ADVANCE WITH US

Ethernet Routing Switch
8000

Virtual Services Platform
4000, 7200, 8000, 9000

Engineering

> Border Gateway Protocol (BGP) Technical Configuration Guide

Extreme Networks
Document Date: November 2020
Part Number: 9036881-00
Revision AA

Notice

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

Documentation disclaimer

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

Link disclaimer

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

Warranty

Extreme Networks provides a limited warranty on Extreme Networks hardware and software. Refer to your sales agreement to establish the terms of the limited warranty. In addition, Extreme Networks’ standard warranty language, as well as information regarding support for this product while under warranty is available to Extreme Networks customers and other parties through the Extreme Networks Support website: http://www.extremenetworks.com/support under the link ““Policies” or such successor site as designated by Extreme Networks. Please note that if You acquired the product(s) from an authorized Extreme Networks Channel Partner outside of the United States and Canada, the warranty is provided to You by said Extreme Networks Channel Partner and not by Extreme Networks.

“Hosted Service” means an Extreme Networks hosted service subscription that You acquire from either Extreme Networks or an authorized Extreme Networks Channel Partner (as applicable) and which is described further in Hosted SAS or other service description documentation regarding the applicable hosted service. If You purchase a Hosted Service subscription, the foregoing limited warranty may not apply but You may be entitled to support services in connection with the Hosted Service as described further in your service description documents for the applicable Hosted Service. Contact Extreme Networks or Extreme Networks Channel Partner (as applicable) for more information.

Hosted Service

THE FOLLOWING APPLIES ONLY IF YOU PURCHASE AN EXTREME NETWORKS HOSTED SERVICE SUBSCRIPTION FROM EXTREME NETWORKS OR AN EXTREME NETWORKS CHANNEL

PARTNER (AS APPLICABLE), THE TERMS OF USE FOR HOSTED SERVICES ARE AVAILABLE ON THE EXTREME NETWORKS WEBSITE, https://extremeportal.force.com OR SUCH SUCCESSOR SITE AS DESIGNATED BY EXTREME NETWORKS, AND ARE APPLICABLE TO ANYONE WHO ACCESSES OR USES THE HOSTED SERVICE. BY ACCESSING OR USING THE HOSTED SERVICE, OR AUTHORIZING OTHERS TO DO SO, YOU, ON BEHALF OF YOURSELF AND THE ENTITY FOR WHOM YOU ARE DOING SO (HEREINAFTER REFERRED TO INTERCHANGEABLY AS “YOU” AND “END USER”), AGREE TO THE TERMS OF USE. IF YOU ARE ACCEPTING THE TERMS OF USE ON BEHALF A COMPANY OR OTHER LEGAL ENTITY, YOU REPRESENT THAT YOU HAVE THE AUTHORITY TO BIND SUCH ENTITY TO THESE TERMS OF USE. IF YOU DO NOT HAVE SUCH AUTHORITY, OR IF YOU DO NOT WISH TO ACCEPT THESE TERMS OF USE, YOU MUST NOT ACCESS OR USE THE HOSTED SERVICE OR AUTHORIZE ANYONE TO ACCESS OR USE THE HOSTED SERVICE.

Licenses

THE SOFTWARE LICENSE TERMS AVAILABLE ON THE EXTREME NETWORKS WEBSITE, https://extremeportal.force.com OR SUCH SUCCESSOR SITE AS DESIGNATED BY EXTREME NETWORKS, ARE APPLICABLE TO ANYONE WHO DOWNLOADS, USES AND/OR INSTALLS EXTREME NETWORKS SOFTWARE, PURCHASED FROM EXTREME NETWORKS, INC., ANY EXTREME NETWORKS AFFILIATE, OR AN EXTREME NETWORKS CHANNEL PARTNER (AS APPLICABLE) UNDER A COMMERCIAL AGREEMENT WITH EXTREME NETWORKS OR AN EXTREME NETWORKS CHANNEL PARTNER. UNLESS OTHERWISE AGREED TO BY EXTREME NETWORKS IN WRITING, EXTREME NETWORKS DOES NOT EXTEND THIS LICENSE IF THE SOFTWARE WAS OBTAINED FROM ANYONE OTHER THAN EXTREME NETWORKS, AN EXTREME NETWORKS AFFILIATE OR AN EXTREME NETWORKS CHANNEL PARTNER; EXTREME NETWORKS RESERVES THE RIGHT TO TAKE LEGAL ACTION AGAINST YOU AND ANYONE ELSE USING OR SELLING THE SOFTWARE WITHOUT A LICENSE. BY INSTALLING, DOWNLOADING OR USING THE SOFTWARE, OR AUTHORIZING OTHERS TO DO SO, YOU, ON BEHALF OF YOURSELF AND THE ENTITY FOR WHOM YOU ARE INSTALLING, DOWNLOADING OR USING THE SOFTWARE (HEREINAFTER REFERRED TO INTERCHANGEABLY AS “YOU” AND “END USER”), AGREE TO THESE TERMS AND CONDITIONS AND CREATE A BINDING CONTRACT BETWEEN YOU AND EXTREME NETWORKS, INC. OR THE APPLICABLE EXTREME NETWORKS AFFILIATE (“EXTREME NETWORKS”).

Extreme Networks grants You a license within the scope of the license types described below. Where the order documentation does not expressly identify a license type, the applicable license will be a Designated System License as set forth below in the Designated System(s) License (DS) section as applicable. The applicable number of licenses and units of capacity for which the license is granted will be one (1), unless a different number of licenses or units of capacity is specified in the documentation or other materials available to You. “Software” means computer programs in object code, provided by Extreme Networks or an Extreme Networks Channel Partner, whether as stand-alone products, pre-installed on hardware products, and any upgrades, updates, patches, bug fixes, or modified versions thereto. “Designated Processor” means a single stand-alone computing device. “Server” means a set of Designated Processors that hosts (physically or virtually) a software application to be accessed by multiple users. “Instance” means a single copy of the Software executing at a particular time: (i) on one physical machine; or (ii) on one deployed software virtual machine (“VM”) or similar deployment.

License type(s)

Designated System(s) License (DS). End User may install and use each copy or an Instance of the Software only: 1) on a number of Designated Processors up to the number indicated in the order; or 2) up to the number of Instances of the Software as indicated in the order, Documentation, or as authorized by Extreme Networks in writing. Extreme Networks may require the Designated Processor(s) to be identified in the order by type, serial number, feature key, Instance, location or other specific designation, or to be provided by End User to Extreme Networks through electronic means established by Extreme Networks specifically for this purpose.

Copyright

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

Virtualization

The following applies if the product is deployed on a virtual machine. Each product has its own ordering code and license types. Note, unless otherwise stated, that each Instance of a product must be separately licensed and ordered. For example, if the end user customer or Extreme Networks Channel Partner would like to install two Instances of the same type of products, then two products of that type must be ordered.

Third Party Components

“Third Party Components” mean certain software programs or portions thereof included in the Software or Hosted Service may contain software (including open source software) distributed under third party agreements (“Third Party Components”), which contain terms regarding the rights to use certain portions of the Software (“Third Party Terms”). As required, information regarding distributed Linux OS source code (for those products that have distributed Linux OS source code) and identifying the copyright holders of the Third Party Components and the Third Party Terms that apply is available in the products, Documentation or on Extreme Networks’ website at:http://www.extremenetworks.com/support/policies/softwarelicensing or such successor site as designated by Extreme Networks. The open source software license terms provided as Third Party Terms are consistent with the license rights granted in these Software License Terms, and may contain additional rights benefiting You, such as modification and distribution of the open source software. The Third Party Terms shall take precedence over these Software License Terms, solely with respect to the applicable Third Party Components to the extent that these Software License Terms impose greater restrictions on You than the applicable Third Party Terms.

The following applies only if the H.264 (AVC) codec is distributed with the product. THIS PRODUCT IS LICENSED UNDER THE AVC PATENT PORTFOLIO LICENSE FOR THE PERSONAL USE OF A CONSUMER OR OTHER USES IN WHICH IT DOES NOT RECEIVE REMUNERATION TO (i) ENCODE VIDEO IN COMPLIANCE WITH THE AVC STANDARD (“AVC VIDEO”) AND/OR (ii) DECODE AVC VIDEO THAT WAS ENCODED BY A CONSUMER ENGAGED IN A PERSONAL ACTIVITY AND/OR WAS OBTAINED FROM A VIDEO PROVIDER LICENSED TO PROVIDE AVC VIDEO. NO LICENSE IS GRANTED OR SHALL BE IMPLIED FOR ANY OTHER USE. ADDITIONAL INFORMATION MAY BE OBTAINED FROM MPEG LA, L.L.C. SEE HTTP://WWW.MPEGLA.COM.

Service Provider

THE FOLLOWING APPLIES TO EXTREME NETWORKS CHANNEL PARTNER’S HOSTING OF EXTREME NETWORKS PRODUCTS OR SERVICES. THE PRODUCT OR HOSTED SERVICE MAY USE THIRD PARTY COMPONENTS SUBJECT TO THIRD PARTY TERMS AND REQUIRE A SERVICE PROVIDER TO BE INDEPENDENTLY LICENSED DIRECTLY FROM THE THIRD PARTY SUPPLIER. AN EXTREME NETWORKS CHANNEL PARTNER’S HOSTING OF EXTREME NETWORKS PRODUCTS MUST BE AUTHORIZED IN WRITING BY EXTREME NETWORKS AND IF THOSE HOSTED PRODUCTS USE OR EMBED CERTAIN THIRD PARTY SOFTWARE, INCLUDING BUT NOT LIMITED TO MICROSOFT SOFTWARE OR CODECS, THE EXTREME NETWORKS CHANNEL PARTNER IS REQUIRED TO INDEPENDENTLY OBTAIN ANY APPLICABLE LICENSE AGREEMENTS,

AT THE EXTREME NETWORKS CHANNEL PARTNER'S EXPENSE, DIRECTLY FROM THE APPLICABLE THIRD PARTY SUPPLIER. WITH RESPECT TO CODECS, IF THE EXTREME NETWORKS CHANNEL PARTNER IS HOSTING ANY PRODUCTS THAT USE OR EMBED THE G.729 CODEC, H.264 CODEC, OR H.265 CODEC, THE EXTREME NETWORKS CHANNEL PARTNER ACKNOWLEDGES AND AGREES THE EXTREME NETWORKS CHANNEL PARTNER IS RESPONSIBLE FOR ANY AND ALL RELATED FEES AND/OR ROYALTIES. THE G.729 CODEC IS LICENSED BY SIPRO LAB TELECOM INC. SEE WWW.SIPRO.COM/CONTACT.HTML. THE H.264 (AVC) CODEC IS LICENSED UNDER THE AVC PATENT PORTFOLIO LICENSE FOR THE PERSONAL USE OF A CONSUMER OR OTHER USES IN WHICH IT DOES NOT RECEIVE REMUNERATION TO: (I) ENCODE VIDEO IN COMPLIANCE WITH THE AVC STANDARD ("AVC VIDEO") AND/OR (II) DECODE AVC VIDEO THAT WAS ENCODED BY A CONSUMER ENGAGED IN A PERSONAL ACTIVITY AND/OR WAS OBTAINED FROM A VIDEO PROVIDER LICENSED TO PROVIDE AVC VIDEO. NO LICENSE IS GRANTED OR SHALL BE IMPLIED FOR ANY OTHER USE. ADDITIONAL INFORMATION FOR H.264 (AVC) AND H.265 (HEVC) CODECS MAY BE OBTAINED FROM MPEG LA, L.L.C. SEE HTTP:// WWW.MPEGLA.COM.

Compliance with Laws

You acknowledge and agree that it is Your responsibility for complying with any applicable laws and regulations, including, but not limited to laws and regulations related to call recording, data privacy, intellectual property, trade secret, fraud, and music performance rights, in the country or territory where the Extreme Networks product is used.

Preventing Toll Fraud

"Toll Fraud" is the unauthorized use of your telecommunications system by an unauthorized party (for example, a person who is not a corporate employee, agent, subcontractor, or is not working on your company's behalf). Be aware that there can be a risk of Toll Fraud associated with your system and that, if Toll Fraud occurs, it can result in substantial additional charges for your telecommunications services.

Security Vulnerabilities

Information about Extreme Networks' security support policies can be found in the Global Technical Assistance Center Knowledgebase at https://gtacknowledge.extremenetworks.com/.

Downloading Documentation

For the most current versions of Documentation, see the Extreme Networks Support website: http://documentation.extremenetworks.com, or such successor site as designated by Extreme Networks.

Contact Extreme Networks Support

See the Extreme Networks Support website:http:// www.extremenetworks.com/support for product or Hosted Service notices and articles, or to report a problem with your Extreme Networks product or Hosted Service. For a list of support telephone numbers and contact addresses, go to the Extreme Networks Support website:http://www.extremenetworks.com/support/contact/ (or such successor site as designated by Extreme Networks), scroll to the bottom of the page, and select Contact Extreme Networks Support.

Contact Avaya Support

See the Avaya Support website: https://support.avaya.com for product or Hosted Service notices and articles, or to report a problem with your Avaya product or Hosted Service. For a list of support telephone numbers and contact addresses, go to the Avaya Support website: https://support.avaya.com (or such successor site as designated by Avaya), scroll to the bottom of the page, and select Contact Avaya Support.

Trademarks

The trademarks, logos and service marks ("Marks") displayed in this site, the Documentation, Hosted Service(s), and product(s) provided by Extreme Networks are the registered or unregistered Marks of Extreme Networks, Inc., its affiliates, its licensors, its suppliers, or other third parties. Users are not

Abstract

This document provides examples on configuring BGP on the Extreme Ethernet Routing Switch 8000 and Virtual Services Platform 4000, 8000, 7200, and 9000. This document covers some of the more popular BGP commands and attributes and the command used to configure them.

Acronym Key

Throughout this guide the following acronyms will be used:

- AS: Autonomous System
- BGP: Border Gateway Protocol
- CIDR: Classless Inter-Domain Routing
- EBGP: External BGP
- EGP: Exterior gateway protocol
- IBGP: Internal BGP
- IGP: Interior gateway protocol

Revision Control

No	Date	Version	Revised By	Remarks
1	24/11/2014	3.0	John Vant Erve	Updated examples using CLI. Added ISIS/BGP redistribution
2	3/12/2014	3.1	John Vant Erve	Edited by Ludovico Stevens – ISIS/BGP redistribution configuration example
3	13/1/2016	3.2	John Vant Erve	Added VSP 7200 and BGP+

Table of Contents

Figures

Tables

Conventions

This section describes the text, image, and command conventions used in this document.

Symbols

Tip – Highlights a configuration or technical tip.

Note – Highlights important information to the reader.

Warning – Highlights important information about an action that may result in equipment damage, configuration or data loss.

Text

Bold text indicates emphasis.

Italic text in a Courier New font indicates text the user must enter or select in a menu item, button or command:

```
ERS5520-48T# show running-config
```

Output examples from Extreme devices are displayed in a Lucida Console font:

```
ERS5520-48T# show sys-info

Operation Mode:          Switch
MAC Address:             00-12-83-93-B0-00
PoE Module FW:           6370.4
Reset Count:             83
Last Reset Type:         Management Factory Reset
Power Status:            Primary Power
Autotopology:            Enabled
Pluggable Port 45:       None
Pluggable Port 46:       None
Pluggable Port 47:       None
Pluggable Port 48:       None
Base Unit Selection:     Non-base unit using rear-panel switch
sysDescr:                Ethernet Routing Switch 5520-48T-PWR
                         HW:02       FW:6.0.0.10  SW:v6.2.0.009
                         Mfg Date:12042004    HW Dev:H/W rev.02
```

1. Border Gateway Protocol (BGP) Overview

The Border Gateway Protocol (BGP) is an exterior gateway protocol that border routers use to exchange network reachability information with other BGP systems. BGP routers form peer relationships with other BGP routers. Using an entity called a BGP Speaker, BGP peers transmit and receive current routing information over a reliable transport layer connection, making periodic updates unnecessary. BGP can be used both within and between autonomous systems.

BGP peers exchange complete routing information only when they establish the peer connection. Thereafter, BGP peers exchange routing information in the form of routing updates. An update includes a network number, a list of autonomous systems that the routing information has passed through (the AS path), and other path attributes that describe the route to a set of destination networks. When multiple paths are available, BGP compares the path attributes to choose the preferred path.

In addition to exchanging BGP information between autonomous systems, you can use BGP to exchange BGP information between routers in the same AS. To differentiate between these uses, the latter is called interior BGP (IBGP).

2. BGP Features Supported by software release

Table 1: BGP Features by Software Release

Feature	ERS 8800	VSP 4000	VSP 7200	VSP 8000	VSP 9000
eBGP (GRT & VRF)	7.0	3.1	4.2.1.0	4.1.0	3.0
eBGP (GRT & VRF) and iBGP (GRT only)	7.0	4.2.0	4.2.1.0	4.2.0	3.0
4-byte AS	7.1.0	3.1.0	4.2.1	4.1.0	3.2.0
BGP+ RFC2545 (GRT only)	7.0	5.0	5.0	5.0	4.1.0

2.1 BGP Configuration Guideline

2.1.1 Configuration Guidelines

When configuring BGP parameters on the Extreme switch, at a minimum it must be configured with the following parameters

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

In addition, BGP Policies can be added to the BGP peer configuration to influence route decisions as we will demonstrate later on in the document.

The BGP Router ID by default is automatically derived from the OSPF Router ID. It is recommended to configure a circuitless IP address (CLIP) and to use this address as the OSPF Router ID. The CLIP address can also be referred to as a loopback address.

It should be noted that once BGP is configured, some parameter changes may require having either the BGP Global state or neighboring admin-state to be disabled/enabled. The CLI prompt will notify you if this is the case.

The BGP policies are dynamically modified. On the global level, the BGP redistribution has an apply command that causes the policy to be applied at that time. The BGP neighbor peer has a CLI command named '*restart soft-reconfiguration <in> <out>*' that allows policies to be applied without bringing down the peer.

The following are some examples of these commands:

To enable/disable BGP globally, enter:

```
router bgp <as> enable
no router bgp enable
```

To enable/disable a BGP neighbor, enter:

```
router bgp
  neighbor <ip address of neighbor or neighbor group name> enable
  no neighbor <ip address of neighbor or neighbor group name> enable
exit
```

To set BGP soft-reconfiguration, enter:

```
neighbor <ip address of neighbor or neighbor group name> soft-reconfiguration-in
enable
```

To restart a BGP peer after adding a route policy to a peer, enter:

```
ip bgp restart-bgp neighbor <ip address of neighbor or neighbor group name> soft-
reconfiguration <in/out>
```

To apply BGP policy redistribution, enter:

ip bgp apply redistribute

ip bgp apply redistribute <direct|isis|ospf|rip|static|vrf>

If using EDM, use the following commands:

EDM: To enable/disable BGP globally, enter:

```
IP->BGP->Generals->AdminStatus <enable/disable>
```

EDM: To disable a BGP neighbor, enter:

```
IP->BGP->Peers->RemoteAddr <IP address of peer> Enable <true/false>
```

EDM: To set BGP soft-reconfiguration, enter:

```
IP->BGP->Peers->RemoteAddr SoftReconfiguration <true/false>
```

3. Basic BGP Fundamentals

There are two types of BGP connections, external BGP (EBGP) and internal BGP (IBGP). Routers belonging to the same autonomous system (AS) and exchange BGP updates are referred to as running IBGP. Routers that belong to a different AS and exchange BGP updates are referred to as running EBGP. Within an AS, routers run an interior gateway protocol such as OSPF.

In Figure 1 shown below, the connections between Router-C in AS 40 to ERS 8000 switch 8008 and VSP 9000 switch 9001 in AS 20 are running EBGP. The connection between 8008 and 9001 is running IBGP.

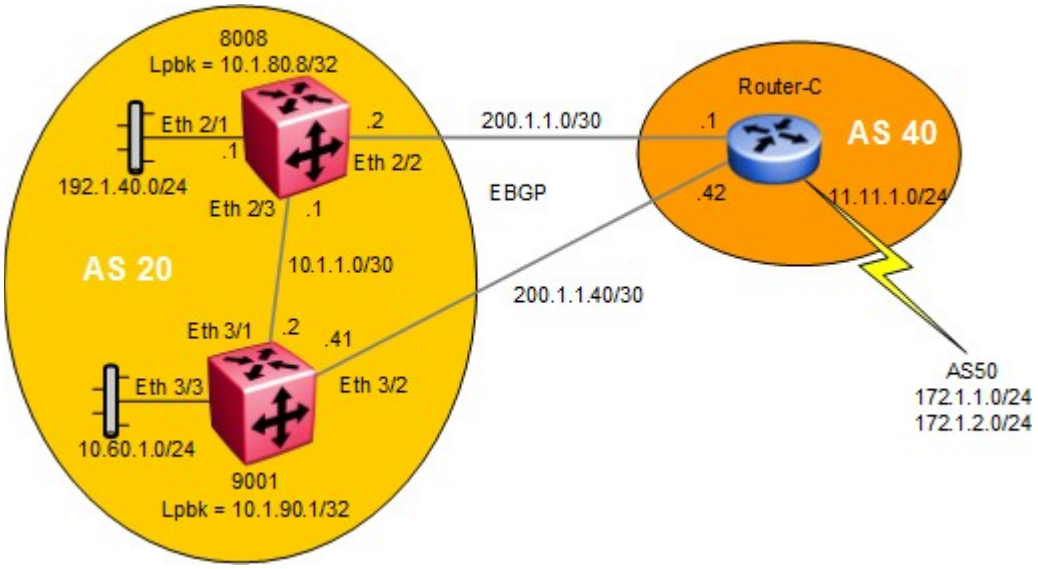

Figure 1: BGP Fundamentals

To configure a router for basic BGP operations, the following parameters must be configured:

- The Local AS number
- The BGP Router ID
    - By default, the BGP Router ID will automatically use the OSPF Router ID. As BGP uses the OSPF router ID, they cannot be different. A change in the router ID will require a BGP restart to take effect.
    - It is recommended to use a loopback IP (also known as a circuitless IP or CLIP) address for the OSPF Router-ID which in turn also becomes the BGP Router-ID. The CLIP address can also be referred to as a loopback address. This IP address is used in BGP Update messages. This will help for trouble-shooting purposes to give you an idea where the updates are coming from.
- The BGP neighbor peer(s) which can be iBGP or/and eBGP.
    - If iBGP, the remote-as will be the same
    - If eBGP, the remote-as will be different

For example, the following commands are used to configure BGP on ERS 8000 switch named 8008 and VSP 9000 switch named 9001

3.1 Basic BGP Configuration Example

3.1.1 Configure 8008 and 9001

3.1.1.1 Configure Loopback IP

For this example, we will simply select loopback instance 1 although any instance number from 1 to 256 can be used.

8008 and 9001: Add a loopback address using instance 1 and enable OSPF

8008:

```
8008:5(config)#interface loopback 1
8008:5(config-if)#interface ip address 10.1.80.8/32
8008:5(config-if)#ip ospf
```

9001:

```
9001:1(config)#interface loopback 1
9001:1(config-if)#interface ip address 10.1.90.1/32
9001:1(config-if)#ip ospf
```

3.1.1.2 Configure Ports with Appropriate IP Address

For this configuration example, we will use brouter ports. If you wish to use more than one port on a switch, it is recommended to add a normal VLAN instead of using brouter ports.

8008 and 9001: Add IP addresses to brouter ports

8008:

```
8008:5(config)#interface gigabitEthernet 2/2
8008:5(config-if)#brouter vlan 2090 subnet 200.1.1.2/30
8008:5(config-if)#exit
8008:5(config)#interface gigabitEthernet 2/3
8008:5(config-if)#brouter vlan 2091 subnet 10.1.1.1/30
8008:5(config-if)#exit
8008:5(config)#interface gigabitEthernet 2/1
8008:5(config-if)#brouter vlan 2092 subnet 192.1.40.1/24
8008:5(config-if)#exit
```

9001:

```
9001:1(config)#interface gigabitEthernet 3/1
9001:1(config-if)#brouter vlan 2090 subnet 10.1.1.2/30
9001:1(config-if)#exit
9001:1(config)#interface gigabitEthernet 3/2
```

```
9001:1(config-if)#brouter vlan 2092 subnet 200.1.1.41/30
9001:1(config-if)#exit
```

In the configuration above, we are using Ethernet Routing Switch 8000 and VSP 9000 brouter ports as the BGP EBGP and IBGP interfaces. Hence, the reason VLAN ID's of 2090, 2091, and 2092 are used. Either a brouter port or a VLAN can be configured as the BGP interface. To display the brouter port VLAN ID's, use the following command:

- `8008:5#show vlan brouter-port`

3.1.1.3 Enable OSPF Globally and on Brouter Ports

8008 and 9001: Enable OSPF AS boundary router (ASBR), add loopback address as the OSPF router-id, and enable OSPF on brouter ports

8008:

```
8008:5(config)#router ospf
8008:5(config-ospf)#as-boundary-router enable
8008:5(config-ospf)#router-id 10.1.80.8
8008:5(config-ospf)#exit
8008:5(config)#router ospf enable
8008:5(config)#interface gigabitEthernet 2/1,2/3
8008:5(config-if)#ip ospf enable
```

9001:

```
9001:1(config)#router ospf
9001:1(config-ospf)#as-boundary-router enable
9001:1(config-ospf)#router-id 10.1.90.1
9001:1(config-ospf)#exit
9001:1(config)#router ospf enable
9001:1(config)#interface gigabitEthernet 3/1,3/3
9001:1(config-if)#ip ospf enable
```

Please note that the BGP router-id is derived from the OSPF router-id. In this example, the BGP router-id will become the CLIP address of 10.1.80.8.

