by this entry matches a single subnet or all its constituent routes.</p> <p>When configured to Exact, the subnet specified by this entry corresponds to a single route.</p> <p>When configured to Range, the subnet specified by this entry corresponds to a range of routes within it.</p> <p>Note This parameter is displayed only when the IP Address is configured with a non-zero value.</p>
Boot Type	Boot BGP RIP to BGP Import Policies

Filtering Action

Range	Permit, Deny
Default	Deny
Description	<p>This parameter specifies the filtering action for the route(s) specified by this entry.</p> <p>When configured to Permit, the route(s) specified by this entry will be imported into BGP routing domain.</p> <p>When configured to Deny, the route(s) specified by this entry will not be imported into BGP routing domain.</p>
Boot Type	Boot BGP RIP to BGP Import Policies

Peer Weight

Range	0-65535
Default	32768

Peer Weight *(continued)*

Description	This parameter specifies the weight associated with this peer. This is used as one of the factors to calculate the degree of preference for all the routes learned from this peer. A higher value of this parameter means that higher preference is given to the routes learned from this peer. A value of 0 indicates that no special preference is given to the routes learned from this peer.
-------------	--

OSPF -> BGP Import Policies

OSPF->BGP Route filtering is necessary to control the OSPF routes being introduced into BGP. When redistribution into BGP is necessary, it is not usually best practice to advertise all OSPF routes into BGP. In particular, cases where BGP is redistributing routes into the IGP table, those routes must not be redistributed back into BGP. This would cause routing loops. The resultant action of the policy is order dependent. This new to release 7.1.

With the release of 7.2R00A, a new parameter, “Peer weight” has been added to allow more control and flexibility over the routing decisions made by BGP regarding OSPF imported routes.s

Entry Number

Range	0-256
Default	1
Description	Entry number used to reference this table record.
Boot Type	Boot BGP OSPF to BGP Import Policies

IP Address

Range	IP Address in dot notation
Default	0.0.0.0
Description	This parameter specifies the IP address of the subnet for which the import policy is specified by this entry. A value of 0.0.0.0 for this parameter indicates that this entry matches any subnet. A value of 0.0.0.0 for this parameter and 255.255.255.255 for the IP Address Mask parameter indicates that this policy entry matches the default route.
Boot Type	Boot BGP OSPF to BGP Import Policies

IP Address Mask

Range	255.0.0.0 to 255.255.255.255 A valid IP Address mask in dotted notation in the format X.X.X.X where X can take a maximum value of 255.
Default	0.0.0.0

IP Address Mask *(continued)*

Description	<p>This parameter specifies the IP address mask of the subnet for which the import policy is specified by this entry.</p> <p>This parameter along with the IP address defines a network/subnet on which the policy specified by this entry will be applied.</p> <p>A value of 0.0.0.0 for both this parameter and the IP Address parameter indicates that this entry matches any subnet.</p> <p>A value of 255.255.255.255 for this parameter and 0.0.0.0 for the IP Address parameter indicates that this policy entry matches the default route.</p>
Boot Type	Boot BGP OSPF to BGP Import Policies

Match Type

Range	Exact, Range
Default	Range
Description	<p>This parameter specifies whether the subnet specified by this entry matches a single subnet or all its constituent routes.</p> <p>When configured to Exact, the subnet specified by this entry corresponds to a single route.</p> <p>When configured to Range, the subnet specified by this entry corresponds to a range of routes within it.</p> <p>Note This parameter is displayed only when the IP Address is configured with a non-zero value.</p>
Boot Type	Boot BGP OSPF to BGP Import Policies

Filtering Action

Range	Permit, Deny
Default	Deny
Description	<p>This parameter specifies the filtering action for the route(s) specified by this entry.</p> <p>When configured to Permit, the route(s) specified by this entry will be imported into BGP routing domain.</p> <p>When configured to Deny, the route(s) specified by this entry will not be imported into BGP routing domain.</p>

Peer Weight

Range	0-65535
Default	32768

Peer Weight (continued)

Description	This parameter specifies the weight associated with this peer. This is used as one of the factors to calculate the degree of preference for all the routes learned from this peer. A higher value of this parameter means that higher preference is given to the routes learned from this peer. A value of 0 indicates that no special preference is given to the routes learned from this peer.
-------------	--

Configuring OSPF Routing Parameters

AS Boundary Routing Parameters

The following tables gives the description of the newly added parameters under the "AS Boundary Routing Parameters" in "Configure OSPF" menu, for providing global control on the import of BGP routes into OSPF domain. The configuration path is: **Configure->Router->Configure OSPF->AS Boundary Routing Parameters**

Import BGP Routes

Range:	Yes, No
Default:	No
Description:	This parameter controls the importing of BGP routes into the OSPF domain. When configured Yes, BGP routes are imported into OSPF domain as per the configured policy. When configured No, BGP routes are not imported into OSPF domain even if allowed by the configured policies.

Default BGP->OSPF Import Policy

Range:	PERMIT, DENY
Default:	DENY
Description:	<p>This parameter specifies the default import policy for importing BGP routes into OSPF routing domain. When a BGP route does not match with any of the policy entries configured in "BGP-OSPF Import Policy Table" then this parameter defines the default policy.</p> <p>When configured "Permit", the route is imported with the following default values for various policy parameters: Metric = 0 Metric Type = Type2 Non-zero Forwarding Address = NOT_ALLOWED OSPF Tag Generation = Disable</p> <p>When configured Deny, BGP routes are not imported into OSPF as a default import policy.</p> <p>This parameter appears only when "Import BGP Routes" parameter is set to Yes.</p>

**BGP to OSPF
Import Policy Table**

This table specifies the policy for importing BGP routes to the OSPF routing domain. BGP routes are imported into OSPF domain as OSPF External routes. The configuration path for the BGP to OSPF Import Policy Table is: **Configure->Router->Configure OSPF->BGP->OSPF Import Policy Table**

Entry Number

Range:	1 - 1024
Default:	1
Description:	This parameter uniquely identifies a BGP-OSPF import policy entry in this table.

IP Address

Range:	A valid IP Address
Default:	0.0.0.0
Description:	This parameter specifies the IP address of the subnet for which the import policy is specified by this entry. A value of 0.0.0.0 for this parameter indicates that this entry matches any subnet.

IP Address Mask

Range:	A valid IP Address Mask
Default:	0.0.0.0
Description:	This parameter specifies the IP address mask of the subnet for which the import policy is specified by this entry. This parameter along with IP Address defines a subnet on which the policy specified by this entry is applied. ■ Note 0.0.0.0 is not a valid value for this parameter.

Match Type

Range:	EXACT, RANGE
Default:	RANGE
Description:	This parameter specifies whether the subnet specified by this entry matches a single subnet or all its constituent routes. When configured EXACT the subnet specified by this entry corresponds to a single route. When configured RANGE the subnet specified by this entry corresponds to a range of routes within it. ■ Note This parameter is displayed only when the IP Address configured with a non-zero value.

Originating AS

Range:	0, 1 - 65535
Default:	0
Description:	This parameter specifies the AS which originated the BGP route specified by this entry. A value of 0 for this parameter indicates that this entry matches any originating AS value of this parameter is ignored while matching this policy entry.

Advertising AS

Range:	0, 1 - 65535
Default:	0
Description:	This parameter specifies the AS which has advertised (Adjacent AS) the route specified by this entry, to the local BGP Speaker. A value of 0 for this parameter indicates that this entry matches any Advertising AS that is the value of this parameter is ignored while matching this policy entry.

Filtering Operation

Range:	PERMIT, DENY
Default:	DENY
Description:	This parameter specifies the filtering action for the route(s) specified by this entry. When configured PERMIT, the route(s) specified by this entry is imported into OSPF routing domain. When configured Deny, the route(s) specified by this entry are not imported into OSPF routing domain.

■ Note

The following parameters listed in this section are displayed only if the “Filtering Operation’s range is set to “PERMIT”.

OSPF Metric

Range:	0, 1 - 0xFFFFFFFF
Default:	0

OSPF Metric (continued)

Description:	<p>This is the OSPF ToS-0 metric with which the BGP route specified by this entry is imported into the OSPF routing domain.</p> <p>A value of 0 for this parameter indicates that the locally calculated DoP is used to derive a OSPF metric automatically, with which BGP route is imported. The OSPF metric of the route is derived from its DoP as follows: OSPF Metric = First 24 bits of (0xFFFFFFFF - DoP - 5).</p> <p>A non-zero value for this parameter indicates the manual translation of BGP DoP to OSPF metric. BGP route is imported into OSPF with the configured metric.</p> <p>■ Note This parameter is displayed only when the Filtering Operation is set to PERMIT.</p>
--------------	---

Metric Type

Range:	TYPE1-TYPE2
Default:	TYPE2
Description:	<p>This parameter specifies the OSPF external metric type with which the BGP route specified by this entry is imported into OSPF.</p> <p>A value of Type2 means the BGP and OSPF metrics are not comparable. A value of Type1 means the BGP and OSPF metrics are comparable.</p> <p>■ Note This parameter is displayed only when the Filtering Operation is set to PERMIT.</p>

Non-zero Forwarding Address

Range:	ALLOWED NOT_ALLOWED
Default:	NOT_ALLOWED

Non-zero Forwarding Address (continued)

Description:	<p>This parameter specifies whether the route specified by this entry can be imported into OSPF with a non-zero Forwarding Address, in cases where the BGP Speaker, the Peer from which the BGP route was learned, and the OSPF router to which the BGP route is being advertised, are on the same subnet.</p> <p>When set to ALLOWED, BGP Route specified by this entry is imported with the forwarding address set to the IP Address present in BGP NextHop path attribute.</p> <p>When set to NOT_ALLOWED, the BGP Route specified by this entry is always imported with the forwarding address set to 0.0.0.0.</p> <p>Note This parameter is displayed only when the Filtering Operation is set to PERMIT.</p>
--------------	---