3.1.1.4 Configure BGP Globally

8008 and 9001: Assign both switch to BGP AS20 and disable synchronization

8008:

```
8008:5(config)#router bgp 20 enable
8008:5(config)#router bgp
8008:5(router-bgp)#no synchronization
```

9001:

```
9001:1(config)#router bgp 20 enable
9001:1(config)#router bgp
9001:1(router-bgp)#no synchronization
```

The BGP synchronization option is set to disable so that it does not require a match for a route prefix in the route table for an IBGP path. By default, BGP synchronization is enabled. Please see section 11 for more details regarding BGP Synchronization.

3.1.1.5 Configure BGP Peers

8008 and 9001: Add BGP peers

8008:

```
8008:5(config)#router bgp
8008:5(router-bgp)#neighbor 200.1.1.1
8008:5(router-bgp)#neighbor 200.1.1.1 remote-as 40 enable
8008:5(router-bgp)#neighbor 10.1.1.2
8008:5(router-bgp)#neighbor 10.1.1.2 remote-as 20 enable
8008:5(router-bgp)#exit
```

9001:

```
9001:1(config)#router bgp
9001:1(router-bgp)#neighbor 200.1.1.42
9001:1(router-bgp)#neighbor 200.1.1.42 remote-as 40 enable
9001:1(router-bgp)#neighbor 10.1.1.1
9001:1(router-bgp)#neighbor 10.1.1.1 remote-as 20 enable
9001:1(router-bgp)#exit
```

3.1.1.6 Configure IGP Network Prefixes

Configure the BGP network prefixes that you want to distribution. The “network” command is used for this purpose. The command format is as follows:

8008 and 9001: Add the appropriate networks which you wish to advertise via BGP

8008:

```
8008:5(config)#router bgp
8008:5(router-bgp)#network 192.1.40.0/24
```

9001:

```
9001:1(config)#router bgp
9001:1(router-bgp)#network 10.60.1.0/24
```

The networks must be present in the routing table before BGP will advertise them. Please see Section 5 for more details regarding the Network command.

By default, the switch will summarize network routes based on class limits (for example, Class A, B, C networks). To disable this feature, use the following command.

- `8008:5(config)#router bgp`
- `8008:5(router-bgp)#no auto-summary`

3.1.1.7 Specifying Number of Routes Learned – Max-Prefix

The BGP implementation has a default number of routes that can be accepted per peer. For the ERS 8000 switch, the default value is 250,000 BGP forwarding routes in its routing information base (RIB) and 500,000 in its forwarding information base (FIB). For the VSP 9000 as of release 4.0 and using generation 2 modules, up to 1 million route operations is supported.

If you wish to set the number of routes, you must change the max-prefix parameter value.

The max-prefix parameter controls the maximum number of routes that a peer can accept. The purpose is to prevent configurations from accepting more routes than it can forward to. Use a setting of 0 to accept an unlimited number of prefixes.

To modify the Max prefix use the following CLI Command:

8008 & 9001:

```
8008:5(config)#router bgp
8008:5(router-bgp)#neighbor <remote peer> max-prefix ?
   <0-2147483647>  Max prefix
```

Example: to allow an unlimited number of prefixes, enter the following command assuming the BGP peer address is 150.1.0.3:

- `8008:5(router-bgp)#neighbor 150.1.0.3 max-prefix 0`

3.1.2 Verify Operations

3.1.2.1 Verify BGP Neighbor State

To verify that the BGP peers are up, use the show ip bgp summary command. Following is the output of this command:

8008:

```
8008:5#show ip bgp summary
================================================================================
                            BGP Summary - GlobalRouter
================================================================================

                                      BGP version - 4
                                         local-as - 20
                                       Identifier - 10.1.80.8
                                   Decision state - Idle
                     The total number of routes is   0

BGP NEIGHBOR INFO :

   NEIGHBOR        RMTAS      STATE     HLDTM  KPALV  HLDCFG KPCFG    WGHT   CONRTY ADVINT
--------------------------------------------------------------------------------------------
10.1.1.2           20      Established  180    60      180    60      100    120    5
200.1.1.1          40      Established  180    60      180    60      100    120    5
Total bgp neighbors: 2

BGP CONFEDERATION INFO :
confederation identifier 0
confederation peer as


BGP NETWORK INFO :

================================================================================
                            BGP Networks - GlobalRouter
================================================================================
192.1.40.0 mask 255.255.255.0 metric 0
```

Via 8008, verify the following information:

Option	Verify
Neighbor	Verify that the BGP neighbors are ***10.1.1.2*** and ***200.1.1.1***.
Rmt AS	Verify that the BGP Remote-AS for each neighbor is correct: • For neighbor ***10.1.1.2***, the local AS should be displayed as ***20***. • For neighbor ***200.1.1.1***, the remote AS should be displayed as ***40***.
State	Verify that the state is ***Established*** for each neighbor. If not, check the configuration on both switches, port state, and any possible mis-configurations in the BGP timers used.

3.1.2.2 Displaying BGP Routes

Assuming the following:

- Router-C is advertising networks 11.11.1.0/24, 172.1.1.0/24, and 172.1.2.0/24
- 9001 is advertising network 10.60.1.0/24
- Default local preference used on all switches

To show routes in the base route table, enter the following command:

8008:

```
8008:5#show ip route
================================================================================
                            IP Route - GlobalRouter
================================================================================
                                                 NH          INTER
DST              MASK             NEXT          VRF    COST FACE   PROT AGE TYPE PRF
--------------------------------------------------------------------------------
10.1.1.0         255.255.255.252 10.1.1.1         -     1     2/3   LOC  0   DB   0
10.1.80.8        255.255.255.255 10.1.80.8        -     1     0     LOC  0   DB   0
10.60.1.0        255.255.255.0   10.1.1.2       Glob~ 0     2/3   BGP  0   IB   175
11.11.1.0        255.255.255.0   200.1.1.1      Glob~ 1     2/2   BGP  0   IB   45
172.1.1.0        255.255.255.0   200.1.1.1      Glob~ 2     2/2   BGP  0   IB   45
172.1.2.0        255.255.255.0   200.1.1.1      Glob~ 2     2/2   BGP  0   IB   45
192.1.40.0       255.255.255.0   192.1.40.1     -     1     2/1   LOC  0   DB   0
200.1.1.0        255.255.255.252 200.1.1.2        -     1     2/2   LOC  0   DB   0
```

To display the full BGP route table, enter the following command:

8008:

```
8008:5#show ip bgp route
The total number of routes is 8

Network/Mask       Peer Rem Addr   NextHop Address Org Loc Pref
------------------------------------ ------------------ ---- ------------
10.60.1.0/24       10.1.1.2        10.1.1.2        IGP 100
       AS_PATH: path-is-empty
10.60.1.0/24       200.1.1.1       200.1.1.1       IGP 100
       AS_PATH: (40)
11.11.1.0/30       200.1.1.1       200.1.1.1       IGP 100
       AS_PATH: (40)
11.11.1.0/30       10.1.1.2        200.1.1.42      IGP 100
       AS_PATH: (40)
172.1.1.0/24       200.1.1.1       200.1.1.1       IGP 100
       AS_PATH: (40 50)
172.1.1.0/24       10.1.1.2        200.1.1.42      IGP 100
       AS_PATH: (40 50)
172.1.2.0/24       200.1.1.1       200.1.1.1       IGP 100
       AS_PATH: (40 50)
172.1.2.0/24       10.1.1.2        200.1.1.42      IGP 100
       AS_PATH: (40 50)
```

For 8008, verify the following information:

Option	Verify
DST Peer Rem Addr NextHop Address	Verify that networks ***10.60.1.0/24, 11.11.1.0/24, 172.1.1.0/24,*** and ***172.1.2.0/24*** are learned via BGP in the common route table. In the BGP route table, both route paths should be displayed with the appropriate NextHop address and AS Path.
PROT	Verify that the BGP routes ***10.60.1.0/24, 11.11.1.0/24, 172.1.1.0/24,*** and ***172.1.2.0/24*** are learned via ***BGP*** in the command route table.
NEXT TYPE	Verify that all routes learned from AS40 (***11.11.1.0/24, 172.1.1.0/24,*** and ***172.1.2.0 /24***) are using the best path: • ***Next = 200.1.1.1*** *and* ***TYPE = IB*** (Indirect & Best) Verify that all routes learned from within AS20 (***10.60.1.0/24***) are using the best path: • ***Next = 10.1.1.2*** *and* ***TYPE = IB***

3.1.2.3 Display BGP Routes Learned via BGP Neighbor

To show routes advertised from neighbor 200.1.1.1, use the following command:

8008:

```
8008:5#show ip bgp neighbors 200.1.1.1 routes
================================================================================
                        BGP Neighbor Routes - GlobalRouter
================================================================================
NETWORK/MASK       PEER-REM-ADDR   NEXTHOP-ADDRESS ORG LOC-PREF    STATUS

--------------------------------------------------------------------------------------------------------------
10.60.1.0/24       200.1.1.1       200.1.1.1       IGP 100         Accepted
       AS_PATH: (40)
11.11.1.0/30       200.1.1.1       200.1.1.1       IGP 100         Used
       AS_PATH: (40)
172.1.1.0/24       200.1.1.1       200.1.1.1       IGP 100         Used
       AS_PATH: (40 50)
172.1.2.0/24       200.1.1.1       200.1.1.1       IGP 100         Used
       AS_PATH: (40 50)

8008:5#show ip bgp neighbors 10.1.1.2 routes
================================================================================
                        BGP Neighbor Routes - GlobalRouter
================================================================================
NETWORK/MASK       PEER-REM-ADDR   NEXTHOP-ADDRESS ORG LOC-PREF    STATUS
--------------------------------------------------------------------------------------------------------------
The total number of accepted routes from the neighbor is 4

Network/Mask       Peer Rem Addr   NextHop Address Org Loc Pref    Status
------------------------------------ ------------------ ---- ----------    -----------
10.60.1.0/24       10.1.1.2        10.1.1.2        IGP 100         Used
       AS_PATH: path-is-empty
11.11.1.0/30       10.1.1.2        200.1.1.42      IGP 100         Accepted
       AS_PATH: (40)
172.1.1.0/24       10.1.1.2        200.1.1.42      IGP 100         Accepted
       AS_PATH: (40 50)
172.1.2.0/24       10.1.1.2        200.1.1.42      IGP 100         Accepted
       AS_PATH: (40 50)
```

Overview of the information displayed:

Option	Verify
Network/Mask	Displays the IP network and mask for the direct route. The best route should have it status displayed as Used as follows: • For BGP peer ***10.1.1.2***, network ***10.60.1.0/24*** should be used • For BGP peer ***200.1.1.1***, networks ***11.11.1.0/24***, ***172.1.1.0/24***,and ***172.1.2.0/24*** should be used.
Peer Rem Addr	Displays, the peer remote address.
NextHop Address	Displays the next-hop address IP address.
Org	Well-known mandatory attribute that specifies the source of a route: • IGP — the route is interior to the originating AS that inserts this route into the BGP table (0 = IGP). • EGP — the route is learned via the Exterior Gateway Protocol (EGP) prior to being inserted into the BGP table (1 = BGP). • Incomplete — the origin of the route is unknown or learned by some other means. For example, these routes could be learned through RIP, OSPF, or static routes (2 = Incomplete).
Local Pref	Displays the local preference attribute.
Status	Displays the route status which will be either Accepted, Best, Used, or Rejected. For this example: • For BGP peer ***10.1.1.2***, network ***10.60.1.0/24*** should be displayed as ***Used*** • For BGP peer ***200.1.1.1***, networks ***11.11.1.0/24***, ***172.1.1.0/24***,and ***172.1.2.0/24*** should be ***Used***.

3.1.2.4 Verify BGP Networks

To display the networks configured, enter the following command:

```
8008:
8008:5#show ip bgp networks

================================================================================
                          BGP Networks - GlobalRouter
================================================================================
192.1.40.0 mask 255.255.255.0 metric 0
```

3.1.2.5 View the BGP Routes Sent out to a Specific Peer

To show routes advertised to a specific peer, in this case, 200.1.1.1 from switch 8008, enter the following command:

```
8008:
8008:5#show ip bgp neighbors 200.1.1.1 advertised-routes

================================================================================
                BGP Neighbor Advertised Routes - GlobalRouter
================================================================================

The total number of routes advertised to the neighbor is 1
NETWORK/MASK       NEXTHOP ADDRESS LOC PREF ORG    STATUS
--------------------------------------------------------------------------------
10.60.1.0/24       10.1.1.2        100      IGP    Used

8008:5#show ip bgp neighbors 10.1.1.2 advertised-routes

================================================================================
                BGP Neighbor Advertised Routes - GlobalRouter
================================================================================

The total number of routes advertised to the neighbor is 3
NETWORK/MASK       NEXTHOP ADDRESS LOC PREF ORG    STATUS
--------------------------------------------------------------------------------
11.11.1.0/24       200.1.1.1       100      IGP    Used
172.1.1.0/24       200.1.1.1       100      IGP    Used
172.1.2.0/24       200.1.1.1       100      IGP    Used
```

Via 8008, verify the following information:

Option	Verify
Network/Mask	Verify that only networks ***192.1.40.0/24,*** and ***10.60.1.0/24*** are advertised to BGP neighbor ***200.1.1.1***. Verify that all networks are advertised to BGP neighbor ***10.1.1.2***.
Status	Verify that the network ***192.1.40.0/24*** is set to ***import*** to indicate a local interface advertised via BGP through the network command. Verified that network ***10.60.1.0/24*** is set to ***Used*** from peer ***200.1.1.1*** and set to ***Accepted*** from peer ***10.1.1.2***. The networks ***11.11.1.0/24, 172.1.1.0/24, and 172.1.2.0/24*** should NOT be advertised back to 200.1.1.1. Verify that networks ***11.11.1.0/24, 172.1.1.0/24, and 172.1.2.0/24*** are set to ***Used*** from peer ***10.1.1.2***.

4. BGP Timers

Every BGP router maintains a KeepAlive and Hold Timer for each BGP session it possesses. These timers are used for peer health check. When the KeepAlive Timer expires, a KEEPALIVE message is sent to the peer router associated with the session. When receiving a KEEPALIVE message or an UPDATE message, the Hold Timer is cleared. When an UPDATE message is sent out, the KeepAlive Timer is also cleared. If the Hold Timer expires, the BGP router assumes that the peer router cannot respond correctly, and thus resets the BGP session.

The following table displays the various timer options available on the Extreme Ethernet Routing Switch 8000. Please see Appendix C – BGP Events regarding details on BGP events and in reference to the timers below.

Table 2: BGP Timers

Parameter	Description
Connect Retry Interval (Sec)	Amount of time in seconds to wait before attempting to reconnect to a BGP neighbor after failing to connect. Router falls back to Connect State after timer expires. • Range 1 to 65535 seconds; default 120.
KeepAlive	Message sent to keep BGP connection alive to ensure Hold Timer does not expire when no Update messages are sent. If the value is zero, no periodical keepalive messages are sent to this neighbor after the BGP connection has been established. • Range 0 to 21845 seconds; default 60
Advertisement Interval (Sec)	Specifies the time interval that transpires in seconds between each transmission of a router advertisement from a BGP neighbor • Range: 1 to 120 seconds; default 5
Hold Timer (Sec)	The hold time is the maximum time allowed between receipt of successive KeepAlive, and/or Update messages. This hold time is reset and counts down upon a successful receipt of a message. The hold time must be either 0 or at least 3 seconds and is 3 x KeepAlive value. • Range 0-65535; default 180

Changing the default timers is performed at the BGP neighbor level using the following commands.

Go to BGP configuration

```
8008:5(config)#router bgp
8008:5(router-bgp)#
```

To change the Keepalive and Holddown Timer:

```
8008:5(router-bgp)#neighbor <neighbor> timers <0-21845> <0-65535>
```

To change the Connect Retry Interval:

```
8008:5(router-bgp)#neighbor <neighbor> retry-interval <1-65535>
```

To change the Advertisement Interval:

```
8008:5(router-bgp)#neighbor <neighbor> advertisement-interval <5-120>
```

The Hold Time is negotiated between peers during session establishment. The smaller value is used. The keepalive is not negotiated and is used at the set value unless the hold timer negotiated is less than the keepalive. Then the keepalive will be 1/3 the hold timer.

5. BGP Network Command

The Extreme Ethernet Routing Switch and Virtual Services Platform uses the *Network* command to specify a list of IGP networks that are advertised as originating from an autonomous system.

To change the Keepalive and Holddown Timer:

```
8008:5(config)#router bgp
8008:5(router-bgp)#network <prefix/len>
```

The prefix/len that is specified must match an active entry in the IP routing table. The route may be local to the switch, configured as a static route, or dynamically learned via an IGP such as RIP or OSPF. The network command cannot be used to aggregate or summarize BGP routes.

When the prefix originated by the *Network* command is advertised via BGP, its Route Origin attribute is set to “IGP”. This indicates that the route is interior to the originating AS.

For example, via 8008 from Section 3.1, if we wish to originate the CLIP address 10.1.80.8./32, enter the following command:

```
8008:5(config)#router bgp
8008:5(router-bgp)#network 10.1.80.8/32
```

On Router C the BGP route table indicates that the network has been learned from 8008. Note that the Route Origin is “IGP”:

```
RouterC:5#show ip bgp route 10.1.80.8/32

================================================================================
                            BGP Routes - GlobalRouter
================================================================================
The total number of routes is 4

NETWORK/MASK       PEER REM ADDR   NEXTHOP ADDRESS ORG LOC PREF
--------------------------------------------------------------------------------------------------------------
10.1.80.8/32        200.1.1.2       200.1.1.2       IGP 100
  AS_PATH: (20)
```

6. Redistribution Policies

Within an AS, BGP update information is distributed between BGP speakers using an Interior Gateway Protocol (IGP) that runs within the AS. The Extreme Ethernet Routing Switch and Virtual Services Platform supports either RIP, ISIS (SPBM), or OSPF for IGP. In regards to SPBM, please see section 6 below.

Under normal operation, the IGP carries no BGP information. Each BGP speaker in an AS uses IBGP exclusively to determine reachability to external networks. In order to inject routes into the IGP, redistribution policies must be created to inject external routes within an AS.

This section provides examples of the commands you use to create redistribution policies that can inject external routes within an AS.

If the AS is running OSPF, the border router must be configured as an AS boundary router (ASBR) in order to accept external routes.

6.1 BGP Redistribution

The ERS or VSP can redistribute routes learned by RIP, ISIS, OSPF, or static route configuration. In addition, it can also redistribute local or direct interfaces. The following command is used to configure BGP route distribution:

The following command is used to configure BGP route distribution:

```
8008:1(config)#router bgp
8008:1(router-bgp)#redistribute ?
  direct Ip bgp redistribute direct command
  ipv6-direct Ip bgp redistribute ipv6-direct command
  ipv6-static Ip bgp redistribute ipv6-static command
  isis Ip bgp redistribute isis command
  ospf         Ip bgp redistribute ospf command
  ospfv3       Ip bgp redistribute ospfv3 command
  rip          Ip bgp redistribute rip command
  static       Ip bgp redistribute static command
```

For example, to redistribute direct interfaces, enter the commands shown below.

```
8008:5(config)#router bgp
8008:5(router-bgp)#redistribute direct
WARNING: Routes not inject until apply command is issued after enable command
8008:5(router-bgp)#redistribute direct enable
8008:5(router-bgp)#exit
8008:5(config)#ip bgp apply redistribute direct
```

Note that when the routes are imported into the BGP route table, a route-policy may be applied in order to suppress specific routes or modify BGP route attributes. For example, this may be useful if you

redistribute all “direct” (i.e. locally connected) routes but do not wish to advertise the IP prefixes of certain interfaces.

Also, the BGP “metric” attribute associated with each prefix, also known as MED, may also be set. Note that if the metric is also set via a route-policy, the route-policy specified metric takes precedence.

When the prefix originated by the “redistribute” command is advertised via BGP its Route Origin attribute is set to “INC”, or incomplete. When BGP selects the best path to a given destination a route with origin “IGP” takes priority over a route with origin “INC”. The following is an example of routes imported on an ERS 8000 after BGP direct redistribute has been enable on a peer ERS 8000. Notice the Route Origin is “INC”.

```
RouterC:5#show ip bgp neighbors 200.1.1.2 routes longer-prefixes 10.1.1.0/30
The total number of accepted  routes from the neighbor is 6

================================================================================
                        BGP Neighbor Routes - GlobalRouter
================================================================================
NETWORK/MASK       PEER-REM-ADDR   NEXTHOP-ADDRESS ORG LOC-PREF    STATUS

--------------------------------------------------------------------------------
10.1.1.0/30        200.1.1.2       200.1.1.2       INC 100         Used
        AS_PATH : (20)
```

6.2 OSPF and BGP Route Distribution

This section describes commands you use to create OSPF and BGP route distribution. The commands used are in reference to Figure 1 used in Section 3.1 above using 8008.

6.2.1 Configuration

To create OSPF and BGP route distribution policies, complete the following steps:

6.2.1.1 Configure OSPF on 8008

Configure 8008 as an OSPF ASBR and enable OSPF

```
8008:1(config)#router ospf
8008:5(config-ospf)#as-boundary-router enable
8008:5(config-ospf)#exit
8008:5(config)#router ospf enable
```

The ERS 8000 must be configured as an OSPF Autonomous System Border Router (ASBR) in order to support other routing protocols other than OSPF.

6.2.1.2 Configure Route Policy to Redistribute BGP Routes into OSPF

Enable BGP into OSPF redistribution

```
8008:5(config)#router ospf
8008:5(config-ospf)#redistribute bgp
8008:5(config-ospf)#redistribute bgp enable
8008:5(config-ospf)#exit
8008:5(config)#ip ospf apply redistribute bgp
```

6.2.1.3 Configure Route Policy Redistribute OSPF Routes into BGP

Enable OSPF into BGP redistribution

```
8008:5(config)#router bgp
8008:5(config-ospf)#redistribute ospf
8008:5(config-ospf)#redistribute ospf enable
8008:5(config-ospf)#exit
8008:5(config)#ip bgp apply redistribute ospf
```

Be very careful enabling BGP redistribution. It could cause learned eBGP routes to be advertised out of your local AS. This would have the effect of other networks routing through the local AS. It is best not to enable this feature if you are peering to an ISP and do not wish to have traffic transit the local AS.

When redistributing OSPF into BGP, route priority will be in effect and you can create routing loops. BGP has a higher route preference than OSPF External 1 & 2. Thus if you redistribute OSPF external 1 & 2 routes into BGP then BGP routes will be used and this could cause a routing loop.

The BGP Router ID automatically uses the OSPF Router ID. If the OSPF Router ID is changed then BGP must be restarted to use the new value. Note that OSPF uses a random Router ID by default. The commands to disable and enable BGP globally are:

```
8008:5(config)#no router bgp enable

8008:5(config)#router bgp <local as> enable
```

6.3 Creating a Policy to Inject Default Route When Using OSPF as an Interior Gateway Protocol

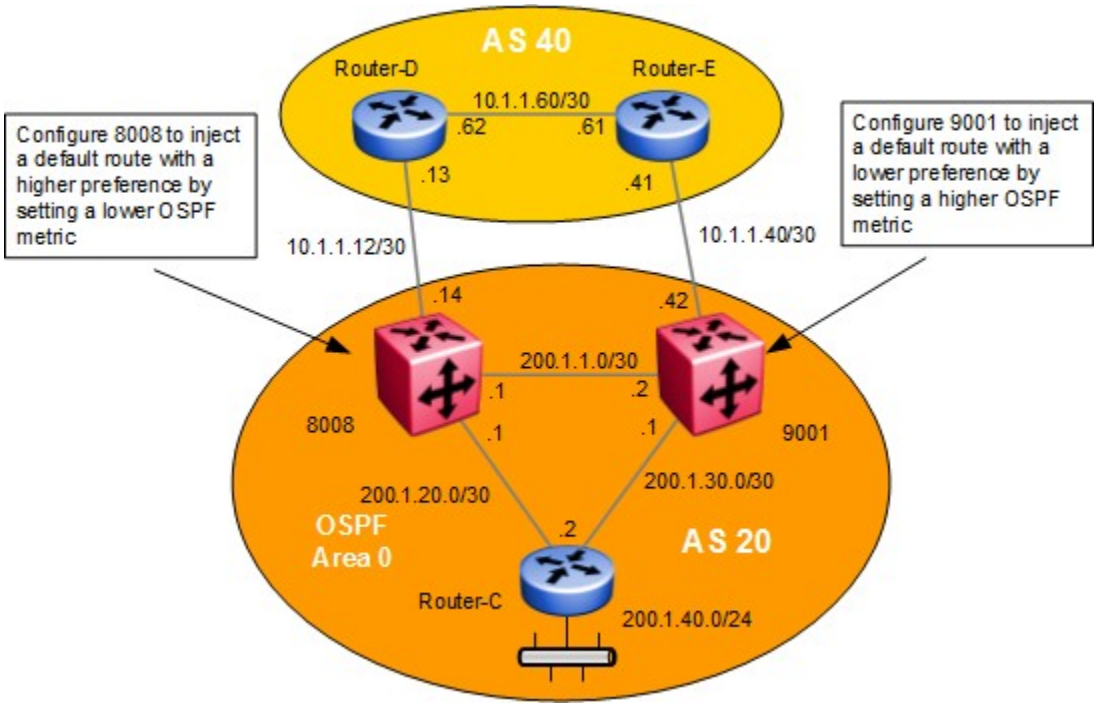

Figure 2: Inject Default Route Configuration Example

In this example, we are going to configure both 8008 and 9001 to inject a default route into IGP. In this example we are using OSPF as the IGP protocol. We can also influence the path of the default route with route metrics so that Router-C can use either 200.1.20.1 or 200.1.30.1 as the next hop. In the following configuration example, we are going to configure the network so that 200.1.20.1 is the default next hop by manipulating the OSPF route metric. Note that by configuring the default route as used in this example, 8008 will always be used at the default route gateway for all outbound traffic outside AS 20 unless of course it should fail.

6.3.1 Configuration

6.3.1.1 Configure the IP Prefix List:

Add IP Prefix with default route using a prefix name of DR

8008 & 9001: Same configuration on both switches

```
8008:5(config)#ip prefix-list DR 0.0.0.0/0
```

6.3.1.2 Configure the IP Route Policy

Add route policy named Default_OSPF using sequence 1, add the IP prefix named DR, set the metric to 100 on 8008, and set the metric to 300 on 9001

8008:

```
8008:5(config)#route-map Default_OSPF 1
8008:5(route-map)#enable
8008:5(route-map)#set injectlist DR
8008:5(route-map)#set metric 100
8008:5(route-map)#exit
```

9001:

```
9001:1(config)#route-map Default_OSPF 1
9001:1(route-map)#permit
9001:1(route-map)#enable
9001:1(route-map)#set injectlist DR
9001:1(route-map)#set metric 300
9001:1(route-map)#exit
```

The policy set-metric value is what will influence the OSPF route decision. The lower the value the higher the route preference. For this example, 8008 is set to a lower metric value than 9001, which results in a higher preference value.

6.3.1.3 Configure Route Redistribution:

Enable BGP redistribution into OSPF

8008 & 9001: Same configuration on both switches

```
8008:5(config)#router ospf
8008:5(config-ospf)#redistribute bgp
WARNING: Routes not inject until apply command is issued after enable command
8008:5(config-ospf)#redistribute bgp route-policy Default_OSPF
8008:5(config-ospf)#redistribute bgp enable
8008:5(config-ospf)#exit
8008:5(config)#ip ospf apply redistribute bgp
```

The end result of this configuration is that Router-C will use the next hop to 8008 for access to the Internet.

7. ISIS and BGP Route Distribution

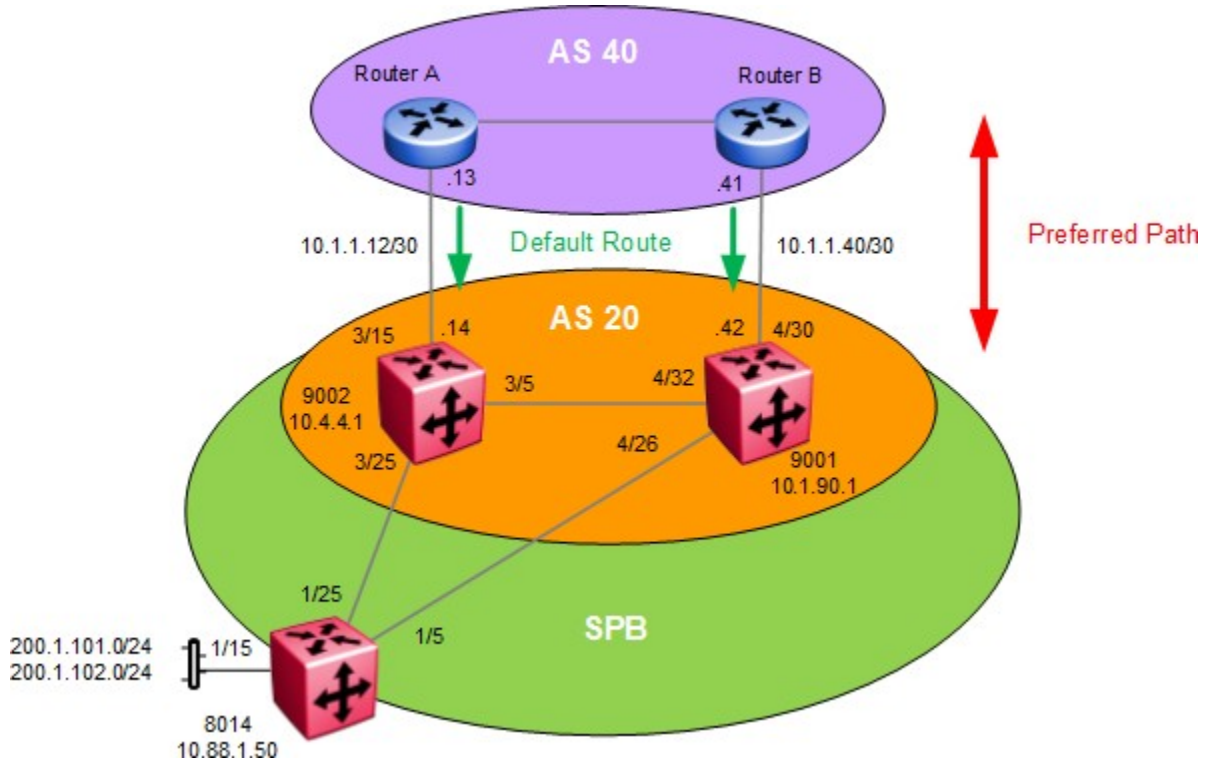

Figure 3: ISIS and BGP Route Distribution

In this example, we are using ISIS (Shortest Path Bridging) as the IGP. The AS40 routers, Router A and Router B, are setup to send a default route to AS20. We are going to configure VSP 9000 switches 9001 and 9002 to accept and inject the default route from AS40 into ISIS. We will also inject ISIS into BGP by summarizing the local 200.1.x.y networks. As an option, we will also influence the path of the default route with BGP Local Preference and ISIS metrics so that ISIS and iBPG will choose 9001 as the preferred switch for the default route.

We will also demonstrate how to use prefix lists and route-maps to selectively only advertise specific networks to the eBGP peer routers in AS40 and reject local networks advertised via BGP between the iBGP peers 9001 and 9002.

Please note that release 4.0 or higher is required on the VSP 9000 to support this configuration example using ISIS accept policy.

7.1 Configuration

7.1.1 ISIS Configuration

7.1.1.1 SPBM Configuration

SPBM base configuration

9001:

```
9001:1(config)#spbm
9001:1(config)#prompt 9001
9001:1(config)#router isis
9001:1(config-isis)#sys-name 9001
9001:1(config-isis)#manual-area 49.0001
9001:1(config-isis)#spbm 1
9001:1(config-isis)#spbm 1 nick-name 0.90.01
9001:1(config-isis)#spbm 1 b-vid 3051-3052 primary 3051
9001:1(config-isis)#exit
9001:1(config)#vlan create 3051 type spbm-bvlan
9001:1(config)#vlan create 3052 type spbm-bvlan
9001:1(config)#router isis enable
9001:1(config)#cfm cmac mepid 901
9001:1(config)#cfm cmac enable
```

9002: Same configuration as 9001 except for the following

```
9002:1(config)# prompt 9002
9002:1(config)#router isis
9002:1(config-isis)#sys-name 9002
9002:1(config-isis)#spbm 1 nick-name 0.90.02
9002:1(config-isis)#exit
9002:1(config)#cfm cmac mepid 902
```

ISIS interface configuration

9001:

```
9001:1(config)#interface GigabitEthernet 4/26,4/32
9001:1(config-if)#no shutdown
9001:1(config-if)#isis
9001:1(config-if)#isis spbm 1
9001:1(config-if)#isis enable
9001:1(config-if)#no spanning-tree mstp force-port-state enable
```

9002: Same configuration as 9001 except for port numbers

```
9002:1(config)#interface GigabitEthernet 3/5,3/25
```

Verify operations

9001:

At this point, ISIS should be operational and ISIS adjacencies should be up

```
9001:1#show isis
================================================================================
                              ISIS General Info
================================================================================
                           AdminState : enabled
                           RouterType : Level 1
                            System ID : d4ea.0efd.e000
                  Max LSP Gen Interval : 900
                               Metric : wide
                 Overload-on-startup : 20
                            Overload : false
                        Csnp Interval : 10
                        PSNP Interval : 2
                   Rxmt LSP Interval : 5
                           spf-delay : 100
                         Router Name : 9001
                   ip source-address : 10.1.90.1
                   Num of Interfaces : 2
                Num of Area Addresses : 1

9001:1#show isis inter
================================================================================
                              ISIS Interfaces
================================================================================
IFIDX      TYPE    LEVEL     OP-STATE  ADM-STATE  ADJ    UP-ADJ  SPBM-L1-METRIC
--------------------------------------------------------------------------------
Port4/26   pt-pt   Level 1   UP        UP         1      1       10
Port4/32   pt-pt   Level 1   UP        UP         1      1       10

9001:1#show isis adj

================================================================================
                              ISIS Adjacencies
================================================================================
INTERFACE L STATE        UPTIME PRI HOLDTIME SYSID            HOST-NAME
--------------------------------------------------------------------------------
Port4/26  1 UP         00:08:20 127       23 0049.0080.1400    8014
Port4/32  1 UP         00:08:20 127       26 d4ea.0efd.a000    9002
```

7.1.1.2 Enable IP Shortcuts

Base ISIS Configuration

9001:

```
9001:1(config)#interface loopback 1
9001:1(config-if)#ip address 1 10.1.90.1/255.255.255.255
9001:1(config-if)#exit
9001:1(config)#router isis
9001:1(config-isis)#ip-source-address 10.1.90.1
9001:1(config-isis)#spbm 1 ip enable
```

9002:

```
9002:1(config)#interface loopback 1
9002:1(config-if)#ip address 1 10.4.4.1/255.255.255.255
9002:1(config-if)#exit
9002:1(config)#router isis
9002:1(config-isis)#ip-source-address 10.4.4.1
9002:1(config-isis)#spbm 1 ip enable
```

Please note ISIS direct interface redistribution is not enabled for this example. This step is only required, for example, if there are local interfaces you wish to advertise to peer SPBM enabled switches. In this example, we only have the one eBGP IP interface on both 9001 and 9002 that does not have to be advertised within ISIS.

7.1.2 Prefix list and route policies

For this example, we will configure switches 9001 and 9001 to only advertise the network range 200.1.x.y range to AS40. In order to do this, we will create a prefix list and route-map to match this range. We will also create a second route map to match this this range used to deny this IP range from every being learned by iBGP between 9001 and 9001.

7.1.2.1 Add route map to summarize local networks

Add prefix list using a name of *local-sub* use for the two route-maps we will create

9001 * 9002: Same configuration on both switches

```
9001:1(config)#ip prefix-list local-sub 200.1.0.0/16 ge 16 le 24
```

Add route-map using a name of *local-nets* used to advertise local network to AS40

9001 * 9002: Same configuration on both switches

```
9001:1(config)#route-map local-nets 1
9001:1(route-map)#permit
9001:1(route-map)#enable
9001:1(route-map)#match network local-sub
9001:1(route-map)#exit
9001:1(config)#route-map local-nets 2
9001:1(route-map)#no permit
9001:1(route-map)#enable
```

Add route-map using a name of *local-deny* to deny local networks between iBGP peers 9001 & 9002

9001 & 9002: Same configuration on both switches

```
9001:1(config)#route-map local-deny 1
9001:1(route-map)#no permit
9001:1(route-map)#enable
9001:1(route-map)#match network local-sub
9001:1(route-map)#exit
9001:1(config)#route-map local-deny 2
9001:1(route-map)#permit
9001:1(route-map)#enable
```

7.1.3 BGP Configuration

7.1.3.1 BGP Base Configuration

Add IP interface – in this example we will use brouter ports

9001:

```
9001:1(config)#interface GigabitEthernet 4/30
9001:1(config-if)#no shutdown
9001:1(config-if)#brouter port 4/30 vlan 2090 subnet 10.1.1.42/30
9001:1(config-if)#auto-negotiate enable
9001:1(config-if)#no spanning-tree mstp force-port-state enable
```

9002:

```
9002:1(config)#interface GigabitEthernet 3/15
9002:1(config-if)#no shutdown
9002:1(config-if)#brouter port 3/15 vlan 2090 subnet 10.1.1.14/30
9002:1(config-if)#no spanning-tree mstp force-port-state enable
```

Depending on the interface type, i.e. SFP+ or SFP for 1 gigabit Ethernet, for SFP interfaces you will need to enable auto negotiation at the interface level – auto-negotiate enable.

Base BGP Configuration – for this example, we will also increase the BGP local preference to 200 on 9001 so that the eBGP 9001 path will be preferred over the eBGP path via 9002 which retains the default local preference of 100. Also, we will add the route-map *local-deny* between the iBGP peering on 9001 and 9002

9001:

```
9001:1(config)#router bgp
9001:1(router-bgp)#no auto-summary
9001:1(router-bgp)#no synchronization
9001:1(router-bgp)#router-id 10.1.90.1
BGP router-id is different from OSPF router-id.
Are you sure you want to continue? (y/n) ? y
9001:1(router-bgp)#bgp default local-preference 200
9001:1(router-bgp)#neighbor 10.1.1.41
9001:1(router-bgp)#neighbor 10.1.1.41 remote-as 40
9001:1(router-bgp)#neighbor 10.1.1.41 enable
9001:1(router-bgp)#neighbor 10.4.4.1
9001:1(router-bgp)#neighbor 10.4.4.1 remote-as 20
9001:1(router-bgp)#neighbor 10.4.4.1 update-source 10.1.90.1
9001:1(router-bgp)#neighbor 10.4.4.1 in-route-map local-deny
```

```
9001:1(router-bgp)#neighbor 10.4.4.1 next-hop-self enable
9001:1(router-bgp)#neighbor 10.4.4.1 enable
9001:1(router-bgp)#exit
9001:1(config)#router bgp 20 enable
9002:
9002:1(config)#router bgp
9002:1(router-bgp)#no auto-summary
9002:1(router-bgp)#no synchronization
9002:1(router-bgp)#router-id 10.4.4.1
BGP router-id is different from OSPF router-id.
Are you sure you want to continue? (y/n) ? y
9002:1(router-bgp)#neighbor 10.1.1.13
9002:1(router-bgp)#neighbor 10.1.1.13 remote-as 40
9002:1(router-bgp)#neighbor 10.1.1.13 enable
9002:1(router-bgp)#neighbor 10.1.90.1
9002:1(router-bgp)#neighbor 10.1.90.1 remote-as 20
9002:1(router-bgp)#neighbor 10.1.90.1 update-source 10.4.4.1
9002:1(router-bgp)#neighbor 10.1.90.1 in-route-map local-deny
9002:1(router-bgp)#neighbor 10.1.90.1 next-hop-self enable
9002:1(router-bgp)#neighbor 10.1.90.1 enable
9002:1(router-bgp)#exit
9002:1(config)#router bgp 20 enable
```

Please note for the iBGP peer configuration between 9001 and 9002, we enabled the next-hop-self parameter. This step is only required if ISIS redistribution of direct interfaces is not enabled as is the case in our example. Hence, the eBGP local interface, for example, from 9001 is not advertised to 9002 and vise-versa. Please see section 10.2 for more details regarding BGP next-hop-self.

7.1.3.2 Verify Operations

At this point, BGP should be up and running with two adjacencies

```
9001:1(config)#show ip bgp summary

================================================================================
                            BGP Summary - GlobalRouter
================================================================================

                                      BGP version - 4
                                         local-as - 20
                                       Identifier - 10.1.90.1
```

```
                                Decision state - Idle
                        The total number of routes is    2


BGP NEIGHBOR INFO :

   NEIGHBOR        RMTAS      STATE     HLDTM  KPALV  HLDCFG KPCFG    WGHT   CONRTY ADVINT
-------------------------------------------------------------------------------------------------------
10.1.1.41          40           Established 180    60     180     60     100    120     5
10.4.4.1           20           Established 180    60     180     60     100    120     5
```

7.1.4 Enable ISIS Route Policy and Disable Alternative Route

7.1.4.1 Enable ISIS Accept Policy

An ISIS policy is required to stop routing loops between ISIS and BGP. The accept policy as shown below ensures that the only ISIS route that 9001 will accept from 9002 is for its loopback interface and vise-versa.

Set ISIS Route Policy allowing only the local networks from each peer using the remote switch SPBM nick-name

9001:

```
9001:1(config)#ip prefix-list local-nets 10.4.4.1/32
9001:1(config)#route-map accept-local 1
9001:1(route-map)#permit
9001:1(route-map)#enable
9001:1(route-map)#match network local-nets
9001:1(route-map)#exit
9001:1(route-map)#route-map accept-local 2
9001:1(route-map)#no permit
9001:1(route-map)#enable
9001:1(route-map)#exit
9001:1(route-map)#router isis
9001:1(config-isis)#accept adv-rtr 0.90.02 enable
9001:1(config-isis)#accept adv-rtr 0.90.02 route-map accept-local
9001:1(config-isis)#exit
9001:1(config)#isis apply accept
```

9002:

```
9002:1(config)#ip prefix-list local-nets 10.1.90.1/32
```

```
9002:1(config)#route-map accept-local 1
9002:1(route-map)#permit
9002:1(route-map)#enable
9002:1(route-map)#match network local-nets
9002:1(route-map)#exit
9002:1(route-map)#route-map accept-local 2
9002:1(route-map)#no permit
9002:1(route-map)#enable
9002:1(route-map)#exit
9002:1(route-map)#router isis
9002:1(config-isis)#accept adv-rtr 0.90.01 enable
9002:1(config-isis)#accept adv-rtr 0.90.01 route-map accept-local
9002:1(config-isis)#exit
9002:1(config)#isis apply accept
```

7.1.4.2 Disable Alternative Route

Alternative routing should always be disabled on the boundary routers.

Disable Alternative Route

9001 & 9002: Same configuration on both switches

```
9001:1(config)#no ip alternative-route
```

7.1.5 Enable BGP to ISIS Redistribution

7.1.5.1 Enable BGP to ISIS Redistribution

Method 1: Redistribute BGP into ISIS – we will also increase the metric on 9002 from a default setting of 100 to 300 to make 9001 more preferred

9001:

```
9001:1(config)#router isis
9001:1(config-isis)#redistribute bgp
9001:1(config-isis)#redistribute bgp enable
9001:1(config-isis)#exit
9001:1(config)#isis apply redistribute bgp
```

9002:

```
9002:1(config)#router isis
9002:1(config-isis)#redistribute bgp
9002:1(config-isis)#redistribute bgp metric 300
9002:1(config-isis)#redistribute bgp enable
9002:1(config-isis)#exit
```

```
9002:1(config)#isis apply redistribute bgp
```

By changing the ISIS metric on 9002 from 100 to 300 we are changing the “external metric” of the ISIS route and the expectation is that switch 9001 should now become the preferred switch for the default route as it will have the lower default metric of 100. However with SPB ISIS IP routes, it is the SPB shortest path to the advertising BEB (“internal” SPB metric) which is more important and the “external” metric is only used as a tie breaker in case the “internal” metric is the same. So the above approach will only work in this example as the ISIS SPB path from 8014 to 9001 and 8014 to 9002 is the same making the metric the tie breaker. However, in networks where the BEB switches do not have equal cost to the border routers, a better way to ensure that ISIS only considers the route to the primary eBGP border router is to make only that border router redistribute that route into ISIS. Since the primary eBGP router holds the route as eBGP and the secondary eBGP router (which has a lower BGP local preference) holds the same route as iBGP, this can be easily done by using a route policy which redistributes only eBGP routes into ISIS and not iBGP routes. BGP local preference will now become the tie breaker. In our example, switch 9001 is set for a higher BGP local preference so it will redistribute the default route into ISIS using the configuration as shown below.

Method 2: Redistribute BGP into ISIS – this time we will create a route-map to match only eBGP so that the BGP local preference will decide who will become the preferred router

9001:

```
9001:1(config)#route-map ebgp_only 1
9001:1(route-map)#permit
9001:1(route-map)#match protocol ebgp
9001:1(route-map)#enable
9001:1(route-map)#exit
9001:1(config)#router isis
9001:1(config-isis)#redistribute bgp
9001:1(config-isis)#redistribute bgp route-map ebgp_only
9001:1(config-isis)#redistribute bgp enable
9001:1(config-isis)#exit
9001:1(config-isis)#isis apply redistribute bgp
```

9002:

```
9002:1(config)#route-map ebgp_only 1
9002:1(route-map)#permit
9002:1(route-map)#match protocol ebgp
9002:1(route-map)#enable
9002:1(route-map)#exit
9002:1(config)#router isis
9002:1(config-isis)#redistribute bgp
9002:1(config-isis)#redistribute bgp route-map ebgp_only
9002:1(config-isis)#redistribute bgp enable
9002:1(config-isis)#exit
9002:1(config)#isis apply redistribute bgp
```

7.1.5.2 Verify Operations

At this point, switch 8014 should point to switch 9001 for the default route

```
8014:5#show ip route

================================================================================================
                                  IP Route - GlobalRouter
================================================================================================
                                                 NH                INTER
DST             MASK            NEXT            VRF          COST FACE PROT AGE TYPE PRF
------------------------------------------------------------------------------------------------
0.0.0.0         0.0.0.0         9001            GlobalRout~ 10   3051 ISIS 0    IBS  7
```

We can also look at the advertised metric on switch 8014 by looking TLV 135 (IP Reachability) using the Method 1 configuration in step 6.1.5

```
8014:5#show isis lsdb tlv 135 detail
================================================================================
                                ISIS LSDB (DETAIL)
================================================================================
--------------------------------------------------------------------------------------------------------------

Level-1 LspID: 0049.0080.1400.00-00     SeqNum: 0x00000c80      Lifetime:  1072
        Chksum: 0x8922  PDU Length: 151
        Host_name: 8014
        Attributes:     IS-Type 1
TLV:135 TE IP Reachability: 3
                Metric: 1       Prefix Length: 32
                UP/Down Bit: FALSE              Sub TLV Bit: FALSE
                IP Address: 10.88.1.50
                Metric: 1       Prefix Length: 24
                UP/Down Bit: FALSE              Sub TLV Bit: FALSE
                IP Address: 200.1.101.0
                Metric: 1       Prefix Length: 24
                UP/Down Bit: FALSE              Sub TLV Bit: FALSE
                IP Address: 200.1.102.0



Level-1 LspID: d4ea.0e10.e465.00-02     SeqNum: 0x0000000e      Lifetime:   877
        Chksum: 0xd572  PDU Length: 45
        Host_name: 9002
        Attributes:     IS-Type 1
TLV:135 TE IP Reachability: 2
                Metric: 1       Prefix Length: 32
                UP/Down Bit: FALSE              Sub TLV Bit: FALSE
                IP Address: 10.4.4.1

                Metric: 300     Prefix Length: 0
                UP/Down Bit: FALSE              Sub TLV Bit: FALSE
                IP Address: 0.0.0.0
```

```
Level-1 LspID: d4ea.0efd.e3df.00-02     SeqNum: 0x0000000b      Lifetime:    882
        Chksum: 0x4379 PDU Length: 45
        Host_name: 9001
        Attributes:     IS-Type 1
TLV:135 TE IP Reachability: 2
                Metric: 1       Prefix Length: 32
                UP/Down Bit: FALSE              Sub TLV Bit: FALSE
                IP Address: 10.1.90.1

                Metric: 1       Prefix Length: 0
                UP/Down Bit: FALSE              Sub TLV Bit: FALSE
                IP Address: 0.0.0.0
```

We can also verify the metric by IP unicast FIB

```
9001:1(config)#show isis lsdb ip-unicast

================================================================================
                          ISIS IP-UNICAST-ROUTE SUMMARY
================================================================================
                             PREFIX          TLV   LSP   HOST
I-SID     ADDRESS            LENGTH  METRIC  TYPE  FRAG  NAME
--------------------------------------------------------------------------------
-         10.88.1.50         32      1       135   0x0   8014
-         200.1.101.0        24      1       135   0x0   8014
-         200.1.102.0        24      1       135   0x0   8014
-         10.4.4.1           32      1       135   0x2   9002
-         0.0.0.0            0       300     135   0x2   9002
-         10.1.90.1          32      1       135   0x2   9001
-         0.0.0.0            0       1       135   0x2   9001
--------------------------------------------------------------------------------
7 out of 7 Total Num of Entries
```

7.1.6 Enable ISIS to BGP Redistribution

For this example, we only wish to send the local networks 200.1.101.0/24 and 200.1.102.0/24 to the eBGP peers in AS40. To do this, we will create two route-maps, one to advertise the local networks via eBGP to the AS40 peers and another to reject the same local networks via iBGP between 9001 and 9002 which we configured in the BGP Configuration step above.