#OSPF Route Tag Generation

Range:	DISABLE, AUTOMATIC, MANUAL
Default:	DISABLE
Description:	<p>This parameter controls the generation of OSPF external route tag with which BGP route specified by this entry is imported into OSPF domain. If allowed, this parameter also specifies the method of tag generation. When configured DISABLE, the BGP Route is imported into the OSPF routing domain with External Route Tag, set to 0. When configured MANUAL, the BGP Route is imported into OSPF with the tag as configured in "OSPF Route Tag" parameter. When configured AUTOMATIC, the BGP route is imported into OSPF with a tag which is automatically generated as per the rules specified by RFC1745.</p> <p>Note This parameter is displayed only when the Filtering Operation is set to PERMIT.</p>

#OSPF Route Tag

Range:	0, 1 - 0x7FFFFFFF
Default:	0
Description:	<p>The parameter specifies the value of OSPF External Route Tag with which the BGP route specified by this entry is imported into the OSPF domain. The tag specified by this parameter is put in the "External Route Tag" field of the corresponding OSPF Route.</p> <p>A value of 0 indicates that the BGP route is imported into OSPF with External Route Tag set to the Originating/Adjacent AS present in BGP AS_PATH Attribute.</p> <p>This parameter is displayed only when the Filtering Operation is set to PERMIT and the OSPF Route Tag Generation parameter is set to MANUAL.</p>

Configuring RIPv2 Route Redistribution Parameters

RIPv2 Redistribution Parameters

The BGP to RIPv2 redistribution is controlled by the users configuration. BGP to RIPv2 redistribution provides menu items in the Vanguard router so that the user can switch the redistribution, configure the redistribution policies, check the redistribution statistics, and receive warning and alarms.

The Enable/Disable parameter and the default policy are included in the IP parameters because they are applicable to the whole BGP to RIPv2 redistribution process. The non-default policy is included as a separate menu item because every single policy is only locally meaningful. The configuration path is:

Configure->Router->Configure IP->IP Parameters

BGP to RIP Enable

Range:	Enable, Disable
Default:	Disable
Description:	This parameter enables or disables route redistribution from BGP to RIPv2.
Boot Effect:	The ASBR imports the BGP routes into RIPv2 according to BGP to RIPv2 redistribution policies and default policy. The imported routes take part in the RIPv2 aggregation and are advertised to other nodes in the local AS. Enable->Disable: The ASBR flushes the imported BGP routes and derived aggregate routes from the global routing table and advertise the unreachable routes via triggered update in the local AS.

BGP to RIP Default Filter

Range:	PERMIT, DENY
Default:	DENY
Description:	This parameter specifies the default import policy for importing BGP routes into RIP2 routing domain. When a BGP route does not match with any of the policy entries configured in "BGP to RIP default filter" table then this parameter defines the default policy. When configured "PERMIT", route is imported with the following default values for policy parameters: RIP2 Tag Generation = Disable When configured "DENY", BGP routes are not imported into RIP2 domain as a default import policy.

BGP to RIP Default Filter *(continued)*

Boot Effect:	<p>Permit->Deny: Any imported BGP route in the global routing table, which uses the default policy, needs to be flushed. The flushed routes are advertised in the local AS.</p> <p>Deny->Permit: All the BGP routes which are denied by the default policy is imported into RIPv2 and advertised in the local AS. The result is that all the BGP routes, which are not originated by local AS, are imported into the RIPv2 domain if no other control method, such as RIP control, is used.</p>
--------------	---

BGP to RIP Nondefault Route Override

Range:	ENABLE, DISABLE
Default:	DISABLE
Description:	This parameter specifies whether the non-default RIP routes are overwritten by imported BGP routes or not. If set to ENABLE the BGP imported routes will overwrite an existing RIP route. If set to DISABLE, any existing RIP routes are not overwritten.

BGP to RIP Default Route Override

Range:	ENABLE, DISABLE
Default:	DISABLE
Description:	This parameter specifies whether a default RIP route is overwritten by imported BGP routes or not. If set to ENABLE the BGP imported route will overwrite an existing RIP Route. If set to DISABLE, any existing RIP routes are not overwritten.

BGP to RIP Default Metric

Range:	1 to 15
Default:	1
Description:	This parameter specifies the RIPv2 metric for any route which matches the default import policy.

BGP to RIP Default Metric (continued)

Boot Effect:	<p>Increase Value: All imported BGP routes need to set the metric to this new value. Increasing metrics make the imported BGP route less preferable and could cause the route to become the secondary if multiple routes share the same destination. In turn, the Access Control/PBR could be effected if configured to use a specific next hop.</p> <p>Decrease Value: All imported BGP routes need to set the metric to this new value. Decreasing metrics make the imported BGP route more preferable and could cause the route to become the primary if multiple routes share the same destination. In turn, the Access Control/PBR could be effected if configured to use a specific next hop.</p>
--------------	---

**BGP to RIPv2
Redistribution
Policy**

These new CMEM records define the non-default policies used in the BGP to RIPv2 route redistribution process. The non-default policy always takes priority over default policy. The configuration path is:

Configure->Router->Configure IP->BGP->RIP Import Policies

<i>Parameters</i>	<i>Description</i>
Entry Number	This parameter uniquely identifies a BGP to RIP Default Filter entry in this table. It also puts an upper bound to the number of entries that can be configured in this table.
IP Address	This parameter specifies the IP address of the subnet for which the import policy is specified by this entry.
IP Address Mask	This parameter specifies the IP address mask of the subnet for which the import policy is specified by this entry. This parameter along with the IP Address defines a subnet on which the policy specified by this entry is applied.
Match Type	This parameter specifies whether the subnet specified by this entry matches a single subnet or all of its constituent routes.
Originating AS	This parameter specifies the AS which originated the BGP route specified by this entry.
Advertising AS	This parameter specifies the AS which has advertised (Adjacent AS) the route specified by this entry, to the local BGP Speaker.
Filtering Action	<p>This parameter specifies the filtering action for the route(s) specified by this entry.</p> <p>■ Note The rest of the policy parameters in this table is applicable only when the "Filter Action" is set to "Permit".</p>

Configuring RIPv2 Route Redistribution Parameters

Parameters	Description
RIP Metric	This is the RIPv2 metric with which the BGP route specified by this entry is imported into the RIPv2 routing domain. A valid value is configured for the manual translation of BGP Dop to RIPv2 metric.
RIP Route Tag	This parameter controls the generation of RIPv2 route tag with which BGP route, specified by this entry, is imported into RIPv2 domain. If allowed, this parameter also specifies the method of tag generation. This parameter is set to Disable when the generation and processing of Route Tag is not supported by local and/or remote RIPv2 ASBRs. This parameter is set to Manual when it is required to communicate some useful information about the route being imported, among the ASBRs. The user determines the meaning of this parameter. This is set to Automatic when it is required to communicate the characteristics of the BGP route being imported, among the ASBRs. This is used to set ORIGIN and AS_PATH attributes of the BGP route intelligently when the corresponding RIPv2 route is exported into BGP domain by remote ASBR.
Tag Value	The parameter specifies the value of RIPv2 Route Tag with which the BGP route, specified by this entry, is imported into the RIPv2 domain. The tag specified by this parameter is put in the "Route Tag" field of the corresponding RIPv2 Route. This parameter is set to a non-zero value when it is required to communicate some useful information about the route being imported, among the ASBRs. This parameter is set to zero when it is required to tag Originating/Advertising AS of the BGP Route along with the corresponding RIPv2 route. This is used to set AS_PATH attributes of the BGP route intelligently, when the corresponding RIPv2 route is exported into BGP domain by remote ASBR, in the cases when automatic tag generation is not supported by local and/or remote RIPv2 ASBRs. Note: This parameter is applicable only when "RIP Route Tag" is set to Manual.
RIP Aggregative	This parameter specifies whether the imported routes can be combined with other routes to form aggregate route.
Import Specific Route	This parameter specifies whether the specific route can be imported when a corresponding aggregate route exists.

Entry Number

Range:	1 to 1,024
Default:	1
Description:	This parameter uniquely identifies a BGP to RIP Default Filter entry in this table.

IP Address

Range:	A valid IP Address
Default:	0.0.0.0
Description:	This parameter specifies the IP address of the subnet for which the import policy is specified by this entry. A value of 0.0.0.0 for this parameter indicates that this entry matches any subnet. ■ Note If configured 0.0.0.0 then the IP Address Mask parameter are not displayed.
Boot Effect:	Refer to the Route Selection Boot Effect section below, for the boot effect of this parameter.

Route Selector Boot Effect

A change in any of the route selectors in this table have the effect of deleting the existing policy entry and adding a new policy entry.

Deleting a Policy entry has the following effect:

- a) If the "Filtering Action" for the existing policy is Deny then routes specified by the entry is imported into RIPv2, if the "BGP to RIP Default Filter" parameter is set to Permit.
- b) If the "Filtering Action" for the existing policy is Permit, then BGP routes specified by the entry is flushed from RIPv2 domain by advertising them with the hop count as INFINITY, if the "BGP to RIP Default Filter" parameter is set to Deny.

Adding a Policy entry has the following effect:

- a) If the "Filtering Action" for the existing policy is Permit, then routes specified by the entry is imported into RIPv2, if the "BGP to RIP Default Filter" parameter is set to Deny.
- b) If the "Filtering Action" for the existing policy is Deny, then BGP routes specified by the entry is flushed from the RIPv2 domain by advertising them with MAXAGE, if the "BGP to RIP Default Filter" parameter is set to Permit.