7.1.6.1 Redistribute ISIS into BGP

Enable ISIS to BGP redistribution and add the route-map named *local-nets* create previously to only advertise the 200.1.x.y networks

9001:

```
9001:1(config)#router bgp
9001:1(router-bgp)#redistribute isis
9001:1(router-bgp)#redistribute isis route-map local-nets
9001:1(router-bgp)#redistribute isis enable
9001:1(router-bgp)#exit
9001:1(config)#ip bgp apply redistribute isis
```

9002:

```
9002:1(config)#router bgp
9002:1(router-bgp)#redistribute isis
9002:1(router-bgp)#redistribute isis route-map local-nets
9002:1(router-bgp)#redistribute isis enable
9002:1(router-bgp)#exit
9002:1(config)#ip bgp apply redistribute isis
```

7.1.6.2 Verify Operations

Verify BGP routes received from each peer

```
9002:1#show ip bgp neighbors 10.1.90.1 route
The total number of accepted  routes from the neighbor is 1


================================================================================
                        BGP Neighbor Routes - GlobalRouter
================================================================================
NETWORK/MASK        PEER-REM-ADDR   NEXTHOP-ADDRESS ORG LOC-PREF    STATUS

--------------------------------------------------------------------------------
0.0.0.0/0           10.1.90.1       10.1.90.1       IGP 200         Used
        AS_PATH : (40)
200.1.101.0/24      10.1.90.1       10.1.90.1       INC 0           Rejected
        AS_PATH: path-is-empty
200.1.102.0/24      10.1.90.1       10.1.90.1       INC 0           Rejected
        AS_PATH: path-is-empty

4001:1(config)#show ip bgp neighbors 10.1.1.13 route
The total number of accepted  routes from the neighbor is 1


================================================================================
                        BGP Neighbor Routes - GlobalRouter
================================================================================
NETWORK/MASK        PEER-REM-ADDR   NEXTHOP-ADDRESS ORG LOC-PREF    STATUS

--------------------------------------------------------------------------------
0.0.0.0/0           10.1.1.13       10.1.1.13       IGP 100         Accepted
        AS_PATH : (40)
```

Verify BGP routes

```
9002:1(config)#show ip bgp route

================================================================================
                            BGP Routes - GlobalRouter
================================================================================
The total number of bgp routes in this Vrf are 2

NETWORK/MASK        PEER REM ADDR   NEXTHOP ADDRESS ORG LOC PREF
--------------------------------------------------------------------------------
0.0.0.0/0           10.1.90.1       10.1.90.1       IGP 200
        AS_PATH: (40)
0.0.0.0/0           10.1.1.13       10.1.1.13       IGP 100
        AS_PATH: (40)
Total number of routes displayed are 2
```

Verify route table – for switch 9002, its next-hop for the default route should point to 9001 based on the metric and BGP preference we configured earlier

```
9002:1(config)#show ip route

================================================================================================
                                    IP Route - GlobalRouter
================================================================================================
                                                  NH                    INTER
DST             MASK            NEXT            VRF/ISID        COST  FACE   PROT AGE TYPE PRF
------------------------------------------------------------------------------------------------
0.0.0.0         0.0.0.0         9001            GlobalRouter    1     3051   BGP  0   IBS  175
10.1.1.12       255.255.255.252 10.1.1.14       -               1      1/15  LOC  0   DB   0
10.1.90.1       255.255.255.255 9001            GlobalRouter    10    3051   ISIS 0   IBS  7
10.4.4.1        255.255.255.255 10.4.4.1        -               1     0      LOC  0   DB   0
10.88.1.50      255.255.255.255 8014            GlobalRouter    10    3051   ISIS 0   IBS  7
10.136.0.0      255.255.0.0     10.136.57.1     GlobalRouter    1     57     STAT 0   IB   5
10.136.57.0     255.255.255.0   10.136.57.41    -               1     57     LOC  0   DB   0
135.0.0.0       255.0.0.0       10.136.57.1     GlobalRouter    1     57     STAT 0   IB   5
200.1.101.0     255.255.255.0   8014            GlobalRouter    10    3051   ISIS 0   IBS  7
200.1.102.0     255.255.255.0   8014            GlobalRouter    10    3051   ISIS 0   IBS  7

9001:1(config)#show ip route

================================================================================================
                                    IP Route - GlobalRouter
================================================================================================
                                                  NH                    INTER
DST             MASK            NEXT            VRF/ISID        COST  FACE   PROT AGE TYPE PRF
------------------------------------------------------------------------------------------------
0.0.0.0         0.0.0.0         10.1.1.41       GlobalRouter    1      4/30  BGP  0   IB   45
10.1.1.40       255.255.255.252 10.1.1.42       -               1      4/30  LOC  0   DB   0
10.1.90.1       255.255.255.255 10.1.90.1       -               1     0      LOC  0   DB   0
10.4.4.1        255.255.255.255 4001            GlobalRouter    10    3051   ISIS 0   IBS  7
10.88.1.50      255.255.255.255 8014            GlobalRouter    20    3051   ISIS 0   IBS  7
200.1.101.0     255.255.255.0   8014            GlobalRouter    20    3051   ISIS 0   IBS  7
200.1.102.0     255.255.255.0   8014            GlobalRouter    20    3051   ISIS 0   IBS  7
```

7.2 BGP 4-Byte Autonomous System Numbers (ASN)

Each Autonomous System (AS) must have its own unique number. In previous releases, BGP used two bytes to assign a number to an AS. This 2-byte AS number provides a pool of 65536 unique Autonomous System numbers and is no longer able to meet the demand so Extreme now supports 4-byte AS numbers. You can configure a BGP peer to operate in either the 2-byte AS mode or in the new 4-byte AS mode, but not both.

BGP supports communication between peers of the same type only. If a new 4–byte AS has to communicate with an old 2-byte AS, assign a 2-byte AS number to the new AS. To summarize, BGP currently supports communication between the following peer types only:

- 2-byte peer to 2-byte peer
- 4-byte peer to 4-byte peer

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

BGP 4-byte AS is not supported with confederations.

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

For 4-byte AS numbers, you can enter them as *plain* (regular expressions) or *asdot* (dot notation).

Configuring 4-byte AS numbering

```
ERS8000:5(config)#router bgp as-4-byte enable
ERS8000:5(config)#router bgp <asnum>
```

Configuring dotted octet notation, 1.0 to 65535.65535

```
ERS8000:5(config)#router bgp as-dot enable
ERS8000:5(config)#router bgp <asnum>
```

8. CIDR and Aggregate Addresses

BGP4 supports Classless interdomain routing. (CIDR) is an addressing scheme (also known as super netting) that eliminates the concept of classifying networks into class types. Earlier addressing schemes identified five classes of networks: Class A, Class B, Class C, Class D, and Class E. An example of CIDR would be an address of 192.3.0.0/16, which normally would be an illegal Class C address.

CIDR makes it easy to aggregate several different routes into a single route. This will considerably help reduce the routing table size.

8.1 Configuration Example

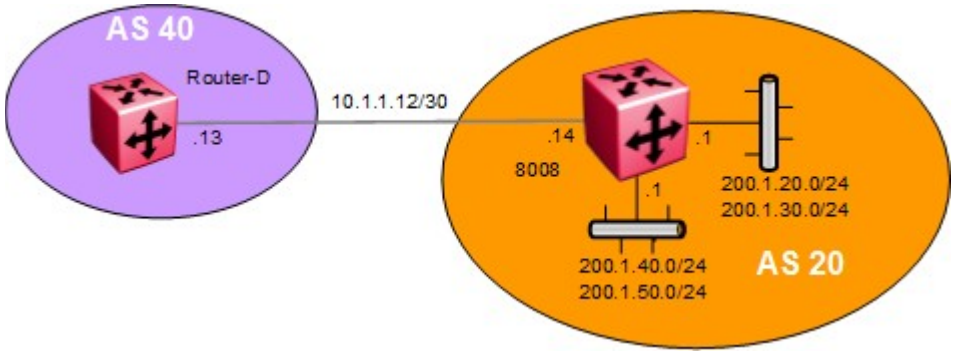

Figure 4: Aggregate Address Configuration Example

In this example, we are going to configure 8008 to summarize all local networks to Router-D in AS 40 with an aggregate route of 200.1.0.0/16.

8.1.1 Configuration

8.1.1.1 Add BGP Neighbor to Router-D

Enable BGP Peer to Router-D in AS 40

```
8008:5(config)#router bgp
8008:5(router-bgp)#neighbor 10.1.1.13
8008:5(router-bgp)#neighbor 10.1.1.13 remote-as 40
8008:5(router-bgp)#neighbor 10.1.1.13 enable
```

8.1.1.2 Add Networks

For this example, we will first add all local networks attached to 8008

Add networks

```
8008:5(config)#router bgp
8008:5(router-bgp)#network 200.1.20.0/24
8008:5(router-bgp)#network 200.1.30.0/24
8008:5(router-bgp)#network 200.1.40.0/24
8008:5(router-bgp)#network 200.1.50.0/24
```

Router-D: verify routes

```
Router-D:5#show ip route -s 200.1.0.0/16
================================================================================
                            IP Route - GlobalRouter
================================================================================
                                                NH          INTER
DST               MASK              NEXT        VRF    COST FACE   PROT AGE TYPE PRF
--------------------------------------------------------------------------------
200.1.1.0         255.255.255.252 10.1.1.14     Glob~ 1      2/2   BGP  0   IB   175
200.1.20.0        255.255.255.0   10.1.1.14     Glob~ 1      2/2   BGP  0   IB   175
200.1.30.0        255.255.255.0   10.1.1.14     Glob~ 1      2/2   BGP  0   IB   175
200.1.40.0        255.255.255.0   10.1.1.14     Glob~ 1      2/2   BGP  0   IB   175
200.1.50.0        255.255.255.0   10.1.1.14     Glob~ 1      2/2   BGP  0   IB   175
```

8.1.1.3 Enable Summarization

Enable summarization of networks you configured in previous step

```
8008:5(config)#router bgp
8008:5(router-bgp)#aggregate-address 200.1.0.0/16 summary-only
8008:5(router-bgp)#network 200.1.50.0/24
```

Router-D: verify routes

```
Router-D:5#show ip route -s 200.1.0.0/16
================================================================================
                            IP Route - GlobalRouter
================================================================================
                                                NH          INTER
DST               MASK              NEXT        VRF    COST FACE   PROT AGE TYPE PRF
--------------------------------------------------------------------------------
200.1.0.0         255.255.0.0     10.1.1.14     Glob~ 1      2/2   BGP  0   IB   175
200.1.1.0         255.255.255.252 10.1.1.14     Glob~ 1      2/2   BGP  0   IB   175
```

9. EBGP Multihop

When two EBGP speakers are directly connected, by default, BGP enforces the one-hop rule for BGP peers. In other words, the remote peer must be located on a directly attached network. However, there may be situations where you may not be able to use the address of the next-hop due to indirect connections such as peering to a circuitless IP address. In this case, BGP multihop is used.

BGP multihop is only used for eBGP connections, not for IBGP connections.

Because the bgp neighbor is not directly connected when using BGP multihop, static routes must also be configured.

By default, the multihop TTL is set for 255. Presently, there is no configuration command to change the TTL setting.

9.1 Configuration Example – BGP Multihop

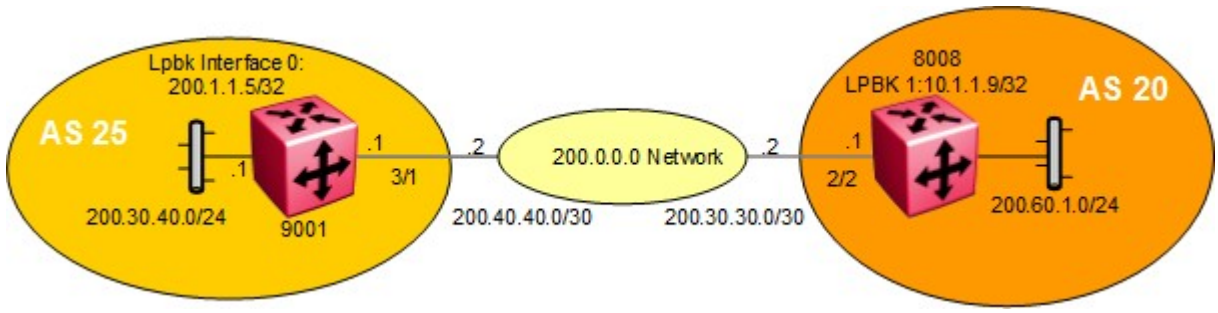

Figure 5: EBGP Configuration Example

For this configuration example, we wish to create a BGP peer between 8008 and 9001. Since they are not directly attached, we will need to enable BGP multihop on 8008 and 9001.

9.1.1 8008 and 9001 Configuration

9.1.1.1 Enable Loopback Address

Add loopback address

8008:

```
8008:5(config)#interface loopback 1
8008:5(config-if)#ip address 10.1.1.9/32
```

9001:

```
9001:1(config)#interface loopback 1
9001:1(config-if)#ip address 200.1.1.5/32
```

9.1.1.2 Add local IP Interfaces

For this configuration example, we will use brouter ports. If you wish to use more than one port on a switch, it is recommended to create add a normal VLAN instead of using brouter ports.

8008 and 9001: Add IP addresses to brouter ports

8008:

```
8008:5(config)#interface gigabitEthernet 2/2
8008:5(config-if)#brouter vlan 2090 subnet 200.30.30.1/30
8008:5(config-if)#exit
8008:5(config)#interface gigabitEthernet 2/3
8008:5(config-if)#brouter vlan 2091 subnet 200.60.1.1/24
8008:5(config-if)#exit
```

9001:

```
9001:1(config)#interface gigabitEthernet 3/1
9001:1(config-if)#brouter vlan 2090 subnet 200.40.40.1/30
9001:1(config-if)#exit
9001:1(config)#interface gigabitEthernet 3/2
9001:1(config-if)#brouter vlan 2091 subnet 200.30.40.1/24
9001:1(config-if)#exit
```

9.1.1.3 Configure BGP local AS

Assign 8008 and 9001 to local AS

8008:

```
8008:5(config)#router bgp 20 enable
```

9001:

```
9001:1(config)#router bgp 25 enable
```

9.1.1.4 Configure BGP Peers

Configure eBGP peer and enable eBGP Multihop

8008:

```
8008:5(config)#router bgp
8008:5(router-bgp)#no synchronization
8008:5(router-bgp)#neighbor 200.40.40.1
8008:5(router-bgp)#neighbor 200.40.40.1 remote-as 25
8008:5(router-bgp)#neighbor 200.40.40.1 ebgp-multihop
8008:5(router-bgp)#neighbor 200.40.40.1 enable
```

9001:

```
9001:1(config)#router bgp
9001:1(router-bgp)#no synchronization
9001:1(router-bgp)#neighbor 200.30.30.1
9001:1(router-bgp)#neighbor 200.30.30.1 remote-as 20
9001:1(router-bgp)#neighbor 200.30.30.1 ebgp-multihop
9001:1(router-bgp)#neighbor 200.30.30.1 enable
```

9.1.1.5 Configure IGP Network Prefixes

Add the appropriate networks which you wish to advertise via BGP

8008:

```
8008:5(config)#router bgp
8008:5(router-bgp)#network 200.60.1.0/24
8008:5(router-bgp)#no  auto-summary
```

9001:

```
9001:1(config)#router bgp
9001:1(router-bgp)#network 200.30.40.0/24
9001:1(router-bgp)#no auto-summary
```

9.1.1.6 Add Static Route

Add a static route to the loopback address to the peer eBGP router

8008:

```
8008:5(config)#ip route 200.40.40.0 255.255.255.252 200.30.30.2 weight 1
```

9001:

```
9001:1(config)#ip route 200.30.30.0 255.255.255.252 200.40.40.2 weight 1
```

10. EBGP Load Balance Using ECMP

Equal Cost Multipath (ECMP) can be used on the Ethernet Routing Switch and Virtual Services Platform to provide load-balance of traffic over 8 paths. A good example of using ECMP is in a dual-home configuration with two connections to two separate routers.

10.1 Configuration Example

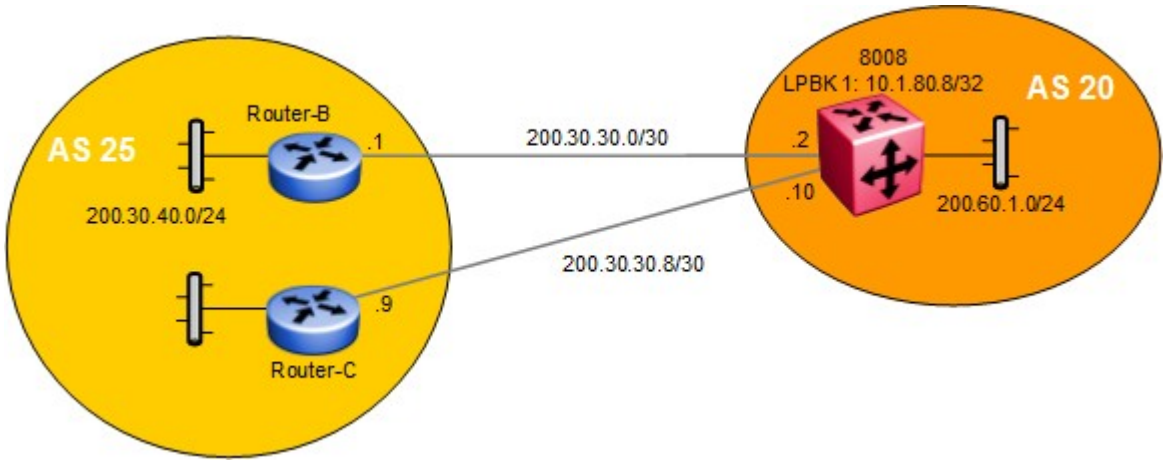

Figure 6: EBGP Configuration Example

In this example, we will configure 8008 to perform EBGP load balance to Router-B and Router-C using ECMP.

10.1.1 8008 Configuration

10.1.1.1 Enable LPBK Address

Add loopback address

```
8008:5(config)#interface loopback 1
8008:5(config-if)#interface ip address 10.1.80.8/32
8008:5(config-if)#exit
8008:5(config)#router ospf
8008:5(config-ospf)#router-id 10.1.80.8
8008:5(config-ospf)#exit
```

10.1.1.2 Configure BGP

BGP configuration with multiple paths set to 2

```
8008:5(config)#router bgp 20 enable
8008:5(config)#router bgp
8008:5(router-bgp)#no synchronization
```

```
8008:5(router-bgp)#bgp multiple-paths 2
```

10.1.1.3 Enable ECMP

Enable ECMP

```
8008:5(config)#ip ecmp
```

10.1.1.4 Add BGP interfaces

Add BGP interfaces assuming we are using brouter ports

```
8008:5(config)#interface gigabitEthernet 2/2
8008:5(config-if)#brouter vlan 2090 subnet 200.30.30.2/30
8008:5(config-if)#exit
8008:5(config)#interface gigabitEthernet 2/5
8008:5(config-if)#brouter vlan 2091 subnet 200.30.30.10/30
8008:5(config-if)#exit
```

10.1.1.5 Configure BGP Peers

Add BGP peers

```
8008:5(config)#router bgp
8008:5(router-bgp)#neighbor 200.30.30.1
8008:5(router-bgp)#neighbor 200.30.30.1 remote-as 25 enable
8008:5(router-bgp)#neighbor 10.30.30.9
8008:5(router-bgp)#neighbor 10.30.30.9 remote-as 25 enable
8008:5(router-bgp)#exit
```

10.1.1.6 Configure IGP Network Prefixes

Add the appropriate networks which you wish to advertise via BGP

```
8008:5(config)#router bgp
8008:5(router-bgp)#network 200.60.1.0/24
```

11. BGP Synchronization and Next-Hop-Self

BGP synchronization, depending on if it is enabled or not, either enables or disables the router from accepting or forwarding routes from BGP peers without waiting for an update from the IGP. It is used when there are routers in the AS not running BGP. With synchronization enabled, the router should not advertise a route until all the routers in the AS have learned the route via IGP.

Normally synchronization should be enabled unless the AS is a stub AS and does not pass traffic from one AS to another or all the routers in the AS run BGP. In other words, if all the routers in your AS are running BGP, there is no need to enable BGP Synchronization.

If there are routers in the AS not running BGP, BGP Synchronization should normally be left enabled. When BGP Synchronization is enabled, a BGP router will wait to learn routes from an IGP, such as OSPF, before advertising routes learned by BGP. If all the routers within AS are expected to forward traffic outside the local AS to other AS's, BGP should be redistributed into the IGP so that the router(s) not running BGP will learn how to forward traffic to the external networks. Note that re distributing BGP routes info IGP should only be done for networks where it is a limited number of eBGP routes that need to be redistributed. Redistributing thousands of routes into IGP, such as OSPF, consumes both CPU and memory resources.

For example, looking at the figure 7 below, assuming the AS20 consists of Router-C not running BGP. There is an iBGP peer between 8008 and 9001 with Synchronization disabled on both routers, and BGP redistribution into OSPF enabled. When 9001 learns a route via eBGP from Router-F, 9001 will propagate this route to 8008 using iBGP. If 8008 now propagates this route to Router-E before Router-C within AS20 has learned this route, then Router-E could start sending traffic for this route before Router-C is ready to forward this traffic. Enabling Synchronization solves this problem by preventing a BGP speaker from advertising a route over eBGP until all routes within an AS have learned the route.

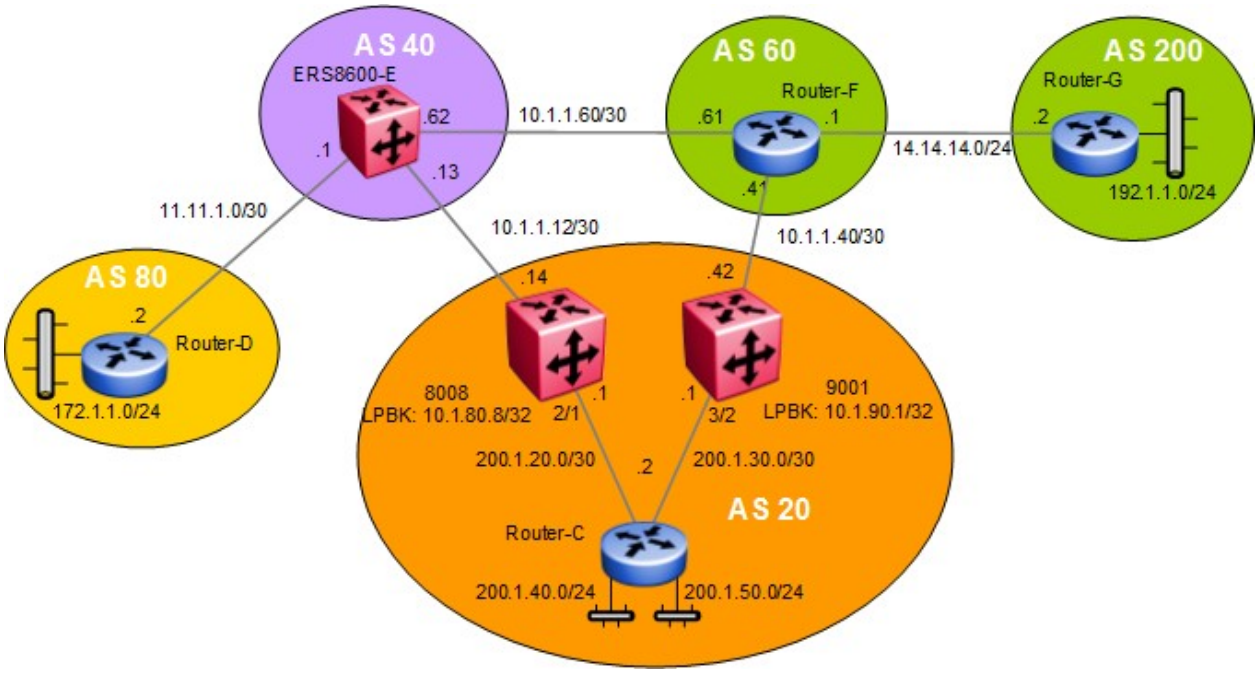

Figure 7: BGP Synchronization and Self Hop Configuration Example

11.1 Configuration Example 1 – Initial Configuration

In this example, we will show the effects of enabling and disabling BGP synchronization and next hop-self by performing the following:

- Enabling BGP synchronization on both 8008 and 9001
- Remove the connection between 9001 and Router-F to see how this affects 8008

BGP route distribution into OSPF is not enabled on either 8008 or 9001. Router-C is not running BGP, only OSPF as an IGP.

11.1.1 Configuration – With BGP Synchronization Enabled

11.1.1.1 Enable OSPF Interface Using Brouter Ports

For this configuration example, brouter ports are used. Either VLAN or brouter ports can be used.

Configure OSPF interfaces

8008:

```
8008:5(config)#interface GigabitEthernet 2/1
8008:5(config)#brouter vlan 2095 subnet 200.1.20.1/30
8008:5(config)#ip ospf enable
```

9001:

```
9001:1(config)#interface GigabitEthernet 3/2
9001:1(config)#brouter vlan 2095 subnet 200.1.30.1/30
9001:1(config)#ip ospf enable
```

11.1.1.2 Configure loopback address

Add lpbk address and enable ospf

8008:

```
8008:5(config)#interface loopback 1
8008:5(config-if)#interface ip address 10.1.80.8/32
8008:5(config-if)#ip ospf
```

9001:

```
9001:1(config)#interface loopback 1
9001:1(config-if)#interface ip address 10.1.90.1/32
9001:1(config-if)#ip ospf
```

11.1.1.3 Configure OSPF

Enable OSPF with ASBR

8008:

```
8008:1(config)#router ospf
8008:5(config-ospf)#as-boundary-router enable
8008:5(config-ospf)#router-id 10.1.80.8
8008:5(config-ospf)#exit
8008:5(config)#router ospf enable
```

9001:

```
9001:1(config)#router ospf
9001:1(config-ospf)#as-boundary-router enable
9001:1(config-ospf)#router-id 10.1.90.1
9001:1(config-ospf)#exit
9001:1(config)#router ospf enable
```

11.1.1.4 Configure BGP Globally

Enable BGP with synchronization enabled (default setting)

8008 & 9001: Same configuration on both switches

```
8008:5(config)#router bgp 20 enable
8008:5(config)#router bgp
8008:5(router-bgp)synchronization
```

11.1.1.5 Add BGP Network Prefixes

Add the appropriate networks which you wish to advertise via BGP, i.e. local interfaces

8008:

```
8008:5(config)#router bgp
8008:5(router-bgp)#network 200.1.20.0/30
8008:5(router-bgp)#network 10.1.1.12/30
```

9001:

```
9001:1(config)#router bgp
9001:1(router-bgp)#network 200.1.30.0/30
9001:1(router-bgp)#network 10.1.1.40/30
```

11.1.1.6 Add BGP peers

Add BGP peers

8008:

```
8008:5(config)#router bgp
8008:5(router-bgp)#neighbor 200.1.30.1
8008:5(router-bgp)#neighbor 200.1.30.1 remote-as 20 enable
8008:5(router-bgp)#neighbor 10.1.1.13
8008:5(router-bgp)#neighbor 10.1.1.13 remote-as 40 enable
8008:5(router-bgp)#exit
```

9001:

```
9001:1(config)#router bgp
9001:1(router-bgp)#neighbor 200.1.20.1
9001:1(router-bgp)#neighbor 200.1.20.1 remote-as 20 enable
9001:1(router-bgp)#neighbor 10.1.1.41
9001:1(router-bgp)#neighbor 10.1.1.41 remote-as 60 enable
9001:1(router-bgp)#exit
```

11.1.2 Verify Operations

11.1.2.1 Verify BGP State

Verify that BGP is established on both 8008 and 9001

8008:

```
8008:5(config)#show ip bgp summary
================================================================================
                          BGP Summary - GlobalRouter
================================================================================
                                  BGP version - 4
                                     local-as - 20
                                   Identifier - 10.1.80.8
                                Decision state - Idle
                  The total number of routes is   3

BGP NEIGHBOR INFO :
   NEIGHBOR       RMTAS     STATE    HLDTM  KPALV  HLDCFG KPCFG   WGHT  CONRTY ADVINT
--------------------------------------------------------------------------------------
200.1.30.1        20     Established  180    60     180    60     100    120    5
10.1.1.13         40     Established  180    60     180    60     100    120    5
```

9001:

```
9001:1(config)#show ip bgp sum
================================================================================
                          BGP Summary - GlobalRouter
================================================================================
                                  BGP version - 4
                                     local-as - 20
                                   Identifier - 10.1.90.1
                                Decision state - Idle
                  The total number of routes is   5

BGP NEIGHBOR INFO :
   NEIGHBOR       RMTAS     STATE    HLDTM  KPALV  HLDCFG KPCFG   WGHT  CONRTY ADVINT
--------------------------------------------------------------------------------------
200.1.20.1        20          Established 180   60    180    60    100   120    5
10.1.1.41         40          Established 180   60    180    60    100   120    5
```

11.1.2.2 Viewing the BGP Route Table on 9001

The following results show the effect of what happens to route table on 8008 when the connection between 9001 and Router-F is broken.

Step 1: If the connection between 9001 and Router-F is removed, by using the show ip bgp route command, the following routes are displayed:

```
9001:1#show ip bgp route
================================================================================
                            BGP Routes - GlobalRouter
================================================================================
The total number of routes is 11

NETWORK/MASK        PEER REM ADDR    NEXTHOP ADDRESS ORG LOC PREF
--------------------------------------------------------------------------------
10.1.1.12/30         200.1.20.1     200.1.20.1       IGP 100
        AS_PATH: path-is-empty
11.11.1.0/30         200.1.20.1     10.1.1.13        IGP 100
        AS_PATH: (40)
14.14.14.0/24        200.1.20.1     10.1.1.13         IGP 100
        AS_PATH: (40 60)
10.1.1.40/30         200.1.20.1     10.1.1.13         IGP 100
        AS_PATH: (40 60)
10.1.1.60/30         200.1.20.1     10.1.1.13         IGP 100
        AS_PATH: (40)
172.1.1.0/24         200.1.20.1     10.1.1.13         IGP 100
        AS_PATH: (40 80)
172.1.2.0/30         200.1.20.1     10.1.1.13         IGP 100
        AS_PATH: (40 80)
192.1.1.0/24         200.1.20.1     10.1.1.13         IGP 100
        AS_PATH: (40 60 200)
200.20.20.0/24     200.1.20.1     200.1.20.1        IGP 100
        AS_PATH: path-is-empty
200.1.20.0/30        200.1.20.1     200.1.20.1        IGP 100
        AS_PATH: path-is-empty
```

Step 2: Notice that because the connection to Router-F is down, all external routes via the next hop address of 10.1.1.13 are not available. If we now look at the regular route table as shown below, notice is no external route information. None of the BGP entries are in the route table because the next hop (10.1.1.13) for these entries is unreachable and are not learned through OSPF.

```
9001:1#show ip route
================================================================================
                                IP Route - GlobalRouter
================================================================================
                                                NH                  INTER
DST           MASK          NEXT                VRF              COST FACE  PROT AGE TYPE PRF
--------------------------------------------------------------------------------
10.1.80.8     255.255.255.255 200.1.30.2    GlobalRouter     21    3/2   OSPF 0   IB   20
10.1.90.1     255.255.255.255 10.1.90.1       -              1     0     LOC  0   DB   0
200.1.20.0    255.255.255.252 200.1.30.2    GlobalRouter     11    3/2   OSPF 0   IB   20
200.1.30.0    255.255.255.252 200.1.30.1      -              -     3/2   LOC  0   DB   0
200.1.40.0    255.255.255.0   200.1.30.2    GlobalRouter     11    3/2   OSPF 0   IB   20
200.1.50.0    255.255.255.0   200.1.30.2    GlobalRouter     11    3/2   OSPF 0   IB   20
```

11.2 Correcting the Next Hop Problem

One method to get around the next hop problem is to enable the BGP next hop-self command. Another alternative method would be to enable OSPF passive on the 8008 interface connecting to Router-C.

11.3 How to Correct the Next Hop Problem from Step 11.1

This section describes how to resolve the next hop problem pointed out in the previous section. To resolve this problem, complete the following steps.

11.3.1 Configuration – Enabling BGP Next Hop-Self and Synchronization

11.3.1.1 Configure 8008 for Next Hop-Self

Enable the nexthop-self parameter.

```
8008:5(config)#router bgp
8008:5(router-bgp)#no neighbor 200.1.30.1 enable
8008:5(router-bgp)#neighbor 200.1.30.1 next-hop-self enable
8008:5(router-bgp)#neighbor 200.1.30.1 enable
```

11.3.2 Verify Operations

11.3.2.1 Viewing the BGP Route Table on 9001

Step 1: If you now view the bgp route table on 9001, it will look like the following.

```
9001:1#show ip bgp route

================================================================================
                          BGP Routes - GlobalRouter
================================================================================
NETWORK/MASK       PEER REM ADDR   NEXTHOP ADDRESS ORG LOC PREF
--------------------------------------------------------------------------------
11.11.1.0/30         200.1.20.1          200.1.20.1          IGP 100
        AS_PATH: (40)
10.1.1.12/30         200.1.20.1          200.1.20.1          IGP 100
        AS_PATH: path-is-empty
14.14.14.0/24        200.1.20.1          200.1.20.1          IGP 100
        AS_PATH: (40 60)
10.1.1.60/30         200.1.20.1          200.1.20.1          IGP 100
        AS_PATH: (40)
172.1.1.0/24         200.1.20.1          200.1.20.1          IGP 100
        AS_PATH: (40 80)
172.1.2.0/30         200.1.20.1          200.1.20.1          IGP 100
        AS_PATH: (40 80)
192.1.1.0/24         200.1.20.1          200.1.20.1          IGP 100
        AS_PATH: (40 60 200)
200.20.20.0/24       200.1.20.1          200.1.20.1          IGP 100
        AS_PATH: path-is-empty
200.1.20.0/30        200.1.20.1          200.1.20.1          IGP 100
        AS_PATH: path-is-empty
```

Step 2: Notice that now the next hop contains the next of 200.1.20.1 for all routes. However, if we look the IP route table, it still will not have changed:

```
9001:1#show ip route
================================================================================
                                  IP Route - GlobalRouter
================================================================================
                                          NH                      INTER
DST           MASK          NEXT          VRF              COST  FACE   PROT AGE TYPE PRF
--------------------------------------------------------------------------------
10.1.80.8     255.255.255.255 200.1.30.2   GlobalRouter     21    3/2    OSPF 0   IB   20
10.1.90.1     255.255.255.255 10.1.90.1     -               1     0      LOC  0   DB   0
200.1.20.0    255.255.255.252 200.1.30.2  GlobalRouter      11    3/2    OSPF 0   IB   20
200.1.30.0    255.255.255.252 200.1.30.1    -               -     3/2    LOC  0   DB   0
200.1.40.0    255.255.255.0   200.1.30.2   GlobalRouter      11    3/2    OSPF 0   IB   20
200.1.50.0    255.255.255.0   200.1.30.2   GlobalRouter      11    3/2    OSPF 0   IB   20
```

As you will notice from the route table above, the BGP entries still do not appear in the IP routing table. In order to get the BGP routes to appear in the IP routing table, BGP synchronization must be disabled. Since IGP is not synchronized with BGP, BGP entries are not put into the IP forwarding table.

11.3.2.2 Disabling Synchronization on 9001 and 8008

Disable BGP synchronization

8008:

```
8008:5(config)#router bgp
8008:5(router-bgp)#no synchronization
```

9001:

```
9001:1(config)#router bgp
9001:1(router-bgp)#no synchronization
```

11.3.3 Verifying Operation

11.3.3.1 Viewing the BGP Route Table on 9001

If you now view the ip route table on 9001, it will look like the following:

```
9001:1#show ip route

================================================================================
                                IP Route - GlobalRouter
================================================================================
                                          NH                    INTER
DST           MASK          NEXT          VRF              COST FACE  PROT AGE TYPE PRF
--------------------------------------------------------------------------------
10.1.1.12     255.255.255.252 200.1.30.2    GlobalRouter     0     3/2   BGP  0   IB   175
10.1.80.8     255.255.255.255 200.1.30.2    GlobalRouter     21    3/2   OSPF 0   IB   20
10.1.90.1     255.255.255.255 10.1.90.1      -               1    0      LOC  0   DB   0
11.11.1.0     255.255.255.252 200.1.30.2  GlobalRouter       1     3/2   BGP  0   IB   175
172.1.1.0     255.255.255.0   200.1.30.2  GlobalRouter       2     3/2   BGF  0   IB   175
172.1.2.0     255.255.255.0   200.1.30.2  GlobalRouter       2     3/2   BGF  0   IB   175
192.1.1.0     255.255.255.0   200.1.30.2  GlobalRouter       3     3/2   BGF  0   IB   175
200.1.20.0    255.255.255.252 200.1.30.2  GlobalRouter       11    3/2   OSPF 0   IB   20
200.1.30.0    255.255.255.252 200.1.30.1    -                 -    3/2   LOC  0   DB   0
200.1.40.0    255.255.255.0   200.1.30.2    GlobalRouter      11    3/2   OSPF 0   IB   20
200.1.50.0    255.255.255.0   200.1.30.2    GlobalRouter      11    3/2   OSPF 0   IB   20
```

We have solved the route table problem in 9001, however; it may still not be able to get to any of the external networks. This is because Router-C has no knowledge of these external routes. One method to correct this problem is to enable BGP to OSPF redistribution. Another method is to enable BGP on Router-E. Please see example above on how to redistribute BGP routes into OSPF.

12. MD5 Authentication Configuration Example

MD5 Authentication is a method to securing TCP connections that support BGP sessions. MD5 assigns an authentication key to each BGP router, which then attaches a computed MD5 signature to each BGP packet.

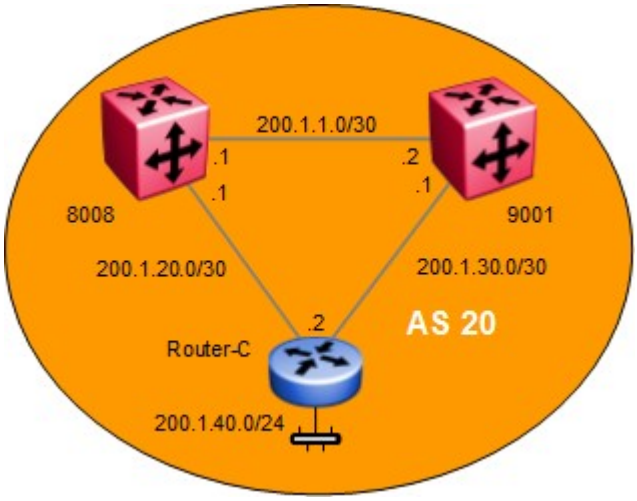

Figure 8: BGP MD5 Configuration Example

In the configuration example below, we will configure MD-5 authentication between 8008 and 9001.

12.1 MD5 Configuration

12.1.1 Configure ERS8000 and VSP 9000 for MD-5 Authentication

12.1.1.1 Configure MD-5 Authentication

Disable the administration state for the peer

8008:

```
8008:5(config)#router bgp
8008:5(router-bgp)#no neighbor 200.1.1.2 enable
```

9001:

```
9001:1(config)#router bgp
9001:1(router-bgp)#no neighbor 200.1.1.1 enable
```

12.1.1.2 Enable MD5 authentication

Add password (the secret key) and enable MD5 authentication

8008:

```
8008:5(config)#router bgp
8008:5(router-bgp)#neighbor 200.1.1.2 md5-authentication enable
8008:5(router-bgp)#neighbor password 200.1.1.2 <password>
8008:5(router-bgp)#neighbor 200.1.1.2 enable
```

9001:

```
9001:1(config)#router bgp
9001:1(router-bgp)#neighbor 200.1.1.1 md5-authentication enable
9001:1(router-bgp)#neighbor password 200.1.1.1 <password>
9001:1(router-bgp)#neighbor 200.1.1.1 enable
```

The MD5 password can have a string length of up to 1536 characters.

13. BGP Peer Group Configuration Example

BGP peer groups are used when a group of BGP neighbors share the same update policies. All the members of the peer group will share the same configuration options.

The following example shows the configuration for 8008 applied to all routers within AS 20

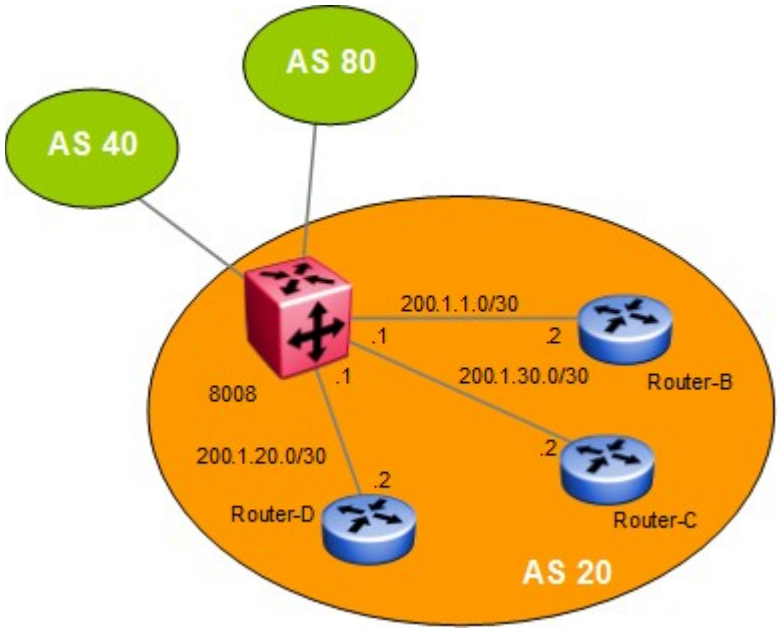

Figure 9: BGP Peer Group Configuration Example

In the example below, Router-B, C, and D all use a BGP Keep Alive timer of 60 and Hold Timer of 180. By using a Peer Group on 8008, we can change the Keep Alive and Hold timer once and apply this Peer Group configuration to the BGP neighbors Router-B, C, and D. If required, route policies can also be added to the Peer Group configuration.

Changes to a peer group only apply to the current members. This means that you have to add the neighbors before making peer group changes.

13.1 BGP Peer Group Configuration

13.1.1 Create the Peer Group (Group_1)

Create peer group using the name Group_1

```
8008:5(config)#router bgp
8008:5(router-bgp)#neighbor Group_1
```

Note that the assigned peer group name is context-sensitive. For example, the name string “Group_1” is Not the same as ‘group_1”.

13.1.2 Create BGP Peers

Create BGP peers to Router B, C and D

```
8008:5(router-bgp)#neighbor 200.1.1.2
8008:5(router-bgp)#neighbor 200.1.20.2
8008:5(router-bgp)#neighbor 200.1.30.2
```

13.1.3 Add Peers as Member of Group_1

Add peer to peer group Group_1

```
8008:5(router-bgp)#neighbor 200.1.1.2 peer-group Group_1
8008:5(router-bgp)#neighbor 200.1.20.2 peer-group Group_1
8008:5(router-bgp)#neighbor 200.1.30.2 peer-group Group_1
```

13.1.4 Assign Peer Group to AS 20

Assign Group_1 to remote_as 20

```
8008:5(router-bgp)#neighbor Group_1 remote-as 20
```

13.1.5 Assign Variables to Peer Group

Change the peer group Group_1 BGP keep-alive timer to 60 seconds and the BGP hold-down time to 180 seconds

```
8008:5(router-bgp)#neighbor Group_1 timers 60 180
```

13.1.6 Enable the Peer Group

Enable the peer group; once enabled, 8008 should peer with Router-B, Router-C, and Router-D

```
8008:5(router-bgp)#neighbor Group_1 enable
```

13.1.7 Assigning Policies to Peer Group

If required, you can assign a policy to a peer group. Assuming you have an existing policy named Pref_20, enter the following command to add this policy to the peer group Group_1

If required, add an existing policy (Pref_20) to Group_1

```
8008:5(router-bgp)#neighbor Group_1 in-route-map Pref_AX20
```

BGP Peer Groups are used to apply changes to all group members. You can still enter specific settings for each peer directly. Unless you set the parameter again in the peer

group, it will not be overwritten on the peer.

14. Route Selection and Traffic Management – BGP Path Attributes

The Extreme Ethernet Routing Switch and Virtual Switching Platforms use route policies to control traffic flow. By using policies, traffic can be controlled over multiple connections for inbound traffic from other ASs and outbound traffic that comes from outside a particular AS.

Overall, policies are created to control routes, work with default routing, control specific and aggregated routes, and manipulate BGP attributes. The rest of this section deals with BGP Path Attributes.

Path attributes fall into four separate categories

1. Well-known mandatory – Attributes are mandatory and must be included in every UPDATE message
2. Well-known discretionary – Discretionary may or may not be sent in a particular UPDATE message
3. Optional transitive – Optional transitive attribute is accepted and passed along to other BGP peers
4. Optional non-transitive – Optional non-transitive attribute must be accepted or ignored and not passed along to other BGP peers.

Path attributes help the border routers to select among paths using built-in algorithms or manually configured polices.

Various attributes are used to decide the path a BGP router will take. The following attributes are used by BGP in deciding what path to take.

- Origin – (well-known mandatory)
- AS_path – (well-known mandatory)
- Next Hop – (well-known mandatory)
- Multi-Exit Discriminator Attribute – (optional non-transitive)
- Local Preference – (well-known discretionary)
- Atomic Aggregate – (well-known discretionary)
- Aggregator – (optional transitive)
- Community Attribute – (optional transitive)

14.1 Origin Attribute (Type 1)

The Origin attribute is a well-known mandatory attribute that specifies the source of a route. The Origin is created by the AS that originates the route and includes the following possible values:

- *IGP* – The route is interior to the originating AS that inserts this route into the BGP table (0 = IGP).
- *EGP* – The route is learned via the Exterior Gateway Protocol (EGP) prior to being inserted into the BGP table (1 = BGP).
- *Incomplete (INC)* – The origin of the route is unknown or learned by some other means. For example, these routes could be learned through RIP, OSPF, or static routes (2 = INComplete)

BGP uses the Origin attribute in its decision making process. BGP prefers the path with the lowest origin type. IGP is the lowest Origin type followed by EGP and Incomplete.

14.1.1 Origin Attribute Configuration Example – Static Route Distribution

In this example, we will configure 9001 to distribute static routes for network 44.44.44.0/24. ERS8000-D route table should display this static route as INC (incomplete) and all other routes as IGP.

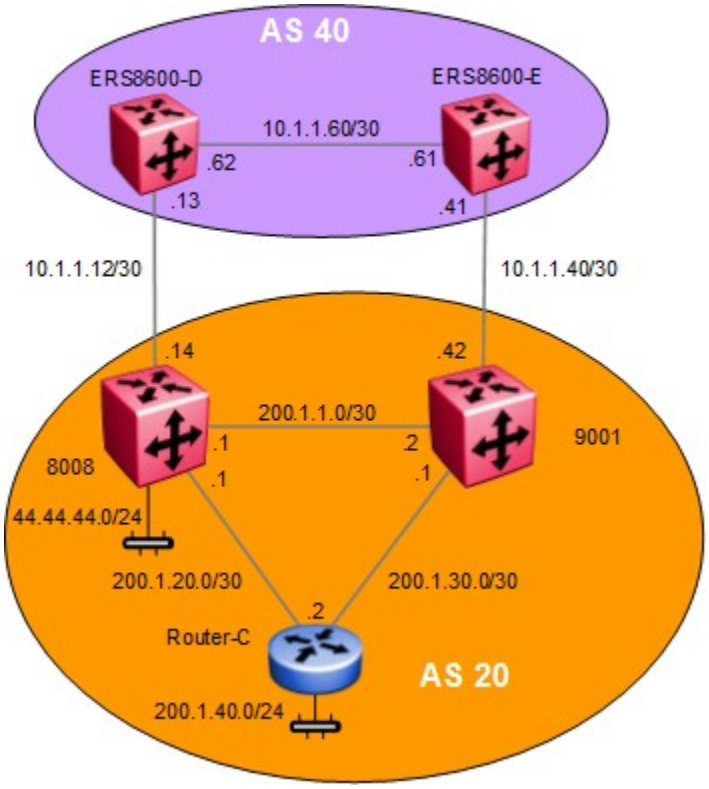

Figure 10: BGP Origin Attribute Configuration Example

14.1.1.1 Configuration

14.1.1.1.1 Add a Static Route to Network 44.44.44.0/24

Create an IP static route on 9001

```
9001:1(config)#ip route 44.44.44.0 255.255.255.0 200.1.1.1 weight 1
```

14.1.1.1.2 Create a Static Route Redistribution Policy

Create static route redistribution policy

```
9001:1(config)#router bgp
9001:1(router-bgp)#redistribute static
9001:1(router-bgp)#redistribute static enable
9001:1(router-bgp)#end
9001:1#ip bgp apply redistribute static
```

14.1.1.1.3 Add BGP Networks you Wish to Advertise

In this example, we will only advertise the network 200.1.1.0/30 via BGP

```
9001:1(config)#router bgp
9001:1(router-bgp)#network 200.1.1.0/30
```

14.1.1.2 Verification

14.1.1.2.1 Verify Static Route Redistribution via ERS8000-D

If we go to ERS8000-D and view its route table, it should display the static route advertised from 9001 with a BGP attribute of INC as shown below and the network 200.1.1.0/30 as IGP

```
ERS8000-D:5#show ip bgp route
Network/Mask        Peer Rem Addr    NextHop Address Org Loc Pref    B/U Sl
------------------------ -------------------- -------------------- ----- -------------- ----- ---
10.1.1.12/30         10.1.1.62       10.1.1.62        IGP      10 B/U   4
        As Path:
10.1.1.40/30         10.1.1.42       10.1.1.42        IGP      100 B     4
        As Path: <20>
10.1.1.40/30         10.1.1.62       10.1.1.62        IGP      0         4
        As Path:
10.1.1.60/30         10.1.1.62       10.1.1.62        IGP      0 B       4
        As Path:
11.11.1.0/30         10.1.1.62       10.1.1.62        IGP      0 B/U   4
        As Path:
44.44.44.0/24        10.1.1.42       10.1.1.42        INC      0         4
        As Path: <20>
172.1.1.0/24         10.1.1.62       11.11.1.2        IGP      10        4
        As Path: <80>
172.1.2.0/30         10.1.1.62       11.11.1.2        IGP      10        4
        As Path: <80>
192.1.1.0/24         14.14.14.2      14.14.14.2       IGP      100 B/U   4
        As Path: <200>
200.1.1.0/30         10.1.1.42       10.1.1.42        IGP      100 B/U   4
       As Path: <20>
200.1.1.0/30         10.1.1.62       10.1.1.14        IGP      10    4
       As Path: <20>
200.1.20.0/30        10.1.1.62       10.1.1.14        IGP      10    4
       As Path: <20>
200.1.30.0/30        10.1.1.42       10.1.1.42        IGP      100 B/U    4
       As Path: <20>
200.1.30.0/30        10.1.1.62       10.1.1.14        IGP      10      4
       As Path: <20>
```

14.1.2 Changing the Origin Type

When using the BGP *network* command, the origin should show up as IGP on each of the remote peers. In some cases, you may wish to change the origin to a lower priority origin type such as INC to influence the route decision at a remote peer. A path with the lowest origin type will be selected, where IGP is lower than EGP and EGP is lower than INComplete.