IP Address Mask

Range:	A valid IP address mask
Default:	0.0.0.0
Description:	This parameter specifies the IP address mask of the subnet for which the import policy is specified by this entry. This parameter along with IP Address defines a subnet on which the policy specified by this entry is applied. ■ Note 0.0.0.0 is not a valid value for this parameter.
Boot Effect:	Refer to "Route Selector Boot Effect" section on page 3-43.

Match Type

Range:	EXACT, RANGE
Default:	RANGE
Description:	<p>This parameter specifies whether the subnet specified by this entry matches a single subnet or all of its constituent routes:</p> <p>When configured EXACT the subnet specified by this entry corresponds to a single route.</p> <p>When configured RANGE the subnet specified by this entry corresponds to a range of routes within it.</p> <p>■ Note This parameter is displayed only when the IP Address configured with a non-zero value.</p>
Boot Effect:	Refer to “Route Selector Boot Effect” section on page 3-43.

Originating AS

Range:	0, 1 to 65,535
Default:	0
Description:	<p>This parameter specifies the AS which originated the BGP route specified by this entry. A value of 0 for this parameter indicates that this entry matches any Originating AS i.e. value of this parameter is ignored while matching this policy entry.</p>
Boot Effect:	Refer to “Route Selector Boot Effect” section on page 3-43.

Advertising AS

Range:	0, 1 to 65,535
Default:	0
Description:	<p>This parameter specifies the AS which has advertised (Adjacent AS) the route specified by this entry, to the local BGP Speaker. A value of 0 for this parameter indicates that this entry matches any Advertising AS i.e. value of this parameter is ignored while matching this policy entry.</p>
Boot Effect:	Refer to “Route Selector Boot Effect” section on page 3-43.

Filtering Action

Range:	PERMIT, DENY
Default:	DENY
Description:	This parameter specifies the filtering action for the route(s) specified by this entry: When configured to Permit, the route(s) specified by this entry is imported into RIPv2 routing domain. When configured to Deny, the route(s) specified by this entry are not imported into RIPv2 routing domain.
Boot Effect:	Permit->Deny Previously imported BGP route(s) is flushed from the RIPv2 routing table and advertised with INFINITY METRIC to other nodes in the local AS. Deny->Permit The route(s) specified by this entry, if present in the BGP routing table, is imported into the RIPv2 routing table and advertised to other nodes in the local AS.

Filtering Action Parameters

The following parameters are displayed only if the “Filtering Action” range is set to “PERMIT”:

Import Specific Route

Range:	Enable, Disable
Default:	Disable
Description:	This parameter specifies whether the specific route could be imported when a corresponding aggregate route exists.
Boot Effect:	Refer to “Route Selector Boot Effect” section on page 3-43.

RIP Aggregative

Range:	Enable, Disable
Default:	Enable
Description:	This parameter specifies whether the imported routes could be combined with other routes to form aggregate route.
Boot Effect:	Enable->Disable: Previously formed aggregate routes based on imported BGP route(s) is flushed from the RIPv2 routing table and advertised with INFINITY METRIC to other nodes in the local AS. Disable->Enable: The route(s) specified by this entry, is taken into RIPv2 aggregation process, the aggregate route, if created, is put in the RIPv2 routing table and advertised to other nodes in the local AS.

RIP Metric

Range:	1 to 15
Default:	1
Description:	This is the RIPv2 metric with which the BGP route specified by this entry is imported into the RIPv2 routing domain. A valid value is configured for the manual translation of BGP Dop to RIPv2 metric.
Boot Effect:	A change in this parameter results in the BGP routes specified by this entry being re-imported into RIPv2 domain as per the newly configured value of this parameter.

RIP Route Tag

Range:	Disable, Manual, AdvertisingAS_Automatic, OriginatingAS_Automatic
Default:	Disable
Description:	This parameter controls the generation of RIPv2 route tag with which BGP route specified by this entry is imported into RIPv2 domain. If allowed, this parameter also specifies the method of tag generation. This parameter is set to Disable when the generation and processing of Route Tag is not supported by local and/or remote RIPv2 ASBRs. This parameter is set to Manual when it is required to communicate some useful information about the route being imported, among the ASBRs. The user determines the meaning of this parameter. AdvertisingAS_Automatic or OriginatingAS_Automatic is used when it is required to communicate the characteristics of the BGP route being imported, among the ASBRs. This is used to set ORIGIN and AS_PATH attributes of the BGP route intelligently when the corresponding RIPv2 route is exported into BGP domain by remote ASBR.
Boot Effect:	A change in this parameter results in the BGP routes specified by this entry getting re-imported into RIPv2 domain as per the newly configured value of this parameter.

Tag Value

Range:	0, 1 to 65,535
Default:	0
Description:	<p>The parameter specifies the value of RIPv2 Route Tag with which the BGP route specified by this entry is imported into the RIPv2 domain. The tag specified by this parameter is put in the " Route Tag" field of the corresponding RIPv2 Route. This parameter is set to a non-zero value when it is required to communicate some useful information about the route being imported, among the ASBRs. This parameter is set to zero when it is required to tag Originating/Advertising AS of the BGP Route along with the corresponding RIPv2 route. This is used to set AS_PATH attributes of the BGP route intelligently, when the corresponding RIPv2 route is exported into BGP domain by remote ASBR, in the cases when automatic tag generation is not supported by local and/or remote RIPv2 ASBRs.</p> <p>■ Note This parameter is applicable only when "RIP Route Tag" is set to Manual.</p>
Boot Effect:	A change in this parameter results in the BGP routes specified by this entry getting re-imported into RIP domain as per the newly configured value of this parameter.

Configuring BGP Aggregation Parameters

BGP Aggregation

The BGP aggregation is controlled by user's configuration. The BGP aggregation provides menu items in Vanguard routers so that users can switch the aggregation, configure the aggregation rules and multi-homed sites, check the aggregation statistics, and receive warning and alarms.

The BGP Aggregation and BGP Proxy Aggregation selections are found under Global Parameters. These parameters are used to control the whole BGP aggregation process. The BGP Aggregate Profiles, IP Map Profiles and BGP Aggregate Attribute Profiles are shown in Figure 3-3.

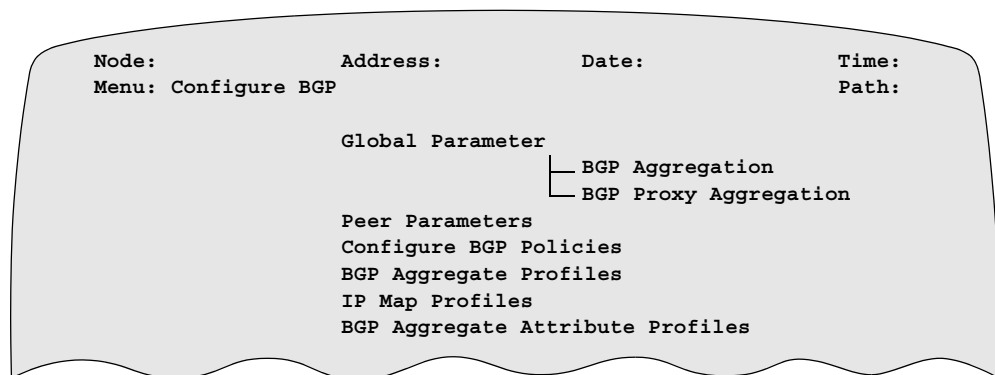


Figure 3-3. Configure BGP Menu

BGP Aggregation Parameters in BGP Global Control Parameter Menu

Figure 3-3 shows the two new parameters added to the BGP Global control parameter menu. They are arranged at the end of the BGP global control table.

BGP Aggregation

Range:	Enable, Disable
Default:	Disable
Description:	<p>This parameter specifies whether BGP aggregation can run in this node or not.</p> <p>Note If no BGP aggregation profiles are configured, this parameter is set to Disable no matter what value configured here.</p>
Boot Effect:	<p>Disable->Enable: BGP Aggregation begins to run in this node. Aggregate routes are created and added in the BGP routing table. These routes can be advertised to its peers.</p> <p>Enable->Disable: BGP Aggregation stops running in this node. Aggregate routes are flushed in the BGP routing table and sent to its peers as withdrawn routes. Specific routes may be sent out as feasible NLRI.</p>

BGP Proxy Aggregation

Range:	Enable, Disable
Default:	Disable
Description:	<p>This parameter specifies whether BGP proxy aggregation can run in the BGP speaker or not. The BGP proxy aggregation is used to enable the BGP speaker to aggregate routes, even if these routes are not originated by the local AS.</p> <p>■ Note This parameter is set to Disable and is not displayed if BGP Aggregation is set to Disable.</p>
Boot Effect:	<p>Disable->Enable: BGP Proxy Aggregation begins to run in this node. Aggregate routes are created and added in the BGP routing table. These routes can be advertised to its peers.</p> <p>Enable->Disable: BGP Proxy Aggregation stops running in this node. Aggregate routes are flushed in the BGP routing table and sent to its peers as withdrawn routes. Specific routes may be sent out as feasible NLRI</p>

BGP Aggregation Profiles

BGP Aggregation Profiles are new CMEM records. They are used to specify which routes can be aggregated to form aggregate routes and how to handle the specific routes.

The following records were created to incorporate these changes:

- BGP Aggregate Profiles
- IP Map Profiles
- BGP Aggregate Attribute Profiles

BGP Aggregation Profiles

Below is a table listing the BGP Aggregate Profiles descriptions:

Parameters	Description
Entry Number	This parameter uniquely identifies a BGP aggregation profile entry in this table. It also puts an upper bound to the number of entries that can be configured in this table.
Aggregate Address	This parameter specifies the IP address of aggregate route which is created based on the component routes in the BGP routing table.
Aggregate Mask	This parameter specifies the IP address mask of aggregate route which is created based on the component routes in the BGP routing table.