For example, let's assume we wish to take the path from 9001 to ERS8000-E to get to network 200.1.40.0/24. To do so, we will configure a route policy on 8008 to advertise network 200.1.40.0/24 with an origin of INC and apply this to the BGP peer containing neighbor 10.1.1.13 or ERS8000-D.

If we look at the BGP route table on ERS8000-D, you will notice the path to 200.1.40.0/24 has two paths both with an origin of IGP with the preferred route via 10.1.1.14.

- ERS8000-D# ***show ip bgp route 200.1.40.0/24***

```
The total number of routes is 22
Network/Mask        Peer Rem Addr    NextHop Address Org Loc Pref
------------------------ -------------------- -------------------- ----- --------------
200.1.40.0/24       10.1.1.14        10.1.1.14        IGP 100
                AS_PATH: (1)
200.1.40.0/24       10.1.1.61        10.1.1.42        IGP 100
                AS_PATH: (1)
```

- ERS8000-D# ***show ip route -s 200.1.40.0/24***

```
================================================================================
                                    Ip Route
================================================================================
            DST             MASK             NEXT COST VLAN  PORT PROT AGE TYPE PRF
--------------------------------------------------------------------------------------------------------
     200.1.40.0   255.255.255.0       10.1.1.14    1 2171  4/6   BGP   0 IB    45

1 out of 14 Total Num of Route Entries, 14 Total Num of Dest Networks displayed.
--------------------------------------------------------------------------------------------------------
TYPE Legend:
I=Indirect Route, D=Direct Route, A=Alternative Route, B=Best Route, E=Ecmp Route
, U=Unresolved Route, N=Not in HW
```

In the configuration steps that follow, we will configure a route policy to advertise from 8008 to ERS8000-D the route to 200.1.40.0/24 with an origin of INC. This will result in ERS8000-D using the next-hop of 10.1.1.61 (via ERS8000-E) to get to network 200.1.40.0/24.

14.1.2.1 Configuration – Changing the Origin Type

14.1.2.1.1 Configure an IP Prefix Named 200.1.40.0 and add IP Prefix 200.1.40.0/24

Configure the IP prefix list named 200.1.40.0

```
8008:5(config)#ip prefix-list 200.1.40.0 200.1.40.0/24
```

14.1.2.1.2 Configure an IP Policy

Configure the IP route policy named BGP_org_routerb

```
8008:5(config)#route-map BGP_org_routerd 1
8008:5(route-map)#match  protocol any
8008:5(route-map)#match network 200.1.40.0
8008:5(route-map)#set origin incomplete
8008:5(route-map)#enable
```

14.1.2.1.3 Add the Policy to the BGP Peer Router-D

Add this policy to the BGP peer Router-D

```
8008:5(config)#router bgp
8008:5(router-bgp)#neighbor 10.1.1.13 out-route-map BGP_org_routerd
8008:5(router-bgp)#end
8008:5#ip bgp restart-bgp neighbor 10.1.1.13 soft-reconfiguration out
```

14.1.2.1.4 Soft Start BGP Peer

Soft restart the BGP peer

```
8008:5#ip bgp restart-bgp neighbor 10.1.1.13 soft-reconfiguration out
```

14.1.2.2 Verify Operations

14.1.2.2.1 Verify Route to 200.1.400/24 via ERS8000-D

Step 1 – Once the policy has been added, the route 200.1.40.0/24 from ERS8000-D's perspective for peer 10.1.1.14 has been changed to INC. Hence, the path to 200.1.40.0/24 should now be via ERS8000-E:

```
ERS8000-D:5#show ip bgp route 200.1.40.0/24
The total number of routes is 13
Network/Mask        Peer Rem Addr    NextHop Address Org Loc Pref
------------------------ -------------------- -------------------- ----- --------------
200.1.40.0/24       10.1.1.61        10.1.1.42        IGP 100
                AS_PATH: (1)
200.1.40.0/24       10.1.1.14        10.1.1.14        INC 100
                AS_PATH: (1)
```

Step 2 – Once the policy has been added, the route 200.1.40.0/24 from ERS8000-D's perspective for peer 10.1.1.14 has been changed to INC. Hence, the path to 200.1.40.0/24 is now via ERS8000-E:

```
ERS8000-D:5#show ip route -s 200.1.40.0/24
================================================================================
                                    Ip Route
================================================================================
             DST             MASK             NEXT COST VLAN   PORT PROT AGE TYPE PRF
--------------------------------------------------------------------------------------------------------
      200.1.40.0    255.255.255.0        10.1.1.61     1 2170  4/1    BGP    0 IB    175

1 out of 13 Total Num of Route Entries, 13 Total Num of Dest Networks displayed.
--------------------------------------------------------------------------------------------------------
TYPE Legend:
I=Indirect Route, D=Direct Route, A=Alternative Route, B=Best Route, E=Ecmp Route
, U=Unresolved Route, N=Not in HW
```

14.2 AS Path Attribute (Type 2)

Whenever a route passes through one AS to another, the new AS prepends its AS number to the update. The ordered list is called the AS Sequence.

The AS Path attribute helps to ensure a loop-free topology. IBGP connections do not change the AS Path as these connections reside within a specific AS.

BGP always prefers the shortest path. Hence, by manipulating the AS Path to a remote EBGP peer, we can influence the incoming route selection where there is more than one path to the local AS.

14.2.1 Config Example: Load Balance Approach using AS Path to Influence Inbound Traffic Flow

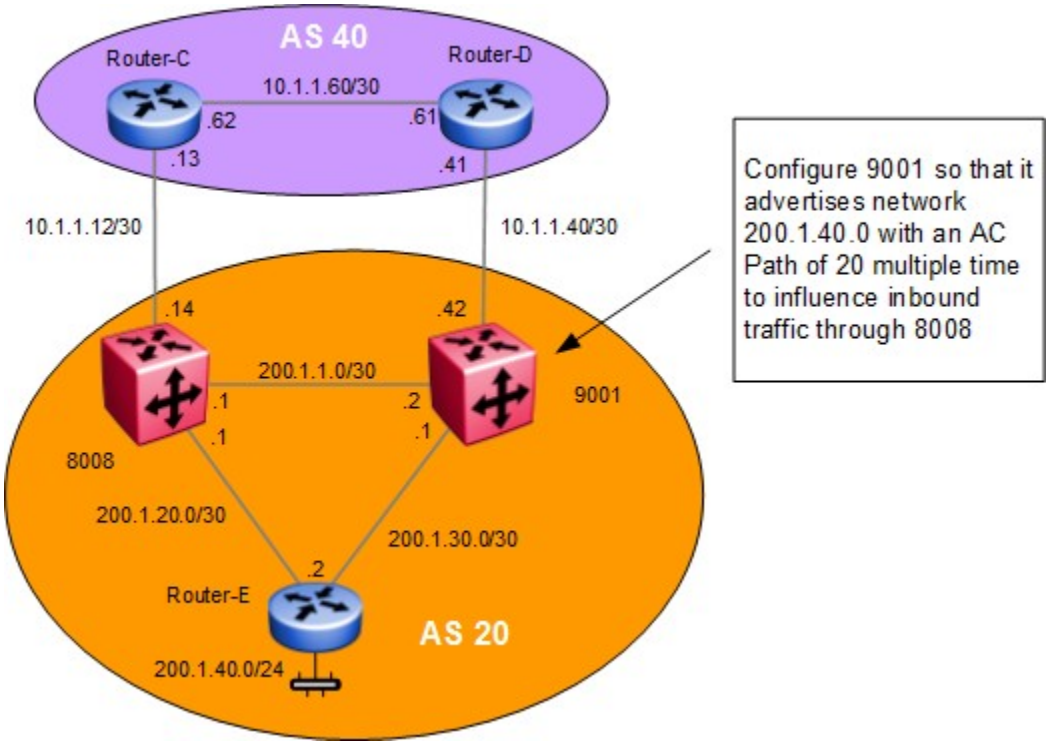

Figure 11: BGP AS Path Configuration Example

In this example, 8008 will advertise the network 200.1.40.0 unchanged. 9001 will be setup to have its internal AS number inserted into the AS Path multiple times. This should influence all inbound traffic destined for the 200.1.40.0 over 8008.

14.2.1.1 Configure an IP Prefix Named 200.1.40.0 and Add IP Prefix 200.1.40.0/24 on 9001

Configure the IP prefix list named 200.1.40.0

```
9001:1(config)#ip prefix-list 200.1.40.0 200.1.40.0/24
```

14.2.1.2 Add IP AS List

Configure IP AS list to advertise AS 20 multiple times, For this example, we will use as-list 1 and advertise AS 20 three times.

```
9001:1(config)#ip as-list 1 memberid 1 permit as-path "20 20 20"
```

14.2.1.3 Add a Route Policy

Configure the IP route policy named AS_Prepend and add sequence 1 to match on network 200.1.40.0 and then append AS path 1 (AS path ’20 20 20’)

```
9001:1(config)#route-map AS_Prepend 1
9001:1(route-map)#match network 200.1.40.0
9001:1(route-map)#set as-path 1
9001:1(route-map)#enable
```

14.2.1.4 Add Policy to BGP Peer Router-D

Add the policy to the BGP peer Router-D

```
9001:1(config)#router bgp
9001:1(router-bgp)#neighbor 10.1.1.41 out-route-map AS_Prepend
9001:1(router-bgp)#end
9001:1#ip bgp restart-bgp neighbor 10.1.1.41 soft-reconfiguration out
```

14.2.2 Configuration Example: AS_Path Filtering

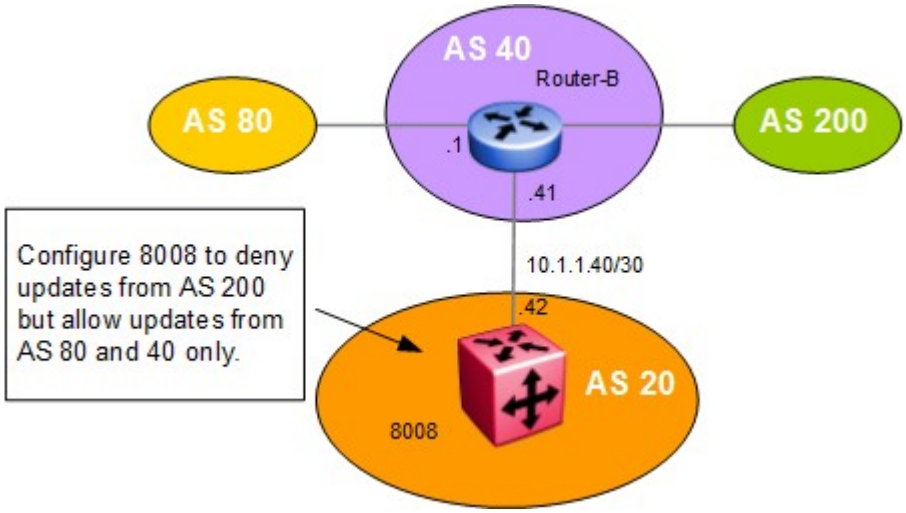

Figure 12: BGP AS Path Filtering Example

In this example, we will set up an access list to deny any updates on 8008 from AS 200 but allows updates from AS 40 and AS 80 only.

14.2.2.1 Configure an IP AS List on 8008

The following commands set up an access list that denies updates from AS200 but still allowing updates from AS40 and AS80

Create as-list 2 to deny AS 200, as-list 3 to allow AS 40, and as-list 4 to allow AS 80

```
8008:5(config)#ip as-list 2 memberid 1 deny as-path 200
8008:5(config)#ip as-list 3 memberid 1 permit as-path 40
8008:5(config)#ip as-list 4 memberid 1 permit as-path 80
```

14.2.2.2 Configure an IP Route Policy Named AS

Create a route policy named AS to match the as-list's above.

```
8008:5#(config)#route-map AS 1
8008:5#(route-map)#match as-path 2
8008:5#(route-map)#no permit
8008:5#(route-map)#enable
8008:5#(route-map)#exit
8008:5#(config)#route-map AS 2
8008:5#(route-map)#match as-path 3
8008:5#(route-map)#permit
8008:5#(route-map)#enable
8008:5#(route-map)#exit
```

```
8008:5#(config)#route-map AS 3
8008:5#(route-map)#match as-path 4
8008:5#(route-map)#permit
8008:5#(route-map)#enable
```

14.2.2.3 Add Policy to BGP Peer Router-B

Add the policy to the BGP peer Router-D

```
8008:5#(config)#router bgp
8008:5#(router-bgp)#neighbor 10.1.1.41 in-route-map AS
8008:5#(router-bgp)#end
8008:5#ip bgp restart-bgp neighbor 10.1.1.41 soft-reconfiguration in
```

14.2.3 Alternative Configuration Method for 8008

The above configuration example is just one method of AS Path configuration. The same configuration can also be accomplished by using the following commands:

14.2.3.1 Configure an IP AS List on 8008

Create as-list 2 to deny AS 200, allow AS 40 and AS 80.

```
8008:5(config)#ip as-list 2 memberid 1 deny as-path 200
8008:5(config)#ip as-list 2 memberid 2 permit as-path 40
8008:5(config)#ip as-list 2 memberid 3 permit as-path 80
```

14.2.3.2 Configure an IP Route Policy Named AS

Add sequence 1 to match as-list 2

```
8008:5#(config)#route-map AS 1
8008:5#(route-map)#match as-path 2
8008:5#(route-map)#no permit
8008:5#(route-map)#enable
```

14.2.3.3 Add Policy to BGP Peer Router-B

Add the policy to the BGP peer Router-D

```
8008:5#(config)#router bgp
8008:5#(router-bgp)#neighbor 10.1.1.41 in-route-map AS
8008:5#(router-bgp)#end
8008:5#ip bgp restart-bgp neighbor 10.1.1.41 soft-reconfiguration in
```

14.3 Local Preference Attribute (Type 5) Configuration Example

Local Preference is a well-known non-transitive attribute that influences the flow of outbound traffic by setting the exit point of an AS. Border routers within an AS calculate Local Preference if the attribute is not configured in a BGP accept policy.

The Local Preference attribute is local to ASs and is exchanged between iBGP peers only; e.g. it does not have any effect on the internal IGP protocol being used.

When BGP must select the best route and there are multiple paths to the same destination, the path with the larger preference is preferred.

In this example, we want to influence the traffic so the link from ERS8000-C to 8008 is used as the preferred path and the link from ERS8000-D to 9001 is used for back up only. 8008 is set with a higher Local Preference while 9001 is set for a lower local preference. We will also configure 8008 to inject a default route with a lower OSPF metric than 9001 resulting in 8008 having a higher preference. With this configuration, all traffic leaving AS 40 will exit via the customer AS will exit via ERS8000-C.

Local Preference can also be used to load balance outbound traffic based on CIDR or network address groups.

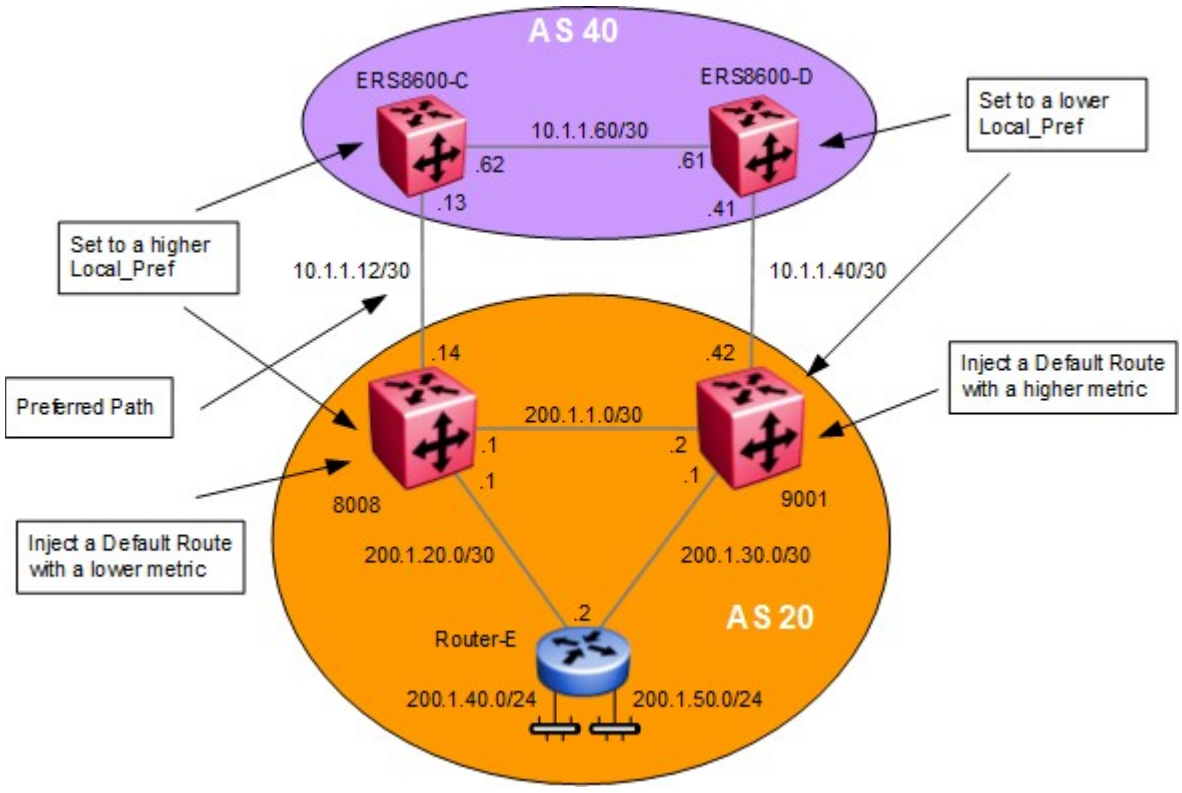

Figure 13: BGP Local Preference Configuration Example

14.3.1 Configuration : Local Preference

Please note that in regards to the OSPF configuration, this configuration example only provides the steps to add an OSPF route policy and enable distribution of BGP into OSPF. The interface and VLAN configuration steps are not included.

14.3.1.1 Local Preference Configuration for iBGP

Configure a local preference value for 8008 such that the value is higher that the local preference value you set for 9001 and configure a local preference value for ERS8000-C such that the value is higher that the local preference value you set for ERS8000-D

8008 & ERS8000-C: Same configuration on both switches

```
8008:5(config)#router bgp
8008:5(router-bgp)#bgp default local-preference 100
```

9001 & ERS8000-D: Same configuration on both switches

```
9001:1(config)#router bgp
9001:1(router-bgp)#bgp default local-preference 10
```

14.3.1.2 Configure the IP Prefix List for the Default Route on 8008 and 9001

Configure an prefix list, in this example named DR, and add the default route prefix of 0.0.0.0/0

8008 & 9001: Same configuration on both switches

```
8008:5(config)#ip prefix-list DR 0.0.0.0/0
```

14.3.1.3 Configure the IP Route Policy on 8008 and 9001

Configure the IP route policy named Default_OSPF, set the metric on 8001 to 100, and set the metric on 9001 to 300

8008:

```
8008:5(config)#route-map Default_OSPF 1
8008:5(route-map)#set injectlist DR
8008:5(route-map)#set metric 100
8008:5(route-map)#enable
```

9001:

```
9001:1(config)#route-map Default_OSPF 1
9001:1(route-map)#set injectlist DR
9001:1(route-map)#set metric 300
9001:1(route-map)#enable
```

The set-metric value directly influences the OSPF route decision. For this example, 8008 is set to a lower metric value than 9001, which results in a higher preference value.

14.3.1.4 Configure Route Redistribution to Redistribute BGP into OSPF Using the Route Policy to Inject a Default Route

Enable BGP redistribution into OSPF and apply the route policy from the previous step

8008 & 9001: Same configuration on both switches

```
8008:5(config)#router ospf
8008:5(config-ospf)#redistribute bgp
8008:5(config-ospf)#redistribute bgp route-map Default_OSPF
8008:5(config-ospf)#redistribute bgp enable
8008:5(config)#ip ospf apply redistribute bgp
```

14.4 Configuration Example: Adding Preference to Specific Routes

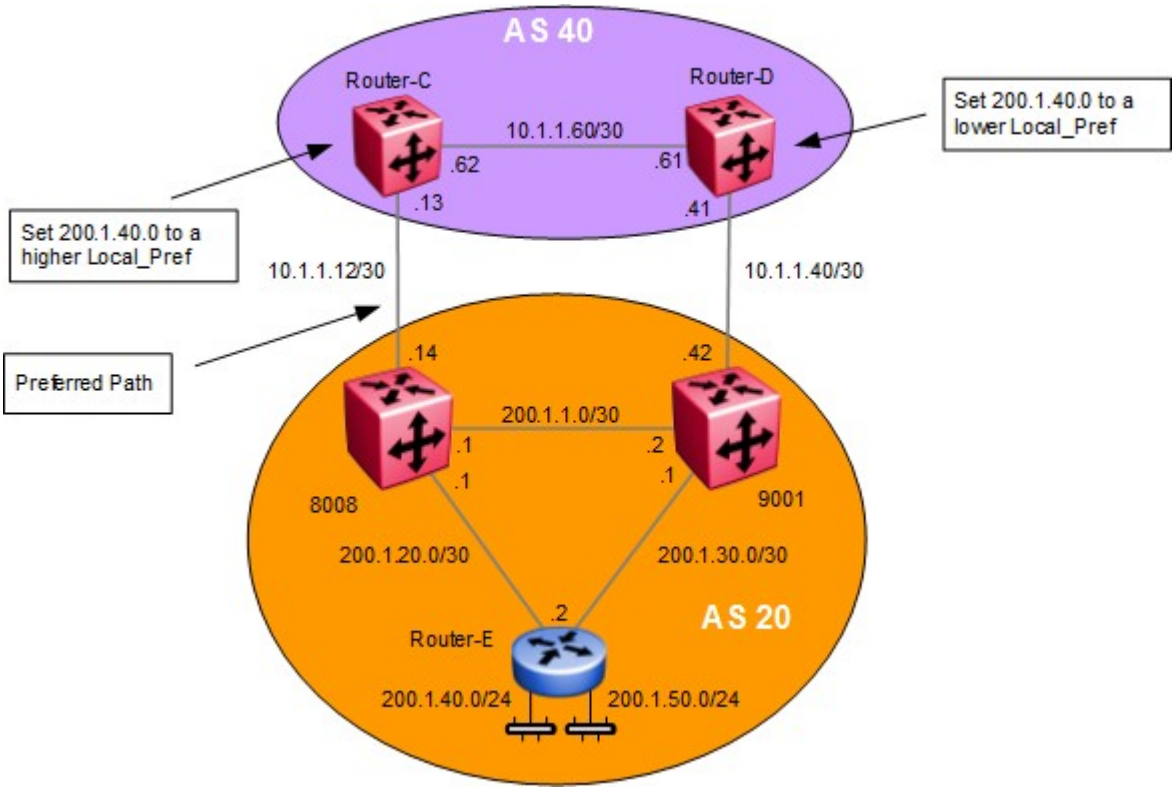

Figure 14: BFP Local Preference to Specific Routes Configuration Example

In the previous example, we configured the default local preference to influence all networks. As an alternative, the router can be configured with a route policy to influence specific networks. In this example, we wish to influence the traffic for network 200.1.40.0 to take the path between ERS8000-C and 8008. This can be accomplished by configuring a policy on ERS8000-C to have a higher Local Preference than ERS8000-D.

14.4.1 Configuration: Preference for Specific Routes

14.4.1.1 Configure the IP Prefix List Named 200.1.40.0 on Router-C and Router-D

Configure a prefix list named 200.1.40.0, and add the route prefix of 200.1.40.0/24

ERS8000-C & ERS8000-D: Same configuration on both switches

```
ERS8000-C:5(config)#ip prefix-list 200.1.40.0 200.1.40.0/24
```

14.4.1.2 Configure the IP Route Policy Named Policy with Sequence 1 to Match Prefix-List 200.1.40.0

Add a route policy to match the prefix list 200.1.40.0, set the local preference to 900 on Router-C and set the local preference to 200on Router-D

ERS8000-C:

```
ERS8000-C:5(config)#route-map Policy 1
ERS8000-C:5(route-map)#match network 200.1.40.0
ERS8000-C:5(route-map)#set local-preference 900
ERS8000-C:5(route-map)#enable
```

ERS8000-D:

```
ERS8000-D:5(config)#route-map Policy 1
ERS8000-D:5(route-map)#match network 200.1.40.0
ERS8000-D:5(route-map)#set local-preference 200
ERS8000-D:5(route-map)#enable
```

14.4.1.3 Assign the Route Policy to the Appropriate BGP Peer

Assign the route policy 'Policy' to the appropriate BGP peer

ERS8000-C:

```
ERS8000-C:5(config)#router bgp
ERS8000-C:5(router-bgp)#neighbor 10.1.1.14 in-route-map Policy
ERS8000-C:5(router-bgp)#ip bgp restart-bgp neighbor 10.1.1.14 soft-reconfiguration in
```

ERS8000-D:

```
ERS8000-D:5(config)#router bgp
ERS8000-D:5(router-bgp)#neighbor 10.1.1.42 in-route-map Policy
ERS8000-D:5(router-bgp)#ip bgp restart-bgp neighbor 10.1.1.42 soft-reconfiguration in
```

14.5 Multi-Exit Discriminator (MED) Attribute (Type 4)

The MED attribute is an optional non-transitive attribute that hints at preferred paths for routes that come from neighbors.