Parameters	Description (continued)
Include Network Profiles	This parameter specifies the entry numbers of IP Map Profile records. The IP address ranges in the designated IP Map Profile records are used to specify which route is eligible to be a component route. The aggregate route is created solely based on the component routes which are falling into the IP address ranges in the corresponding IP Map Profiles records.
Exclude Network Profiles	This parameter specifies the entry numbers of IP Map Profile records. The IP address ranges in the designated IP Map Profile records are used to specify which route is not eligible to be a component route. The aggregate route is created solely based on the component routes which are not falling into the IP address ranges in the corresponding IP Map Profiles records.
Peer List	This parameter specifies a list of BGP peers to which specific routes can be advertised with an aggregate route. To be successfully advertised to a BGP peer, a specific route must match a BGP outbound policy and the action type of the policy must be permit.
Specific Map Profiles	This parameter specifies the entry numbers of IP Map Profile records. The IP address ranges in the designated IP Map Profile records are used to specify which route is eligible to be advertised with the aggregate route.
Surpress Map Profiles	This parameter specifies the entry numbers of IP Map Profile records. The IP address ranges in the designated IP Map Profile record are used to specify which route is not eligible to be advertised with the aggregate route.
Path Setting	This parameter specifies the attribute setting method in the aggregate route.
Attribute Profile	This parameter specifies the entry number of a BGP Aggregate Attribute Profile record. The content of the designated record is applied to the aggregate route.
Append Community Profiles	This parameter specifies the well-known community attributes for the component routes.

Entry Number

Range:	Vanguard 7300 Series: 1 to 256 Vanguard 6435/6455: 1 to 64 Vanguard 34x: 1 to 32
Default:	1
Description:	This parameter uniquely identifies a BGP aggregation profile entry in this table.

Aggregate Address

Range:	Any valid IP address
Default:	0.0.0.0
Description:	This parameter specifies the IP address of an aggregate route which is created based on the component routes in the BGP routing table.
Boot Effect:	Any old aggregate route generated by this profile is flushed and a new aggregate route is created. This implies BGP route updates of withdrawing the old aggregate route and advertising the new aggregate route.

Aggregate Address Mask

Range:	Any valid IP address mask
Default:	255.0.0.0
Description:	This parameter specifies the IP address mask of the aggregate route which is created based on the component routes in the BGP routing table.
Boot Effect:	Any old aggregate route entry generated by this profile needs to be flushed and a new aggregate route is created. This implies BGP route updates of withdrawing the old aggregate route and advertising the new aggregate route.

Include Network Profiles

Range:	A list of entry number ranges of IP Map Profile records.
Default:	None
Description:	This parameter specifies the entry numbers of IP Map Profile record. The IP address ranges in the designated IP Map Profile records are used to specify which route is eligible to be a component route. The aggregate route is created solely based on the component routes which are falling into the IP address ranges in the corresponding IP Map Profiles records. If None is the input, there is no Include Network Profiles. Example: 1,2 - 5,6
Boot Effect:	The aggregate route defined by this profile needs to be checked, flushed if required and a new aggregate route is created. This implies BGP route updates of withdrawing the old aggregate route if required, and advertising a new aggregate route.

Exclude Network Profiles

Range:	A list of entry number ranges of IP Map Profile records, or None.
Default:	None
Description:	This parameter specifies the entry numbers of IP Map Profile records. The IP address ranges in the designated IP Map Profile records are used to specify which route is not eligible to be a component route. The aggregate route is created solely based on the component routes which are not falling into the IP address ranges in the corresponding IP Map Profiles records. If None is input, there is no Exclude Network Profiles. Example: 1,2 - 5,6
Boot Effect:	The aggregate route defined by this profile needs to be checked, flushed if required and a new aggregate route is created. This implies BGP route updates of withdrawing the old aggregate route if required, and advertising a new aggregate route.

Peer List

Range:	A list of ranges of BGP peer numbers, or “All”.
Default:	All
Description:	This parameter specifies a list of BGP peers to which specific routes can be advertised along with an aggregate route. To be successfully advertised to a BGP peer, a specific route must match a BGP outbound policy and the action type of the policy must be permit. If All is input, then all BGP peers are included. Example: 1,2,3-5
Boot Effect:	The aggregate route and specific routes need to be flushed in the old BGP peers, and advertised to the new defined peers.

Specific Map Profiles

Range:	A list of entry number ranges of IP Map Profiles records, or “None”.
Default:	None
Description:	This parameter specifies the entry numbers of IP Map Profile records. The IP address ranges in the designated IP Map Profile records are used to specify which route is eligible to be advertised with the aggregate route. If None is input, there is no Specific Map Profiles. Example: 1,2,3-5
Boot Effect:	The specific routes defined by this profile need to be flushed in the BGP peers and the new specific routes are advertised.

Suppress Map Profiles

Range:	A list of entry number ranges of IP Map Profiles records, or “None”.
Default:	None
Description:	This parameter specifies the entry numbers of IP Map Profile records. The IP address ranges in the designated IP Map Profile record are used to specify which route is not eligible to be advertised with the aggregate route. If None is input, there is no Suppress Map Profiles. Example: 1,2,3-5
Boot Effect:	The suppressed routes newly defined by this profile need to be flushed in the peers and the old defined routes can be advertised to the peers.

Path Setting

Range:	Automatic, Originator
Default:	Automatic
Description:	This parameter specifies the attribute setting method in the aggregate route. Automatic: The attribute in the aggregate route is calculated based the component routes automatically. No attribute in the component routes are lost. Originator: The attribute in the aggregate route is populated as a self-originated route.
Boot Effect:	The old aggregate route needs to be flushed and the new aggregate route is created in the BGP routing table. BGP updates are sent to peers to reflect these changes.

Attribute Profile

Range:	Entry number of BGP Aggregate Attribute Profiles record, or “None”.
Default:	None
Description:	This parameter specifies the entry number of BGP Aggregate Attribute Profile record. The content of the designated record is applied to the aggregate route. If None is the input, there is no Attribute Profile.
Boot Effect:	The old aggregate route needs to be flushed and a new aggregate route is created in BGP routing table. BGP updates are sent to peers to reflect these changes.

Append Community Profiles

Range:	NO_EXPORT, NO_ADVERTISE, NO_EXPORT_SUBCONFED, NO_PEER, NONE
Default:	NONE
Description:	This parameter specifies the well-known community attributes for the component routes.

IP Map Profiles

The IP Map Profiles are:

- Entry Number
- IP Networks

Entry Number

Range:	Vanguard 7300 Series: 1 to 512 Vanguard 6435/6455, Vanguard 34x: 1 to 256
Default:	1
Description:	This parameter uniquely identifies an IP Map Profile entry in this table. It also puts an upper bound to the number of entries that can be configured in this table.

IP Networks

Range:	A list of any valid IP address (ranges) with a valid mask length.
Default:	0.0.0.0/8
Description:	This parameter specifies a list of networks ranges. Each element of the list specifies one network range in IP Address/Mask Length format (e.g. X.X.X.X/yy). As many as 16 network elements can be specified in this parameter. Example: 150.6.1.0/24-150.6.16.0/24, 150.6.24.0/24
Boot Effect:	The BGP Aggregate Profiles records, which reference this IP Map Profiles record, need to be booted (table boot).

BGP Aggregate Attribute Profiles

The BGP Aggregate Attribute Profiles are:

- Entry Number
- MED
- Local_Pref
- Origin
- Next_hop
- AS_Path_Prepend

Entry Number

Range:	Vanguard 7300 Series: 1 to 128 Vanguard 6435/6455 and Vanguard 34x: 1 to 32
Default:	1
Description:	This parameter uniquely identifies a BGP Aggregate Attribute Profile entry in this table. It also puts an upper bound to the number of entries that can be configured in this table.

MED

Range:	0-0xFFFFFFFF, Automatic
Default:	Automatic
Description:	This parameter specifies the MED attribute of the corresponding aggregate route.If Automatic is input, the internally calculated MED value is used.
Boot Effect:	The BGP Aggregate Profiles records, which reference this profiles record, need to be booted (table boot).

Local_Pref

Range:	0-0xFFFFFFFF, Automatic
Default:	Automatic
Description:	This parameter specifies the Local Preference attribute of the corresponding aggregate route.If Automatic is input, the internally calculated Local_Pref are used.
Boot Effect:	The BGP Aggregate Profiles records, which reference this profiles record, need to be booted (table boot).

Origin

Range:	IGP, EGP, Incomplete, Automatic
Default:	Automatic
Description:	This parameter specifies the Origin attribute of the corresponding aggregate route. If Automatic is input, the internally calculated Origin value is used.
Boot Effect:	The BGP Aggregate Profiles records, which reference this profiles record, need to be booted (table boot).

Next_Hop

Range:	Any valid IP address
Default:	0.0.0.0
Description:	This parameter specifies the Next Hop attribute of the corresponding aggregate route. A value of 0.0.0.0 (default) implies use of the calculated next hop address.
Boot Effect:	The BGP Aggregate Profiles records, which reference this profiles record, need to be booted (table boot).

AS_Path_Prepend

Range:	List of up to ten AS numbers to be prepended to the AS_Path, or "None".
Default:	None
Description:	This parameter specifies the AS number(s) which are prepended to the calculated AS_Path attribute in the corresponding aggregate route. If None is input, no AS number is prepended in the AS_Path attribute. Example: 200,300,150
Boot Effect:	The BGP Aggregate Profiles records, which reference this profiles record, need to be booted (table boot).

Configuring BGP Community Attribute

BGP Community Attribute

BGP community attribute is configured through BGP community attribute profile. The profile ID should be used in the BGP policy configuration and BGP aggregation profiles. The BGP default policy should not include any community attribute configuration since the default value of BGP community attribute should be NONE.