MEDs are only used with multiple connections to a neighboring AS in order to select a path for return traffic. A lower MED value indicates a stronger MED than a higher value.

One AS sets the MED value and a different AS uses that value to select a path. When an UPDATE message enters an AS with a MED value, the value is used to help the AS make routing decisions.

MED can also be used to load balance inbound traffic. For example, different MED values can be used for control different CIDR blocks.

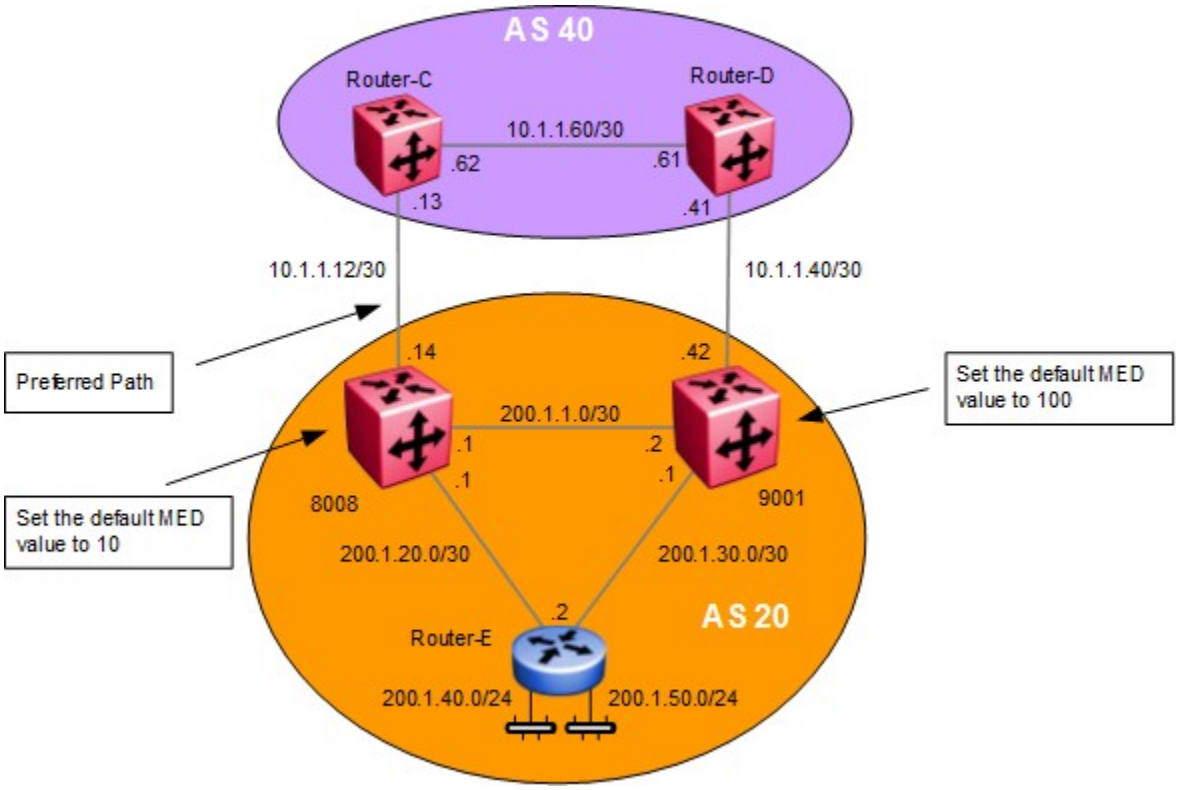

Figure 15: BGP MED Configuration Example

14.5.1 MED Configuration – Example 1

For this example, we will set the default MED setting on 8008 with a value of 10 and 9001 with a value of 100. The overall effect will result in 8008 advertising all routes with a MED setting of 10 whereas 9001 will advertise all routes with a MED setting of 100. This should result in all traffic destined for AS 20 to transverse over 10.1.1.12 network via 8008.

14.5.1.1 Set MED Value

Configure a MED value for 8008 to a value that is lower than the MED value assigned to 9001.

Set the MED value

8008:

```
8008:5(config)#router bgp
8008:5(router-bgp)#default-metric 10
```

9001:

```
9001:1(config)#router bgp
9001:1(router-bgp)#default-metric 100
```

A lower MED value indicates a stronger path preference than a higher MED value.

14.6 MED Configuration – Example 2

For this example, we will set a policy so that 8008 will advertise the 200.1.40.0/24 network with a MED setting of 10 and 200.1.50.0/24 with a MED setting of 100. 9001 will be configured with a policy to advertise 200.1.40.0/24 with a MED setting of 100 and 200.1.50.0/24 with a MED setting of 10. This should result in all traffic destined for network 200.1.40.0/24 to transverse over 8008 while 200.1.50.0/24 will transverse over 9001.

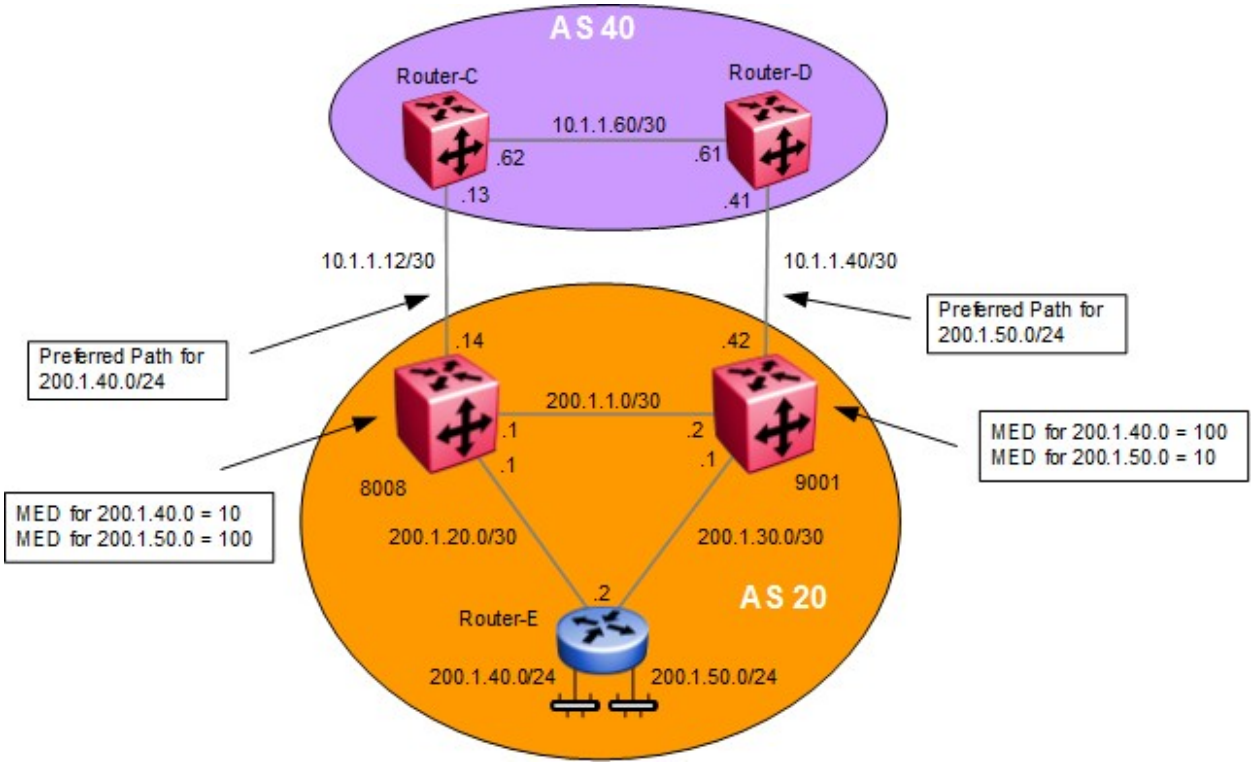

Figure 16: BFP MED Configuration Example 2

14.6.1 Configuration

14.6.1.1 Configure the IP Prefix List Named 200.1.40.0 and 200.1.50.0

Add a prefix list and network

8008 & 9001: Same configuration on both switches

```
8008:5(config)#ip prefix-list 200.1.40.0 200.1.40.0/24
8008:5(config)#ip prefix-list 200.1.50.0 200.1.50.0/24
```

14.6.1.2 Configure the IP Route Policy Name MED

Step 1: Add a route policy named MED to match network 200.1.40.0 and set the MED value to 10 on 8008 and 100 on 9001

8008:

```
8008:5(config)#route-map MED 1
8008:5(route-map)#match network 200.1.40.0
8008:5(route-map)#set metric 10
8008:5(route-map)#enable
```

9001:

```
9001:1(config)#route-map MED 1
9001:1(route-map)#match network 200.1.40.0
9001:1(route-map)#set metric 100
9001:1(route-map)#enable
```

Step 2: Add a second sequence to the policy named MED to match network 200.1.50.0 and set the MED value to 100 on 8008 and 10 on 9001

8008:

```
8008:5(config)#route-map MED 2
8008:5(route-map)#match network 200.1.50.0
8008:5(route-map)#set metric 100
8008:5(route-map)#enable
```

9001:

```
9001:1(config)#route-map MED 2
9001:1(route-map)#match network 200.1.50.0
9001:1(route-map)#set metric 10
9001:1(route-map)#enable
```

14.6.2 Other MED Commands

14.6.2.1 Always Compare MED

In the configuration examples in this section, the *AlwaysCompareMed* setting was left to the default setting of disable. When disabled, the MEDs are only compared among paths from the same autonomous system. In the two examples above, since we only have two autonomous systems, the default setting can be used. If you have multiple autonomous systems, this parameter should be enabled to allow MED to compare among paths among multiple autonomous systems. To enable or disable this parameter, enter the following command:

```
8008:5(config)#router bgp
8008:5(config)#bgp always-compare-med
```

14.6.2.2 Deterministic-med

The BGP deterministic MED command is used to compare MED variable of different routes that are advertised by peers in the same autonomous system (AS). When this command is enabled, only routes from the same AS are considered, and the route with the lowest MED is used.

The CLI syntax for this command is:

```
8008:5(config)#router bgp
8008:5(config)#bgp deterministic-med enable
```

By default it is disabled. This feature is not supported in Device Manager.

Example

Consider the following routes received for the network 200.1.40.0/24 from different peers:

route-1 :AS Path (40), Peer 200.1.1.1, MED 100

route-2 :AS Path (50), Peer 172.10.1.1, MED 110

route-3 :AS Path (40), Peer 10.12.1.4, MED 80

In this example, route-2 is from a different AS, and is the only one in its group. Thus, it is chosen as the best from its group. Route-1 and route-3 are grouped together, and route-3 is chosen because it has the lowest MED. In the comparison between route-2 and route-3, the MED is ignored, and the best entry is chosen based on other factors.

14.6.2.3 No MED Path is Worst

When set to enable (the default value), BGP treats an update that is missing, a multi-exit discriminator (MED) attribute, as the worst path. To enable or disable this parameter, enter the following command:

```
8008:5(config)#router bgp
8008:5(router-bgp)#no no-med-path-is-worst
```

14.6.2.4 MED Compare within a Confederation

When enabled, allows you to compare multi-exit discriminator (MED) attributes within a confederation. The default value is disabled. To enable to disable this parameter, enter the following command:

```
8008:5(config)#router bgp
8008:5(router-bgp)#comp-bestpath-med-confed enable
8008:5(router-bgp)#ip bgp restart-bgp
```

14.7 Community Attribute (Type 8)

Community is an optional transitive attribute that groups destinations into communities to simplify policy administration in a BGP network. A community is a group of destinations that share a common administrative property.

With the Community attribute, customers can control their own routing policies with respect to destinations. Communities are a common practice in cases where a customer has more than one destination and wishes to share some common attribute.

The following are specific community types:

- *Internet* – Advertise this route to the Internet community.
- *No-export* – do not advertise any destinations outside of a BGP confederation
- *No Advertise* – do not advertise to any BGP peer including IBGP peers
- *No Export Subconfed* – do not advertise to external BGP peers even within the same confederation.

For the community type ‘no export subconfed’, the ERS 8000 uses a setting of ‘local-as’.

By using the community attribute, you can control what routing information to accept, prefer, or distribute to other BGP neighbors. If you specify the append option in the route policy, the specified community value is added to the existing value of the community attribute. Otherwise, the specified community value replaces any community value that was set previously.

14.7.1 Community Attribute Configuration Example

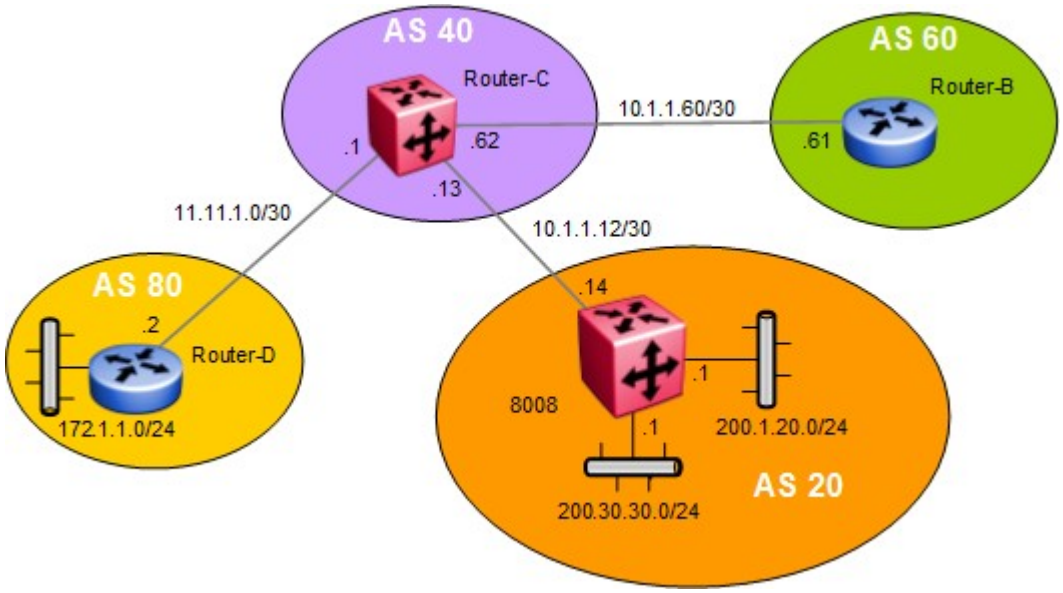

Figure 17: BGP Community Configuration Example

The ERS and VSP switches uses an IP Community List policy to specify the community. In this example, we will configure 8008 with a community attribute of 'no-export' to Router-C in AS 40 for network 200.30.30.0. For all other networks, we will configure 8008 with a community attribute of 'internet'. This will indicate ERS8000-C to not propagate the 200.30.30.0 network but advertise all other routes learned from 8008.

To enable the BGP peer to send the community attribute, use the following commands:

14.7.1.1 Configure the IP Prefix List Named 200.30.30.0

Add a prefix list for network 200.30.30.0/24

```
8008:5(config)ip prefix-list 200.30.30.0 200.30.30.0/24
```

14.7.1.2 Configure the IP Community Lists

Add community list 1 using community string 55:55 and with a community attribute of no-export and community list 2 using community string 55:55 with a community attribute of internet

```
8008:5(config)#ip community-list 1 memberid 1 permit community-string 55:55
8008:5(config)#ip community-list 1 memberid 2 permit community-string no-export
8008:5(config)#ip community-list 2 memberid 1 permit community-string 55:55
8008:5(config)#ip community-list 2 memberid 2 permit community-string internet
```

The community-string is an alphanumeric string value with a string length between 0 and 1536 characters (asnum:community-value) or (well-known community string).

14.7.1.3 Configure the IP Route policy name community

Step 1: Add a route policy named community to match network 200.30.30.0 and add community list 1 with a mode of additive. This will have the effect of announcing network 200.30.30.0/24 with a community attribute of no-export.

```
8008:5(config)#route-map community 1
8008:5(route-map)#match network 200.30.30.0
8008:5(route-map)#set community 1
8008:5(route-map)#set community-mode additive
8008:5(route-map)#enable
```

Step 2: Add sequence 2 to policy named community with a community mode of additive. This will have the effect of announcing all other routes with a community attribute of internet

```
8008:5(config)#route-map community 2
8008:5(route-map)#set community 2
8008:5(route-map)#set community-mode additive
8008:5(route-map)#enable
```

The following options are available for the set-community-mode in a route policy:

- ***config ip route-policy <name> seq <#> set-community-mode <unchanged|additive|none>***
  - ***unchanged*** — do not change an existing community
  - ***additive*** — append the community to the exiting community
  - ***none*** — remove the community

14.7.1.4 Assign ERS8000-C as a Peer to 8008 Enable Community

Send community to ERS8000-C

```
8008:5(config)#router bgp
8008:5(router-bgp)#no neighbor 10.1.1.13 enable
8008:5(router-bgp)#neighbor 10.1.1.13 send-community enable
8008:5(router-bgp)#neighbor 10.1.1.13 out-route-map community
8008:5#router-bgp)#neighbor 10.1.1.13 enable
```

14.7.2 Verification

If we look at the route table on ERS8000-C:

To verify that the BGP community operation, via ERS8000-C, enter the following command:

```
ERS8000-C:5#show ip bgp route community enable
The total number of routes is 11

 Network/Mask Peer Rem Addr NextHop Address      Org    Loc    Pref
-------------------------------------------------------------------
16.16.16.16/30      10.1.1.61            10.1.1.61           IGP    100
        AS_PATH: (60)
        COMMUNITY: no-community-attr
10.1.1.12/30        10.1.1.14            10.1.1.14           IGP    100
        AS_PATH: (20)
        COMMUNITY: 40:100 internet
14.14.14.0/24       10.1.1.61            10.1.1.61           IGP    100
        AS_PATH: (60)
        COMMUNITY: no-community-attr
10.1.1.40/30        10.1.1.61            10.1.1.61           IGP    100
        AS_PATH: (60)
        COMMUNITY: no-community-attr
172.1.2.0/30        11.11.1.2            11.11.1.2           IGP    100
        AS_PATH: (80)
        COMMUNITY: no-community-attr
192.1.1.0/24        10.1.1.61            10.1.1.61           IGP    100
        AS_PATH: (60 200)
        COMMUNITY: no-community-attr
200.30.30.0/24      10.1.1.14            10.1.1.14           IGP    100
        AS_PATH: (20)
        COMMUNITY: 40:100 no-export
200.1.1.4/30        10.1.1.61            10.1.1.61           IGP    100
        AS_PATH: (60 20)
        COMMUNITY: no-community-attr
200.1.30.0/30       10.1.1.61            10.1.1.61           IGP    100
        AS_PATH: (60 20)
        COMMUNITY: no-community-attr
200.1.20.0/24       10.1.1.14            10.1.1.14           IGP    100
```

```
        AS_PATH: (20)
        COMMUNITY: 40:100 internet
```

The end result is, network 200.30.30.0 will not be advertised outside AS 40 while the 200.1.20.0 will be advertised outside AS 40.

15. EBGP Scalability Issues

15.1 Using Policies to Limit EBGP Routes

Route policies can be used to limit the number of EBGP routes to reduce the number of BGP forwarding routes. Within a customer AS, a default route can be injected into the local IGP by the EBGP peering router or routers to reduce the route table size for all non-BGP routers. Please note that the ERS8000 supports up to 250,000 BGP forwarding routes in its routing information base (RIB) and 500,000 in its forwarding information base (FIB). The VSP9000 as of release 4.0 and using generation 2 modules can support up to 1 million route operations.

15.1.1 Configuration Example: Using AS List to Limit Route Table Size

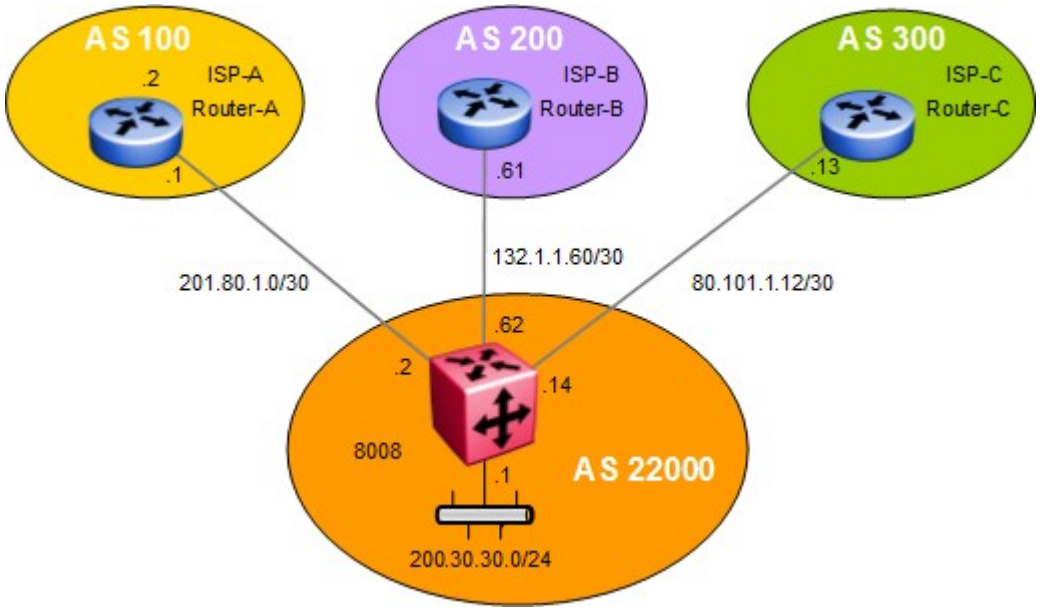

Figure 18: BGP AS Path Filtering Example

8008 is connected to three separate AS's all supplying full internet routes. For this example, we wish to reduce the number of routes accepted from each EBGP connection on switch 8008. One method to do this is to apply a policy to each EBGP connection with an IP prefix list to suppress routes. For example, if we add a route policy with an IP prefix list mask length of 18 bits to each EBGP connection, this will reduce the number of routes on each connection to approximately 22,000 routes for a total of 66,000 routes.

For this example, three route policies will be added. We will add a route policy to inject a default route into AS 22000 assuming the IGP is OSPF. Another route policy will be added and applied to each EBGP peer to reduce the number of routes learned by applying an IP prefix with a mask length of 18. In addition, we will add another policy to block 8008 from sending out any routing information from any of the three EBGP peers; we do not want 8008 from becoming a transit router for the other ASs.

Please note that the VLAN and/or brouter port and OSPF configuration is not provided for this configuration example.

15.1.1.1 Configure an IP Prefix List

The following commands add an IP prefix list for a default route and another to limit the prefix mask length to 18 bits

Create prefix lists

```
8008:5(config)#ip prefix-list default 0.0.0.0/0
8008:5(config)#ip prefix-list Limit_18 0.0.0.0/0 ge 0 le 18
```

15.1.1.2 Configure AS List

Add access list used to deny AS 100, 200, and 300

```
8008:5(config)#ip as-list 1 memberid 1 deny as-path "100"
8008:5(config)#ip as-list 1 memberid 2 deny as-path "200"
8008:5(config)#ip as-list 1 memberid 3 deny as-path "300"
```

15.1.1.3 Configure IP Route Policies

Add a route map named DR to inject a default route

```
8008:5(config)#route-map "DR" 1
8008:5(route-map)#enable
8008:5(route-map)#set injectlist "default"
8008:5(route-map)#exit
```

Add a route map named sub_18 to match on the IP prefix list named limit_18 and match on protocol EBGP

```
8008:5(config)#route-map "sub_18" 1
8008:5(route-map)#enable
8008:5(route-map)#match network "limit_18"
8008:5(route-map)#match protocol ebgp
8008:5(route-map)#exit
```

Add a route map named as_out to deny routes learned from AS 100, 200 and 300 by matching AS list 1

```
8008:5(config)#route-map "as_out" 1
8008:5(route-map)#no permit
8008:5(route-map)#enable
8008:5(route-map)#match as-path 1
8008:5(route-map)#route-map "as_out" 2
8008:5(route-map)#enable
8008:5(route-map)#exit
```

15.1.1.4 Configure BGP Globally on 8008

Assign both switch to BGP AS2200, disable synchronization, and add BGP network

```
8008:5(config)#router bgp 20 enable
8008:5(config)#router bgp
8008:5(router-bgp)#no synchronization
8008:5(router-bgp)#network 200.30.30.0/24
8008:5(router-bgp)#exit
```

15.1.1.5 Add BGP Peers

Add BGP peers. Add in-route-maps *sub_18* and out-route-map *as_out*

```
8008:5(config)#router bgp
8008:5(router-bgp)#no synchronization
8008:5(router-bgp)#neighbor "80.101.1.13"
8008:5(router-bgp)#neighbor "132.1.1.61"
8008:5(router-bgp)#neighbor "201.80.1.1"
8008:5(router-bgp)#neighbor 80.101.1.13 remote-as 100
8008:5(router-bgp)#neighbor 80.101.1.13 max-prefix 0
8008:5(router-bgp)#neighbor 80.101.1.13 in-route-map "sub_18"
8008:5(router-bgp)#neighbor 80.101.1.13 out-route-map "as_out"
8008:5(router-bgp)#neighbor 132.1.1.61 remote-as 100
8008:5(router-bgp)#neighbor 132.1.1.61 max-prefix 0
8008:5(router-bgp)#neighbor 132.1.1.61 in-route-map "sub_18"
8008:5(router-bgp)#neighbor 132.1.1.61 out-route-map "as_out"
8008:5(router-bgp)#neighbor 201.80.1.1 remote-as 100
8008:5(router-bgp)#neighbor 201.80.1.1 max-prefix 0
8008:5(router-bgp)#neighbor 201.80.1.1 in-route-map "sub_18"
8008:5(router-bgp)#neighbor 201.80.1.1 out-route-map "as_out"
8008:5(router-bgp)#exit
```

15.1.1.6 Enable OSPF Redistribution of BGP

Enable BGP redistribution into OSPF and add route-policy *DR*

```
8008:5(config)#router ospf
8008:5(config-ospf)#redistribute bgp
8008:5(config-ospf)#redistribute bgp route-policy "DR"
8008:5(config-ospf)#redistribute bgp enable
8008:5(config-ospf)#exit
8008:5(config)#ip ospf apply redistribute bgp
```

16. IBGP Scalability Issues

In order to preserve and update BGP attributes, IBGP connections between border routers must be “fully-meshed”. Any external routing information must be re-distributed to all other routers within the AS. As the number of IBGP speaker’s increases, this full mesh requirement does not scale very well. With many border routers and 1,000s of routes, IBGP peering can become an issue for resources such as CPU, bandwidth, and configuration management.