Community Attribute Profiles are shown in Figure 3-4.

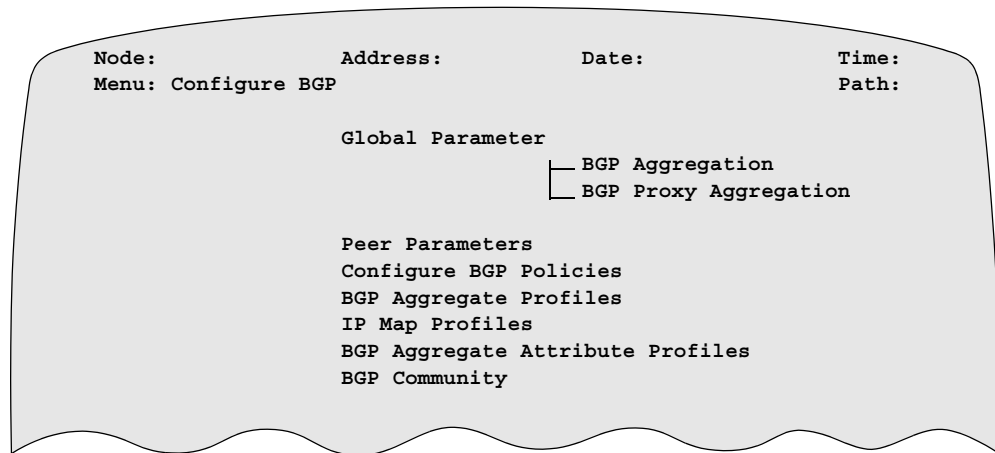


Figure 3-4. Configure BGP Menu

BGP Community Attribute in BGP Global Control Parameter Menu

Figure 3-4 shows BGP Community added to the end of the BGP Global control parameter menu.

User Interface

Community attribute profile is identified by record index number. The parameters in this record include the ASN and attribute value. Booting the community profile causes all the policies, which use this community profile, to be rebooted. Below are additional parameters:

Entry Number

Range:	1 to 128
Default:	
Description:	Index of this record.

AS Number

Range:	1 to 65534
Default:	MY_AS
Description:	This parameter specifies the AS number of a community attribute value. If default value, MY_AS, is configured, the AS number configured in BGP Global Parameter will be used. The wildcard (*) matching any AS number with the same community value, is only applicable for DELETE community profiles in Policy configuration.

Community Value

Range:	1 to 65535
Default:	1
Description:	This parameter specifies the value for the community attribute. The previous parameter, AS Number and this parameter, are combined together to define one element in the set of community value. The wildcard (*) matching any community value with the same AS number, is only applicable for DELETE community profiles in Policy configuration.

Chapter 4

Statistics and Diagnostics

Overview

Introduction

This chapter describes BGP routing statistics and diagnostics. These statistics provide the user with useful information related to BGP functional operation, which can aid network administrators to manage the network where BGP is deployed. This includes information regarding the BGP routes and functional status with the configured peers.

Using the BGP Peer Statistics

Introduction

The options under the BGP Peer Statistics menu are shown below. See Figure 4-2.

Status/statistics->Router Stats->BGP Statistics->BGP Peer Statistics




Figure 4-2. BGP Peer Statistics

Description of Terms

This table gives a description of the parameters displayed in the following BGP Peers Statistics screens.

Term	Description
State	<p>The parameter indicates the BGP connection state for the specified Peer. BGP Connection with a Peer can be in one of the following states.</p> <ol style="list-style-type: none"> 1. Disable: Peer is disabled by the Administrator. 2. Idle: An error has happened while establishing a BGP connection and the local BGP speaker is waiting for the expiry of the Restart timer to restart the connection establishing process. 3. Connecting: Waiting for a TCP connection to get established. 4. Active: An Error occurred while establishing the TCP connection. Still trying to establish the connection. 5. OpenSent: OPEN message has been sent to the peer, Waiting for an OPEN/NOTIFICATION from the peer. 6. OpenConfirm: KEEPALIVE has been sent against peer's OPEN. Waiting for a KEEPALIVE/NOTIFICATION from peer. 7. Established: BGP connection established. BGP messages (UPDATE, KEEPALIVE, NOTIFICATION) can be exchanged with the peer.

Term	Description (continued)
BGP ID	This parameter indicates the BGP Identifier of the specified Peer. This is the IP Address which belongs to the Peer and is received by local BGP speaker in OPEN message sent from the specified Peer. This Address uniquely identifies a Peer in the network. This address need not be an IP Address configured for the specified Peer at local BGP speaker.
IP Address	This specifies the IP Address belonging to the specified Peer, which is being used for the current BGP connection with this peer. BGP connection with a peer can be formed using any of the configured IP Addresses for it at local BGP speaker.
UP Since	This parameter specifies the time (in Days: Hours: Minutes: Seconds) since the BGP connection with the specified Peer is active. This is the instance of time at which the BGP connection was established with the Peer.
Total Number Peers	This parameter specifies the total number of BGP Peers currently configured at local BGP speaker. This includes both active and idle peers.

Display BGP Peers Summary

This option will be used to display the summary of all the BGP peerings with the configured BGP Peers. This summary includes information like BGP Connection state, BGP Peer Identifier, and the time since the BGP connection is established. See Figure 4-3.

```

Node: <Node name> Address: <node address> Date: <date>

BGP Peer Statistics Summary Page 1 of 1

Total Number of Peers: 3

PeerNo Peer AS State BGP ID Peer IP Addr Up Since
(DAY:HH:MM:SS)
1 100 Established 150.1.1.1 1.1.1.1 000:00:48:00
2 200 Idle 0.0.0.0 0.0.0.0 000:00:00:00
3 Local_AS Established 160.1.1.1 20.1.1.1
121:23:56:55

Press any key to continue (ESC to exit)...
    
```

Figure 4-3. Display BGP Peers Summary

Display BGP Peer Detailed Statistics

This option is used to display the detailed information about the BGP Peering with the specified Peer. Apart from the BGP Peer summary, the some of the peering details which are displayed under this options are TCP Connection details, Count for BGP packets exchanged, Count of BGP routes exchanged, and Count of BGP Connection Errors. See Figure 4-4.

```

Node: <Node name> Address: <node address> Date: <date>

BGP Peer Detailed Statistics

PeerNo Peer AS State BGP ID Peer IP Addr Up Since
(DAY:HH:MM:SS)
128 100 Established 150.1.1.1 1.1.1.1 000:00:48:00

TCP Connection Stats:
Connection Type : Active TCP Connection Error : xxxx
SPort : 1026 TCP State Transition : xxxx
DPort : 179 Hold Down/KeepAlive Time : 12/4
Press any key to continue (ESC to exit)..

```

Figure 4-4. Display BGP Peer Detailed Statistics

Reset BGP Peer Statistics

This option will be used to reset the BGP Peer Statistics for all or a particular Peer(s) specified by the user. This resetting operation will clear the following counters of the BGP Peer Statistics.

- TCP Connection Counters
- BGP Messages Transmit/Receive Counters
- BGP Routes Transmit/Receive Counters
- BGP Connection Error Counters

This option does not reset the BGP connection itself and therefore does not reset the rest of the BGP Peer statistics parameters specific to the current BGP connection.

Using the BGP Routing Table Statistics

Introduction

The options under the BGP Routing Table Statistics menu are shown below. See Figure 4-5.

Status/statistics->Router Stats->BGP Statistics->BGP Routing Table Statistics




Figure 4-5. BGP Routing Table Statistics

Description of Terms

Following table gives the description of the parameters displayed in the following mentioned BGP Routing Table Statistics screens.

Term	Description
Network	This is the IP Address of the network for which BGP route is displayed.
Mask	This is the IP Address mask of the network for which BGP route is displayed.
Next Hop	This is the address of the next hop router to which the packet should be forwarded for the specified destination network. A value of 0.0.0.0 indicates that packet should be forwarded to the local BGP speaker itself.
Origin	This specifies the 'Origin' attribute of the BGP route being displayed. The value of this parameter could be one of the following. IGP: for the routes originated by an IGP for example RIP or OSPF. EGP: for the routes originated by an EGP for example EGP. INC: for the routes generated through other means for example Static.

Term	Description (continued)
AGR_AS	<p>This specifies the Aggregating AS Number for the specified BGP route i.e. the AS Number of the BGP speaker which performed the aggregation for the specified route.</p> <p>A value of '0' indicates that specified route is not an aggregated route.</p> <p>A string 'Local_AS' indicates that the aggregation was done by the local BGP speaker.</p>
AAG	<p>This specifies whether the BGP route was atomic aggregated i.e. it was chosen as a less specific route in presence of more specific routes.</p>
AS PATH	<p>This specifies the 'AS Path' attribute of the BGP route.</p> <p>An AS Sequence is represented as AS numbers listed within '<>' while an AS Set is represented as AS numbers listed within '[']'.</p> <p>A string '<Local_AS>' is displayed for this parameter against the BGP routes, originated by the local BGP speaker.</p>
Total No. of routes	<p>This specifies the total number of routes present in the BGP routing table which are currently selected as the best path by the local BGP speaker after applying BGP Inbound Policies and running its decision process.</p>
Total No. of Route Changes	<p>This specifies the number of times the BGP routing table has changed as a result of addition or deletion of a BGP route.</p>
Total No. of Internal routes	<p>This specifies the total number of routes present in the BGP routing table which were originated by the local BGP speaker.</p>
Total No. of External routes	<p>This specifies the total number of routes present in the BGP routing table which were learned from BGP peers.</p>
Total No. of Routes Sent	<p>This specifies the total number of routes which were advertised by the local BGP speaker to a particular/any BGP Peer(s), after applying BGP Outbound Policies.</p>
Total No. of Paths Sent	<p>This specifies the total number of paths which were advertised by the local BGP speaker to a particular/any BGP Peer(s), after applying BGP Outbound Policies.</p>

Display All BGP Routes

This option will be used to display all the routes present in the BGP routing table. Routes displayed under this option are the best routes, currently being used by the local BGP speaker, for the respective networks. BGP Speaker computes the best BGP route by applying BGP Inbound Policies and running its decision process. See Figure 4-6.

```

Node: <Node name> Address: <node address> Date: <date> Time: <Time>
BGP Routing Table Page 1 of 1

Dest_Addr Mask Next_Hop Origin AGR_AS AAG
AS_PATH
150.1.1.0 ffffff00 200.1.1.1 IGP 0 No
<100,200,300>
180.1.1.0 ffffff00 < " >

120.1.0.0 ffff0000 200.1.1.1 INC 400 No
<100,200>[400,500]

Total No.of Routes Sent: 1234
Total No.of Paths Sent: 123

Press any key to continue (ESC to exit)...
    
```

Figure 4-6. Display All BGP Routes

Display All Internal BGP Routes

This option is used to display all the internal BGP routes, originated by the local BGP speaker. See Figure 4-7.

```

Node: <Node name> Address: <node address> Date: <date> Time:
<Time>
BGP Routing Table Page 1 of 1

Total No.of Internal Routes: 2823

Dest_Addr Mask Next_Hop Origin AGR_AS AAG
AS_PATH
170.1.1.0 ffffffff00 0.0.0.0 IGP 0 No
<Local AS>
100.1.0.0 ffff0000 0.0.0.0 IGP Local_AS No
<Local_AS>

Press any key to continue (ESC to exit)...
    
```

Figure 4-7. Display All Internal BGP Routes

Display All External BGP Routes

This option is used to display all the external BGP routes learned, by local BGP speaker from its Peers. Routes displayed under this option are the best external routes, currently being used by the local BGP speaker, for the respective networks. The BGP Speaker computes the best BGP route by applying BGP Inbound Policies and running its decision process. See Figure 4-8.

```

Node: <Node name> Address: <node address> Date: <date> Time:
<Time>
BGP Routing Table Page 1 of 1

Total No.of External Routes: 7000

Dest_Addr      Mask      Next_Hop      Origin  AGR_AS  AAG
AS_PATH
150.1.1.0      ffffffff00  200.1.1.1     IGP     0        No
<100,200,300>
120.1.0.0      ffff0000   200.1.1.1     INC     60       No
<100,200>[400,500]

```

Figure 4-8. Display All External BGP Routes

Display BGP Route for a Particular Network

This option is used to display the most specific BGP route for a destination IP Address, specified by the user. Route displayed under this option is the best BGP route (as computed by the local BGP speaker) for the specified destination. See Figure 4-9.

```

Node: <Node name> Address: <node address> Date: <date> Time:
<Time>
BGP Routing Table Page 1 of 1

Total No.of External Routes: 7000

Dest_Addr      Mask      Next_Hop      Origin  AGR_AS  AAG
AS_PATH
150.1.1.0      ffffffff00  200.1.1.1     IGP     0        No
<100,200,300>
120.1.0.0      ffff0000   200.1.1.1     INC     60       No
<100,200>[400,500]

Press any key to continue (ESC to exit)...

```

Figure 4-9. Display BGP Route for a Particular Network

Display BGP Routes Advertised to a Peer

This option is used to display all the BGP routes advertised to a particular BGP Peer as specified by the user. This option shows users all the BGP routes which are advertised to particular BGP Peers, after applying BGP Outbound Policies. See Figure 4-10.

```

Node: <Node name> Address: <node address> Date: <date> Time:
<Time>
BGP Routing Table Page 1 of 1

Dest_Addr      Mask      Next_Hop      Origin      AGR_AS      AAG
                AS_PATH
150.1.1.0      ffffffff00  200.1.1.1     IGP         0            No
                <100,200,300>
180.1.1.0      ffffffff00  <              "            >
120.1.0.0      ffff0000   200.1.1.1     INC         400          No
                <100,200>[400,500]

Total No.of Routes Sent: 1234
Total No.of Paths Sent: 123

Press any key to continue (ESC to exit)...
    
```

Figure 4-10. Display BGP Routes Advertised to Peer

BGP Community Attributes

BGP Community Attribute statistics are displayed in the BGP routing Table:

```

Node: Top      Address: 200      Date: 14-MAR-2003      Time: 10:15:31
BGP Aggregate Routes      BGP Aggregation Profiles No. 123

Dest_Addr      Mask      Next_Hop      Origin      AGR_AS      AAG
                AS_Path      Community
150.1.0.0      ffff0000   200.1.1.1     IGP         LOCAL_AS    Yes
                <100>      NO_EXPORT
150.1.1.0      ffffffff00  200.1.1.1     IGP         0            No
                <100>
150.1.2.0      ffffffff00  200.1.1.1     IGP         0            No
                <100>
150.1.4.0      ffffffff00  200.1.1.1     IGP         0            No
                <100>
150.1.5.0      ffffffff00  200.1.1.1     IGP         0            No
                <100>
    
```

Figure 4-11. BGP Routing Table with Community Attributes

Community attribute statistics are displayed by listing all the applicable community values. For user-defined values, the ASN:value format is used (for example: 2003:1115). Well-known community values:

- 0xFFFFFFFF01 is NO_EXPORT
- 0xFFFFFFFF02 is NO_ADVERTISE
- 0xFFFFFFFF03 is NO_EXPORT_SUBCONFED

BGP Path attribute shows up in the same style.

Using BGP AS Path Database Display

Introduction

The options under the BGP AS Path Database Display menu are shown below. See Figure 4-12.

Status/statistics->Router Stats->BGP Statistics->BGP AS Path Database Display

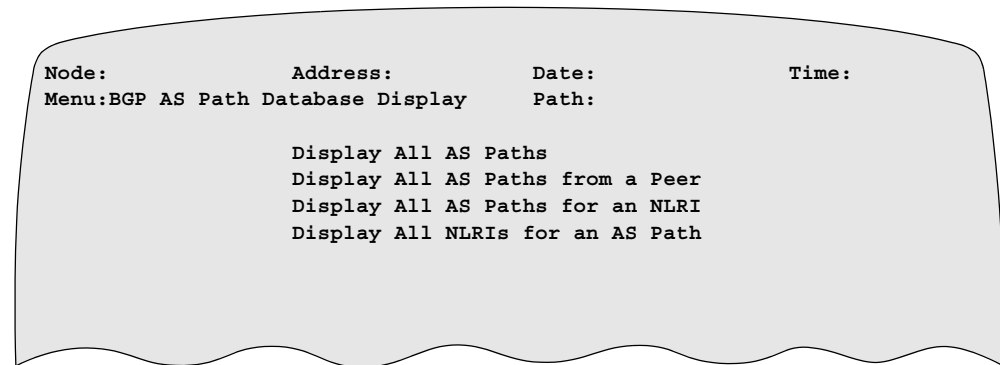


Figure 4-12. BGP AS Path Database Display

Description of Terms

This table gives a description of the parameters displayed in the following BGP AS Path Database Display screens.

Term	Description
Path ID	This parameter specifies a unique identifier for an AS Path in BGP AS Path Database. Path ID is assigned by the local BGP speaker for the purpose of uniquely identifying an AS Path in the AS Path Database.
NextHop	This is the address of the next hop router to which the packet should be forwarded for the destination network using the specified path. A value of 0.0.0.0 indicates that packet should be forwarded to the local BGP speaker itself.
Origin	This specifies the 'Origin' attribute of the AS Path being displayed. The value of this parameter could be one of the following. IGP: for the routes originated by an IGP for example RIP or OSPF. EGP: for the routes originated by an EGP for example EGP. INC: for the routes generated through other means for example Static.

Term	Description (continued)
AGR_AS	<p>This specifies the Aggregating AS Number for the specified AS Path i.e. the AS Number of the BGP speaker which generated the aggregated AS Path after aggregating the AS Paths of the corresponding NLRIs.</p> <p>A value of '0' indicates that specified AS path is not an aggregated AS Path.</p> <p>A string 'Local_AS' indicates that the aggregation was done by the local BGP speaker.</p>
AAG	<p>This specifies whether the BGP routes using the specified AS Path were atomic aggregated i.e. corresponding NLRI was chosen as a less specific route in presence of more specific routes.</p>
Peer No.	<p>This indicates the Peer from which the specified path was learned. A value of 'Local' is displayed for a locally originated path.</p>
MED	<p>This specifies the Multi Exit Discriminator(MED) attribute for the specified path. Locally originated paths are displayed with a value of 0 against this parameter.</p>
Local_Pref	<p>This specifies the Local Preference attribute for the specified path. Locally originated paths are displayed with a value of 0 against this parameter.</p>
DoP	<p>This parameter indicates the local BGP speaker's Degree of Preference for the specified path. Locally originated paths are displayed with a value of 0 against this parameter.</p>
RefCnt	<p>This parameter indicates the number of NLRIs referring to the specified AS Path.</p>
AS_PATH	<p>This specifies the 'AS Path' attribute of the specified AS Path.</p> <p>An AS Sequence is represented as AS numbers listed within '<>' while an AS Set is represented as AS numbers listed within '[']'.</p> <p>A string '<Local_AS>' is displayed for this parameter against the AS Paths, originated by the local BGP speaker.</p>
Total Number of AS Paths	<p>This specifies the total number of AS Paths present in the BGP AS Path Database which are currently stored by local BGP speaker after applying the BGP Inbound Policies.</p>
No.of AS Paths from this Peer	<p>This specifies the number of AS Paths learned from the specified Peer, after applying BGP Inbound Policies for that peer.</p>
No.of AS Paths for this NLRI	<p>This specifies the number of AS Paths available in local BGP Speaker's AS Path Database, for the specified NLRI.</p>
No.of NLRIs using this AS Path	<p>This specifies the number of NLRIs using the specified AS Paths.</p>

Display All AS Paths

This option is used to display all the AS Paths present in the BGP AS Path Database. AS Paths displayed under this option are the ones which are stored by the local BGP speaker after applying the BGP Inbound Policies. The paths displayed may or may not be the best paths for the corresponding NLRIs. See Figure 4-13.

```

Node: <Node name> Address: <node address> Date:
AS Path Database Display Page 1 of 1

No.of AS Paths from this Peer:2 * Best Path for atleast one network

Path Id  Next_Hop      Origin  AGR_AS  AAG  PeerNo  MED
         Local_Pref DoP      RefCnt  AS_PATH
*1       2.1.1.1       IGP     0        No    5        30
         50          50      4        <100,200,300>
100     10.1.1.1      INC     60       No    5        0
         30          25      3        <100,200>[400,500]

Press any key to continue (ESC to exit)...
    
```

Figure 4-13. Display All AS Paths

Display All AS Paths from a Peer

This option is used to display all the AS Paths learned from a particular BGP Peer as specified by the user. Paths displayed under this option are the ones which are stored by the local BGP speaker after applying the BGP Inbound Policies for the specified Peer. See Figure 4-14.

```

Node: <Node name> Address: <node address> Date: <date> Time:
<Time>
AS Path Database Display Page 1 of 1

No.of AS Paths from this Peer:2 * Best Path for atleast one network

Path Id  Next_Hop      Origin  AGR_AS  AAG  PeerNo  MED
         Local_Pref DoP      RefCnt  AS_PATH
*1       2.1.1.1       IGP     0        No    5        30
         50          50      4        <100,200,300>
100     10.1.1.1      INC     60       No    5        0
         30          25      3        <100,200>[400,500]

Press any key to continue (ESC to exit)...
    
```

Figure 4-14. Display All AS Paths from a Peer

Display All AS Paths for an NLRI

This option is used to display all the AS Paths present in the AS Path Database, for an NLRI, specified by the user. One of the AS Paths displayed under this option is the best path for the specified NLRI. See Figure 4-15.

```

Node: <Node name> Address: <node address> Date: <date> Time:
<Time>
AS Path Database Display Page 1 of 1

No. of AS Paths for this NLRI:2 * Best Path for atleast one network

Path Id Next_Hop Origin AGR_AS AAG PeerNo MED
Local_Pref DoP RefCnt AS_PATH

*1 2.1.1.1 IGP 0 No 5 30
50 50 4 <100,200,300>
100 10.1.1.1 INC 60 No 5 0
30 25 3 <100,200>[400,500]

Press any key to continue (ESC to exit)...
    
```

Figure 4-15. Display All AS Paths for an NLRI

Display All NLRIs for an AS Path

This option is used to display all the NLRIs for an AS Path specified by the user. User has to specify the "Path ID" associated with AS Path for which NLRIs need to be displayed. The specified path may or may not be the best path for the corresponding NLRIs displayed under this option. See Figure 4-16.

```

Node: <Node name> Address: <node address> Date: <date> Time:
<Time>
AS Path Database Display Page 1 of 1

No. of NLRIs using this AS Path:4 * Best Path for this NLRI

Path Id Next_Hop Origin AGR_AS AAG PeerNo MED
Local_Pref DoP RefCnt AS_PATH

1 2.1.1.1 IGP 0 No 5 30
50 50 4 <100,200,300>

<IP Address IP Mask> <IP Address IP Mask>
<100.1.0.0 FFFF0000> <200.1.0.0 FFF00000>
* <150.1.0.0 FFFF0000> <190.23.42.0 FFFFFF00>

Press any key to continue (ESC to exit)...
    
```

Figure 4-16. Display All NLRIs for an AS Path

Using BGP Aggregation Statistics

The statistics of BGP Aggregation could be viewed in two groups:

- BGP Routing Table
- BGP Aggregation Stats

The BGP routing table includes the aggregate routes which can be viewed by listing the routing table. The BGP routing table is displayed as it is now. The BGP aggregation statistics can be used to display a specific aggregate route and all the specific routes.

Status/statistics->Router Stats->BGP Statistics->BGP Aggregation

```

Enter Aggregation No.: 123
      1      2      3      4      5      6      7
123456789012345678901234567890123456789012345678901234567890
1  Node: Top      Address: 200      Date: 14-MAR-1998      Time: 0:15:31
2  BGP Aggregate Routes      BGP Aggregation Profiles No. 123
3
4  Dest_Addr      Mask      Next_Hop      Origin      AGR_AS      AAG
5                AS_Path
6  150.1.0.0      ffff0000      200.1.1.1      IGP      LOCAL_AS      Yes
7                <100>
8  150.1.1.0      fffffff0      200.1.1.1      IGP      0      No
9                <100>
10 150.1.2.0      fffffff0      200.1.1.1      IGP      0      No
11                <100>
12 150.1.4.0      fffffff0      200.1.1.1      IGP      0      No
13                <100>
14 150.1.5.0      fffffff0      200.1.1.1      IGP      0      No
15                <100>
18 Press any key to continue (ESC to exit) ...
    
```

Figure 4-17. BGP Aggregation Statistics

BGP to RIPv2 Redistribution Statistics

The statistics of **BGP->RIPv2** Redistribution is organized as following:

Status/Statistics->Router Stats->BGP->RIP Stats

In the IP Routing Table, all the imported BGP routes are displayed if the routes are the best routes for the respective networks. The metric of imported BGP route is used to decide which is the best if more than one route exists for a specific network.

The BGP RIPv2 Redistribution Stats can be further grouped as following:

```

Node: 6.3_TEST Address: 63 Date: 31-JAN-1999 Time: 20:52\
Menu: BGP->RIP Stats Path: (Main.5.16.4)

1. BGP->RIP General Stats
2. BGP->RIP Imported Routes
3. BGP->RIP Policy Stats
4. BGP->RIP General Stats Reset
    
```

Figure 4-18. BGP RIP Statistics

The General Stats displays the general statistics, and debug information is displayed if `TURN_ON_DEBUG` is set in the software key table. The Imported BGP Routes Stats displays the redistribution information for a specific imported BGP route. The Redistribution Policy Stats displays the imported routing entry information by policy and all the imported routing entries if required. The General Stats Reset, resets the statistics. Figure 4-19 displays an IP Routing Table Screen.

```

Node: <Node name> Address: <node address> Date: <date> Time:
<Time>
IP Routing Table
 * Static/Direct Route
 % RIP Route Control

Type Dest Net Mask MetricAge Next Hop
Sbnt 150.30.00.0 ffff0000 1 0 None
BGP 150.30.33.0 ffffff00 3 20 150.30.83.1
BGP 150.30.34.0 ffffff00 3 20 150.30.83.1

Press any key to continue (ESC to exit)...
    
```

Figure 4-19. IP Routing Table

■ **Note**

Imported BGP routes can be identified in the IP Route Table by the “Type” field shown in 4-19.

```
1 2 3 4 5 6 7
1234567890123456789012345678901234567890123456789012345678901234567890
1 Node: Top      Address: 200      Date: 14-MAR-1998  Time: 0:15:31
2 BGP->RIPv2 Route Redistribution Stats
3 Route Stats:
4   Total Number of Imported & Rejected Route: 45
5   Imported Routes: 34                Rejected Routes: 11
6   BGP Aggregated Routes: 10         BGP Specific Routes: 24
7   Aggregative Routes: 20            Non-aggregative Routes: 14
8 Policy Stats:
9   Total Number of Policies: 56       Unusable Policies: 12
10  Most Recently Used Policy (Permit): 12  Route: 198.154.12.0
11  Most Recently Used Policy (Deny): 11   Route: 190.132.0.0
12 Debug:
13  Most Recently Imported Route In BGP Routing Table: 0xab01234
14  Most Recently Imported Route In IP Routing Table: 0x01234568
15  Most Recently Rejected Route In BGP Routing Table: 0xfd543210
16  MRU Policy (Permit): 0x34432312      MRU Policy (Deny): 0x12435680
17
18 Press any key to continue (ESC to exit) ...
```

Figure 4-20. General Statistics

Note

The Debug information is displayed only if configured.

```

Enter Destination of IP address: 150.1.1.2
      (press RETURN key or input 0.0.0.0 means any imported routes)
      1         2         3         4         5         6         7
1234567890123456789012345678901234567890123456789012345678901234567890

1  Node: Top      Address: 200      Date: 14-MAR-1998  Time: 0:15:31
2  BGP->RIPv2 Imported Routes
3
4  Dest_Addr  Mask      Metric  Next_Hop  Tag  Policy No. BGP
5              AS_Path              Aggr
6  150.1.1.0  ffffffff00  5      200.1.1.1  65501 121      No
7              <100,200,300>
8
9
10
11
18 Press any key to continue (ESC to exit) ...
    
```

Figure 4-21. Imported BGP Route Statistics

```

Enter the policy number: 121 (press RETURN key or input 0 means any
policy)
      1         2         3         4         5         6         7
1234567890123456789012345678901234567890123456789012345678901234567890

1  Node: Top      Address: 200      Date: 14-MAR-1998  Time: 0:15:31
2  BGP->RIPv2 Redistribution Policy Stats
3
4  Policy No. Dest_Addr  Mask      Metric  Next_Hop  Tag  BGP
5              AS_Path              Aggr
6  121        150.1.1.0  ffffffff00  5      200.1.1.1  65501 No
7              <100,200,300>
8  121        150.1.2.0  ffffffff00  5      200.1.1.1  300  No
9              <200,400,300>
10 121        150.1.0.0  ffff0000  5      200.1.1.1  400  Yes
11              [100,200,400,300]
12
18 Press any key to continue (ESC to exit) ...
    
```

Figure 4-22. Redistribution Policy Statistics

Display Parameter Description

The following table gives the description of the parameters displayed in figures 4-16 to 4-20.

Parameters	Description
Dest_Addr	This is the IP address of the network for which imported route is displayed.
Mask	This is the IP address mask of the network for which imported route is displayed.
Next Hop	This is the address of the next hop router to which the packet should be forwarded for the specified destination network. A value of 0.0.0.0 indicates that packet should be forwarded to the local router itself.
AS Path	This specifies the 'AS Path' attribute of the original BGP route. An AS Sequence is represented as AS numbers listed within '<>' while an AS Set is represented as AS numbers listed '[']'.
Metric	This is the metric of the imported route.
Policy No.	This is the entry number of the BGP->RIPv2 Redistribution policy, which is used to import a BGP route into RIPv2. A value of 0 indicates the default import policy.
Tag	This is the tag value of the RIPv2 route. This value can be set through user's configuration or through automatic generation. A value of 0 means the user disable the setting of tag value.
BGP Aggregated	This specifies if the original BGP route is aggregate route or not.
Total No. of imported and rejected routes	This specifies the total number of routes that satisfy the conditions of policy selectors, either permit to import into RIPv2 routing domain or deny to enter RIPv2 routing domain.
Imported Routes	This specifies the total number of routes which have been imported.
Rejected Routes	This specifies the total number of routes which have been rejected.
BGP Aggregated Routes	This specifies the total number of BGP aggregate routes which have been imported.
BGP Aggregation Profiles No.	This specifies the BGP aggregation Profiles which are used to form the aggregate route.
BGP Specific Routes	This specifies the total number of BGP specific routes which have been imported.

Parameters	Description
Aggregative Routes	This specifies the total number of imported routes that are eligible to be aggregated.
Non-Aggregative Routes	This specifies the total number of imported routes that are not eligible to be aggregated.
Total number of policy	This specifies the total number of configured policies, valid or invalid, used or unused.
Unusable Policy	This specifies the total number of usable policies.
Most Recently Used Policy	This specifies the entry number of the most recently used policy whose action type is either Permit or Deny.
Route	This specifies the BGP route which is the object of the most recently used policy. The route could be imported or rejected.
Path ID	This parameter specifies a unique identifier for an AS Path in BGP AS Path Database. Path ID is assigned by the local BGP speaker for the purpose of uniquely identifying an AS Path in it's AS Path Database.
Debug	This is the memory address of the useful debug information.
Origin	This specifies the 'Origin' attribute of the BGP route being displayed. The value of this parameter could be one of the following. IGP: for the routes originated by an IGP for e.g. RIP, OSPF etc. BGP: for the routes originated by an EGP for e.g. BGP etc. INC: for the routes generated through other means for e.g. Static etc.
AGR_AS	This specifies the Aggregating AS Number for the specified BGP route i.e. the AS Number of the BGP speaker which performed the aggregation for the specified route. A value of '0' indicates that specified route is not an aggregated route. A string 'Local_AS' indicates that the aggregation was done by the local BGP speaker.
AAG	This specifies whether the BGP route was atomic aggregated i.e. it was chosen as a less specific route in presence of more specific routes.
UP Since	This parameter specifies the duration in time (in Days: Hours: Minutes: Seconds) for which the BGP connection with the specified Peer is active. The time is accumulated from the instance at which the BGP connection was established with the Peer.

Parameters	Description
Total Number Peers	This parameter specifies the total number of BGP Peers currently configured at local BGP speaker. This includes both active & idle peers.
DoP	This parameter indicates the local BGP speaker's Degree of Preference for the specified path. Locally originated paths are displayed with a value of 0 against this parameter.
RefCnt	This parameter indicates the number of NLRIs referring to the specified AS Path.
Local_Pref	This specifies the Local Preference attribute for the specified path. Locally originated paths are displayed with a value of 0 against this parameter.
MED	This specifies the Multi Exit Discriminator (MED) attribute for the specified path. Locally originated paths are displayed with a value of 0 against this parameter.
Peer No.	This indicates the Peer from which the specified path was learned. A value of 'Local' is displayed for a locally originated path.

BGP Diagnostics

Diagnostics

A BGP Debug menu is available to create, modify and delete a BGP route manually. Diagnostics can be performed on BGP RIPv2 Route Redistribution and BGP Aggregation functions. The three menu selections are listed below in Figure 4-23.

Main->Debug->Debug BGP




Figure 4-23. Debug BGP

Border Gateway Protocol Glossary:

AS	Autonomous System.
ASBR	Autonomous System Border Router.
ASN	AS number. A community value is normally represented as “ASN:value”, value is the community attribute value.
BGP	Border Gateway Protocol.
BRT	BGP Routing Table.
CIDR	Classless Inter-Domain Routing.
Community	A group of destinations which share some common property.
Community Attribute	A 32 bits data block, which defines an optional transitive attribute for community. Normally the first two octets are the AS number, the last two octets are user-defined value for the community. Some well-known community attributes are also defined.
DBR	Destination Based Routing.
DOP	Degree of Preference.
EBGP	External BGP.
EGP	Exterior Gateway Protocol.
Extended Community Attribute	An upgraded version of BGP Community Attribute. An 8 octets data block for an extended community, type and data are included in the data block.
GRT	Global Routing Table.
IBGP	Internal BGP.
IDRP	Inter Domain Routing Protocol.
IGP	Interior Gateway Protocol.
IP	Internet Protocol.
ISP	Internet Service Provider.
Load Sharing	BGP route could have multiple paths to a destination. The traffic to the destination is shared among the multiple paths in a predefined manner.
MED	Multi Exit Discriminator.
MIB	Management Information Base.
NLRI	Network Layer Reachability Information.

NO_ADVERTISE	A well-known community attribute. All routes carrying a community attribute containing this value must not be advertised to other BPG peers.
NO_EXPORT	A well-known community attribute. All routes carrying this community attribute must not be advertised outside of a BGP confederation boundary. In Vanguard BGP implementation, this value instructs the BGP speaker not to advertise a route over an EBPG peer.
NO_EXPORT_SUBCONFED	A well-known community attribute. All routes carrying a community attribute containing this value must not be advertised to EBGP peers. In our BGP implementation, this value instructs the BGP speaker not to advertise a route over an EBPG peer.
NO_PEER	A well-known community attribute. All routes carrying this community value to constrain their propagation only to transit providers and not peers.
OSPF	Open Shortest Path First.
PBR	Policy Based Routing.
PIB	Policy Information Base.
RIB	Routing Information Base.
RIP	Routing Information Protocol.
RIPv2	Routing Information Protocol Version Two.
Route Backup	BGP route could have multiple paths to a destination. Once the primary one is out of order, the next path (backup path) will be chosen to keep connectivity.
SNMP	Simple Network Management Protocol.

A

Accessing BGP Parameter Records 3-2
Aggregation 2-19
AS Boundary Routing Parameters 3-33, 3-39
AS Paths 4-13
AS Paths for an NLRI 4-14
AS Paths from a Peer 4-13
AS Support 1-15
AS topology 2-4
AS Traffic 1-14

B

Basic Routing Configuration 3-1
BGP 4-5
BGP Community Attribute
 Configuring 3-57
BGP Connection 1-8
BGP Record Parameters
 Accessing 3-2
BGP Routes Advertised to Peer 4-10
BGP Routing Parameters
 Configuring 3-4
BGP Speaker 1-7
Border Gateway Protocol 1-1
Branch Network Scenario 2-5

C

CIDR enabled 1-2
Community Attribute 2-28
 Application Examples 2-30
Complex AS Topology Support 1-3
Complex Policy Support 1-3
Configuration
 Aggregation Parameters 3-48
 BGP routing parameters 3-4
 Community Attribute Parameters 3-57
 OSPF routing parameters 3-33, 3-39, 3-48
 RIPv2 Route Redistribution Parameters 3-39
Configuring BGP Community Attribute 3-57
Configuring BGP Routing Parameters 3-4

D

Diagnostics 4-22
Display All BGP Routes 4-8
Display All External BGP Routes 4-9
Display All Internal BGP Routes 4-8
Display BGP Route for a Destination 4-9

E

External Peers 1-7

I

Import BGP Routes 3-33
Import Policy 3-5
inite State Machine Error 1-10
Inter ISP Communication 2-4

Internal Peers 1-7
IP Address 3-15
IP Address Mask 3-15

K

Keep Alive Time 3-11

M

Multihome AS 1-15
Multihoming 2-6

N

NLRI 1-9, 4-14

O

Opening and confirming a BGP connection with the
 neighbor 1-9
OSPF 2-5
OSPF Import Policy Table 3-34
OSPF routing parameters
 Configuration 3-33, 3-48
 description 3-4
OSPF routing paramters
 Configuration 3-39

P

PATH Attributes 1-9
Path Selection 1-5
Path Vector Protocol 1-2
Path Weight 3-17
Peer Weight 3-13
Physical Connection 1-8
Policy Enforcement 1-5

R

RIPv2 Redistribution 2-7
Routing Table 3-5

S

Solution 2-2
Statistics 4-2
 BGP Aggregation Statistics 4-15
 BGP AS Path Database Display 4-6, 4-11
 BGP Community Attributes Statistics 4-10
 BGP Peer Statistics 4-3
 BGP Routing Table Statistics 4-6
 BGP to RIPv2 Redistribution Statistics 4-16