Because of scalability, BGP speakers within an AS must maintain n*(n-1)/2 unique IBGP sessions.

Route reflectors and BGP Confederations can be used to eliminate the full-mesh scaling problem by minimizing the number of necessary peer sessions.

16.1 BGP Confederations

Confederations reduce the number of peers required within the AS. BGP confederations are used to divide an AS into multiple smaller ASs and assign these sub-system ASs to a confederation. The IBGP speakers within the sub-system AS only need to establish peer sessions with the other speakers in their own sub-system and one speaker from each sub-system establishes EBGP peer sessions with a speaker from each of the other sub-systems.

Although there is multiple smaller sub system ASs with the BGP confederation, to the outside world, the confederation looks like a single AS.

16.2 Confederation Configuration Example

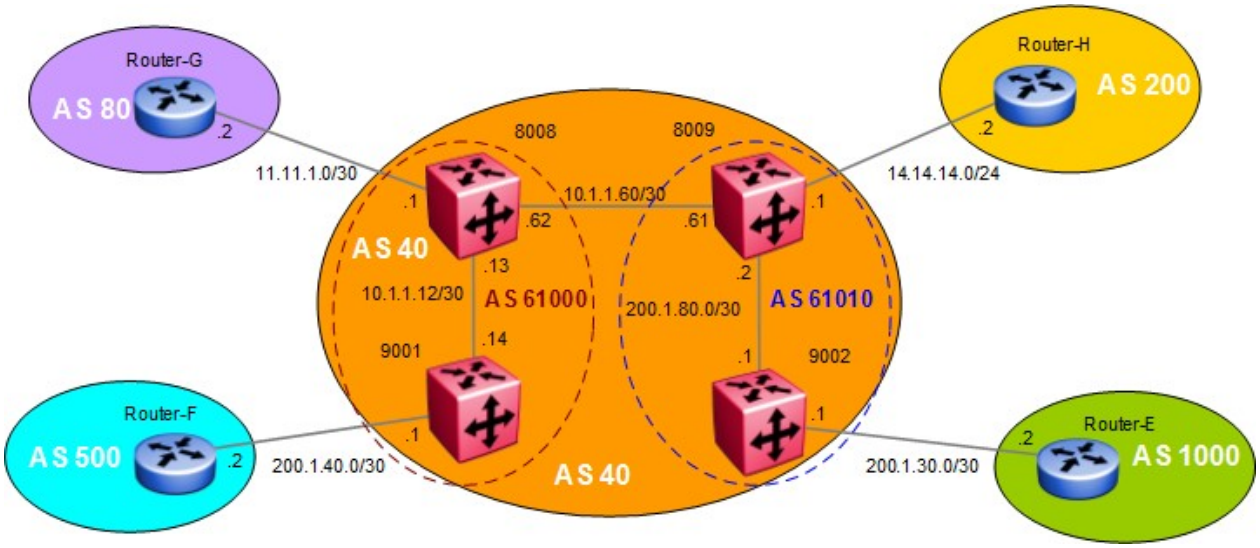

Figure 19: BGP Confederation Configuration Example

In this example, confederations are used to reduce the number of IBGP peers. Without confederations, all the routers in AS 40 must be fully meshed. Confederations will reduce the number of peers within the AS by dividing AS into multiple smaller confederation ASs.

All routers within the confederation AS are fully meshed. Each confederation AS has a connection to the other confederation ASs and use EBGP to exchange routing updates. Even though EBGP is used between confederation ASs, the routing information exchanged is treated as if they are using IBGP. This preserves all the various IBGP information such as local preference and MED.

16.2.1 Configuration

16.2.1.1 Configure BGP Confederation

Assign local AS 40 as the confederation identifier and add the peer AS 61010 to the confederation configuration

8008 & 8009: Same configuration on both switches

```
8008:5(config)#router bgp
8008:5(router-bgp)#bgp confederation identifier 40
8008:5(router-bgp)#bgp confederation peers 61010
```

Assign local AS 40 as the confederation identifier

9001 & 9002: Same configuration on both switches

```
9001:1(config)#router bgp
9001:1(router-bgp)#bgp confederation identifier 40
```

16.2.1.2 Enable BGP

Enable BGP

8008, 8009, 9001, and 9002: Same configuration on all switches

```
8008:5(config)#router bgp 61000 enable
```

16.2.1.3 Assign BGP Peers

Enable BGP peers

8008:

```
8008:5(config)#router bgp
8008:5(router-bgp)#neighbor 11.11.1.2
8008:5(router-bgp)#neighbor 11.11.1.2 remote-as 80
8008:5(router-bgp)#neighbor 11.11.1.2 enable
8008:5(router-bgp)#neighbor 10.1.1.61
8008:5(router-bgp)#neighbor 10.1.1.61 remote-as 61010
8008:5(router-bgp)#neighbor 10.1.1.61 enable
8008:5(router-bgp)#neighbor 10.1.1.14
8008:5(router-bgp)#neighbor 10.1.1.14 remote-as 61000
8008:5(router-bgp)#neighbor 10.1.1.14 enable
```

9001:

```
9001:1(config)#router bgp
9001:1(router-bgp)#neighbor 10.1.1.13
9001:1(router-bgp)#neighbor 10.1.1.13 remote-as 61000
9001:1(router-bgp)#neighbor 10.1.1.13 enable
9001:1(router-bgp)#neighbor 200.1.40.2
9001:1(router-bgp)#neighbor 200.1.40.2 remote-as 500
9001:1(router-bgp)#neighbor 200.1.40.2 enable
```

8009:

```
8009:5(config)#router bgp
8009:5(router-bgp)#neighbor 14.14.14.2
8009:5(router-bgp)#neighbor 14.14.14.2 remote-as 200
8009:5(router-bgp)#neighbor 14.14.14.2 enable
8009:5(router-bgp)#neighbor 10.1.1.62
8009:5(router-bgp)#neighbor 10.1.1.62 remote-as 61000
8009:5(router-bgp)#neighbor 10.1.1.62 enable
8009:5(router-bgp)#neighbor 200.1.80.1
8009:5(router-bgp)#neighbor 200.1.80.1 remote-as 61010
8009:5(router-bgp)#neighbor 200.1.80.1 enable
```

9002:

```
8009:5(router-bgp)#neighbor 200.1.80.1
8009:5(router-bgp)#neighbor 200.1.80.1 remote-as 61010
8009:5(router-bgp)#neighbor 200.1.80.1 enable
9002:1(router-bgp)#neighbor 200.1.30.2
9002:1(router-bgp)#neighbor 200.1.30.2 remote-as 1000
9002:1(router-bgp)#neighbor 200.1.30.2 enable
```

16.3 Route Reflectors

Route reflectors are another alternative to reduce the number of IBGP peering within an AS. Route reflectors allow routers to advertise or reflect IBGP routes to other IBGP speakers.

The internal peers of Route Reflectors are divided into two groups, client peers and non-client peers. A route reflector reflects routes between these two groups. The non-client peer must be fully meshed while the client peers do not need to be fully meshed.

16.3.1 Route Reflector Configuration Example

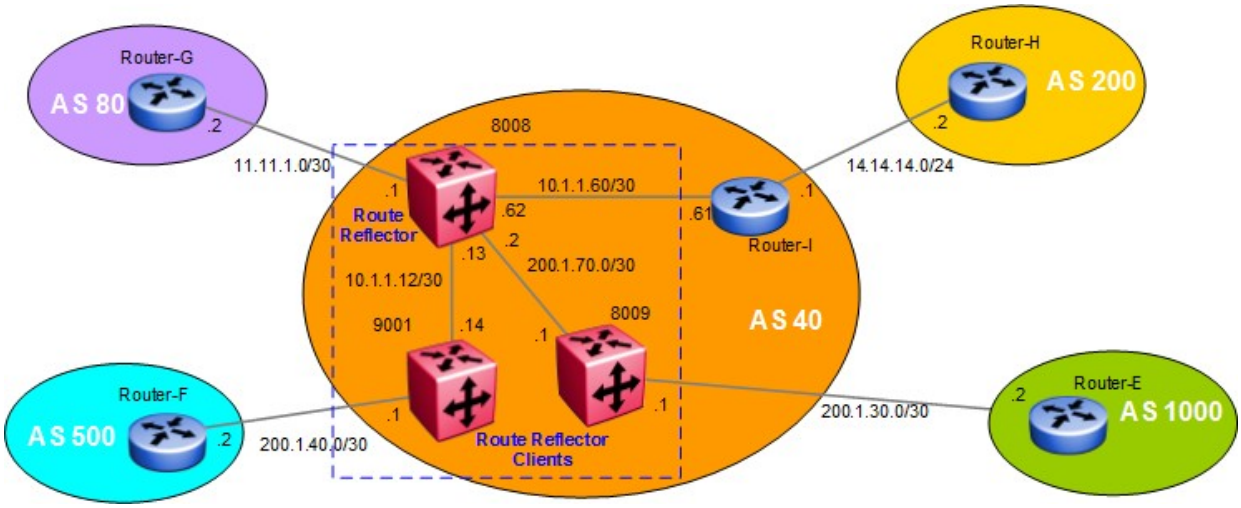

Figure 20: BGP Route Reflector Configuration Example

Without a route reflector, all routers in AS 40 will require full IBGP mesh. For example, ERS8000-C will require IGMP peering with 9001, 8008, and Router I. With route reflection configured on 8008, IBGP peering on ERS8000-C is no longer required to 9001 and Router-I.

The router whose configuration includes a route reflector also includes the router reflector client configuration. The route reflector can also be configured to allow or not allow routes learned by a client to be forwarded to other clients. A route reflector and all its clients as a whole are called a cluster. Other IBGP peers of the route reflector that are not route reflector clients are called non-clients. In this example, 8008 is the route reflector. 9001 and ERS8000-C are route reflector clients while Router-I is a non-client.

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

16.3.1.1 Configure BGP Local AS and enable BGP

Configure BGP local AS 40

8008, 9001, and 8009: Same configuration on all switches

```
8008:5(config)#router bgp 40 enable
```

16.3.1.2 Disable BGP Synchronization

Disable Synchronization

8008, 9001, and 8009: Same configuration on all switches

```
8008:5(config)#router bgp
8008:5(router-bgp)#no synchronization
```

16.3.1.3 Route Reflection Configuration

Enable route reflector client to client route reflection on 8008

8008:

```
8008:5(config)#router bgp
8008:5(router-bgp)#bgp client-to-client reflection
8008:5(router-bgp)#route-reflector enable
```

16.3.1.4 Assign BGP Peers

Add BGP peers and enable route reflector client on 8008 for the 9001 and 8009 peers

8008:

```
8008:5(router-bgp)#neighbor 10.1.1.14
8008:5(router-bgp)#neighbor 10.1.1.14 remote-as 40
8008:5(router-bgp)#neighbor 10.1.1.14 route-reflector-client
8008:5(router-bgp)#neighbor 10.1.1.14 route-reflector-client enable
8008:5(router-bgp)#neighbor 200.1.70.1
8008:5(router-bgp)#neighbor 200.1.70.1 remote-as 40
8008:5(router-bgp)#neighbor 200.1.70.1 route-reflector-client
8008:5(router-bgp)#neighbor 200.1.70.1 route-reflector-client enable
8008:5(router-bgp)#neighbor 11.11.1.2
8008:5(router-bgp)#neighbor 11.11.1.2 remote-as 80
8008:5(router-bgp)#neighbor 11.11.1.2 enable
8008:5(router-bgp)#neighbor 10.1.1.61
8008:5(router-bgp)#neighbor 10.1.1.61 remote-as 40
```

```
8008:5(router-bgp)#neighbor 10.1.1.61 enable
```

9001:

```
9001:1(router-bgp)#neighbor 10.1.1.13
9001:1(router-bgp)#neighbor 10.1.1.13 remote-as 40
9001:1(router-bgp)#neighbor 10.1.1.13 enable
```

8009:

```
8009:5(router-bgp)#neighbor 200.1.70.2
8009:5(router-bgp)#neighbor 200.1.70.2 remote-as 40
8009:5(router-bgp)#neighbor 200.1.70.2 enable
8009:5(router-bgp)#neighbor 200.1.30.2
8009:5(router-bgp)#neighbor 200.1.30.2 remote-as 1000
8009:5(router-bgp)#neighbor 200.1.30.2 enable
```

16.4 Configuration Example using Cluster List

Normally, in a route reflector cluster there is only one route reflector and is identified by the router ID. To increase resilience, a second route reflector can be installed. When installing more than one route reflector in a cluster, the cluster must be configured with a 4-octet cluster ID. The cluster ID allows the router reflectors to recognize updates from other route reflectors in the same cluster. The cluster ID is also appended to all routes sent outside its cluster. If a route reflector receives an update with a cluster ID the same as the local customer ID the update is dropped, hence preventing route loops.

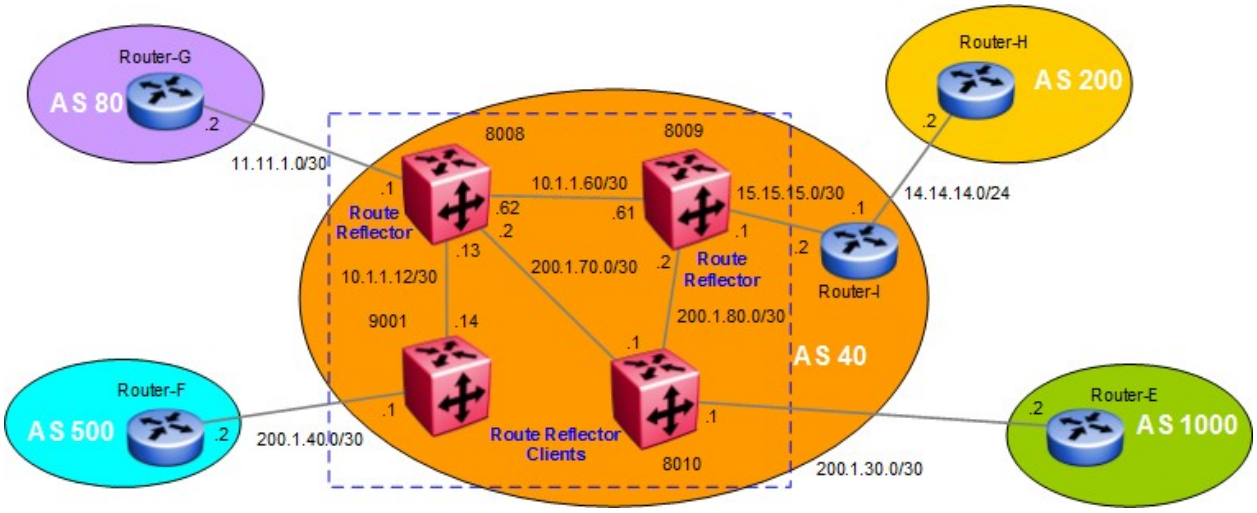

Figure 21: BGP Route Reflector with Cluster List Configuration Example

16.4.1 Configuration

16.4.1.1 Configure BGP Local AS and enable BGP

Create prefix lists

8008, 8009, 8010, and 9001: Same configuration on all switches

```
8008:5(config)#router bgp 40 enable
```

16.4.1.2 Disable Synchronization

Disable Synchronization

8008, 8009, 8010, and 9001: Same configuration on all switches

```
8008:5(config)#router bgp
8008:5(router-bgp)#no synchronization
```

16.4.1.3 Route Reflection Configuration

Disable Synchronization

8008, 8009: Same configuration on both switches

```
8008:5(config)#router bgp
8008:5(router-bgp)#bgp client-to-client reflection
8008:5(router-bgp)#bgp cluster-id 0.0.0.20
8008:5(router-bgp)#route-reflector enable
```

16.4.1.4 Assign BGP Peers

Assign BGP peers

8008: For neighbors 8010 and 9001, enable the peer as a route reflector client

```
8008:5(router-bgp)#neighbor 100.1.1.14
8008:5(router-bgp)#neighbor 100.1.1.14 remote-as 40
8008:5(router-bgp)#neighbor 100.1.1.14 route-reflector-client
8008:5(router-bgp)#neighbor 100.1.1.14 route-reflector-client enable
8008:5(router-bgp)#neighbor 200.1.70.1
8008:5(router-bgp)#neighbor 200.1.70.1 remote-as 40
8008:5(router-bgp)#neighbor 200.1.70.1 route-reflector-client
8008:5(router-bgp)#neighbor 200.1.70.1 route-reflector-client enable
8008:5(router-bgp)#neighbor 11.11.1.2
8008:5(router-bgp)#neighbor 11.11.1.2 remote-as 80
8008:5(router-bgp)#neighbor 11.11.1.2 enable
```

8009: For neighbor 8010, enable the peer as a route reflector client

```
8009:5(router-bgp)#neighbor 10.1.1.62
8009:5(router-bgp)#neighbor 10.1.1.62 remote-as 40
8009:5(router-bgp)#neighbor 10.1.1.62 enable
8009:5(router-bgp)#neighbor 200.1.80.1
8009:5(router-bgp)#neighbor 200.1.80.1 remote-as 40
8009:5(router-bgp)#neighbor 200.1.80.1 route-reflector-client
8009:5(router-bgp)#neighbor 200.1.80.1 route-reflector-client enable
8009:5(router-bgp)#neighbor 15.15.15.2
8009:5(router-bgp)#neighbor 15.15.15.2 remote-as 40
8009:5(router-bgp)#neighbor 15.15.15.2 enable
```

8010:

```
8010:5(router-bgp)#neighbor 200.1.70.2
8010:5(router-bgp)#neighbor 200.1.70.2 remote-as 40
8010:5(router-bgp)#neighbor 200.1.70.2 enable
8010:5(router-bgp)#neighbor 200.1.30.2
8010:5(router-bgp)#neighbor 200.1.30.2 remote-as 1000
8010:5(router-bgp)#neighbor 200.1.30.2 enable
8010:5(router-bgp)#neighbor 200.1.80.2
8010:5(router-bgp)#neighbor 200.1.80.2 remote-as 40
8010:5(router-bgp)#neighbor 200.1.80.2 enable
```

9001:

```
9001:1(router-bgp)#neighbor 10.1.1.13
9001:1(router-bgp)#neighbor 10.1.1.13 remote-as 40
9001:1(router-bgp)#neighbor 10.1.1.13 enable
9001:1(router-bgp)#neighbor 200.1.40.2
9001:1(router-bgp)#neighbor 200.1.40.2 remote-as 500
9001:1(router-bgp)#neighbor 200.1.40.2 enable
```

17. Configuring EBGP Route Flap Dampening

The frequent change of network reachability information that can be caused by an unstable route is commonly referred to as route flap. Route flap dampening is a technique for suppressing information about unstable routes.

NOTE: Dampening is only applied to routes learned via EBGP. This prevents routing loops and prevents IBGP peers having a higher penalty for routes that are external to the AS.

The diagram below demonstrates route flap from 8008 and Router-C.

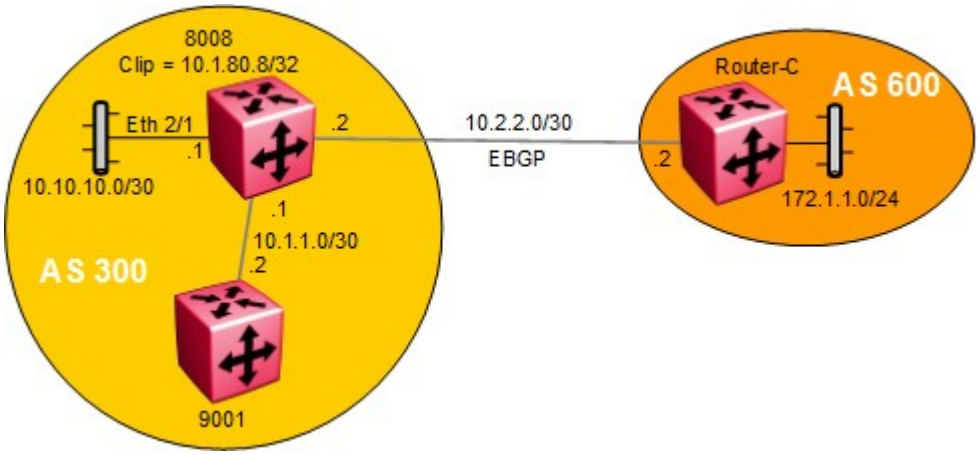

Figure 22: BGP Route Flap Damping Configuration Example

17.1 Configuration: Route Flap Damping

17.1.1 Enabling BGP Route Flap Damping

To enable BGP route flap damping on 8008, use the following command

```
8008:5(config)#router bgp
8008:5(router-bgp)#flap-dampening enable
```

17.2 Verification

17.2.1 Viewing Damping Configuration

Step 1: After enabling flap-damping on 8008, you can view the damping configuration by performing the following command:

```
8008:5#show ip bgp flap-damp-config
================================================================================
                        BGP Flap Dampening - GlobalRouter
================================================================================
                                          Status - enable
                                      PolicyName - N/A
                                 CutoffThreshold - 1536
                                  ReuseThreshold - 512
                                           Decay - 2
                                     MaxHoldDown - 180
```

Step 2: Initially, assuming the EBGP link is stable between 8008 and ERS8000-C, performing the following command should display no dampened paths:

```
8008:5#show ip bgp dampened-paths 10.2.2.2
================================================================================
                        BGP Dampened Paths - GlobalRouter
================================================================================
The total number of hist routes:  0

NETWORK/MASK       PEER REM ADDR   NEXTHOP ADDRESS ORG LOC PREF
--------------------------------------------------------------------------------
```

Step 3: Assuming the link to Router-C goes from an up to down and to up state

```
8008:5#show ip bgp dampened-paths 10.2.2.2
================================================================================
                        BGP Dampened Paths - GlobalRouter
================================================================================
The total number of hist routes:  0

NETWORK/MASK       PEER REM ADDR   NEXTHOP ADDRESS ORG LOC PREF
--------------------------------------------------------------------------------
172.1.1.0/24       10.2.2.2        N/A             IGP 0
        AS_PATH: no-AS_PATH-attr
        MED:0
        DAMPEN INFO:Penalty:1024  Count:1  Status:announced  hist-del time:set:180,
remain:173
```

Via 8008, verify the following information:

Option	Verify
Count	The value shown indicates the number of times this route has flapped. If we perform another admin-state disable/enable on ERS8000-C, the count will go up to 2.
Remain	This value indicates the amount of hold down time left. Notice when we displayed the flap damped configuration, the maximum hold down time displayed a value of 180 seconds – this is initially set upon a new route flap. This counter will continue to count down to zero unless of course there is another flap in which case the counter will go back up to 180 and count down again.
Penalty	If the penalty count is greater than the Cut Off Threshold, the route will be suppressed even if the route is up.

17.3 BGP Quick-Start Feature

The quick-start feature which, when enabled, avoids flap penalty on the peers. This feature will force the peers to transition from IDLE state immediately when the switch is reset. This will result in no flap penalty imposed on the peers. The default setting is set to disable by default. To enable or disable this feature, enter the command below:

To enable or disable this feature, enter the command below:

```
8008:5(config)#router bgp
8008:5(router-bgp)#quick-start enable
8008:5(router-bgp)#no quick-start enable
```

18. BGP+

BGP+ enables the exchange of IPv6 routes using IPv4 peering. BGP+ is an extension of BGPv4 for IPv6 where 6-in-4 tunnels are required for BGP+ to work over IPv4 transport. BGP+ is the extension of BGP protocol to standards RFC 4760 (multi-protocol extensions to BGP) and RFC 2545 (MP-BGP for IPv6). These extensions allow BGPv4 peering to be enabled with IPv6 address family capabilities. The BGP+ implementation includes support for BGPv6 policies, including redistributing BGPv6 into OSPFv3/RIPng/ISISv6 and advertising OSPFv3, RIPng, ISISv6, Static and Local routes into BGPv6 (through BGP+). It also supports the aggregation of global unicast IPv6 addresses.

The current support is not an implementation of BGPv6. Native BGPv6 peering uses the IPv6 Transport layer (TCPv6) for establishing the BGPv6 peering, route exchanges, and data traffic. Native BGPv6 peering is not supported.

When you use BGP+ you must configure an IPv6 tunnel and static routes at BGP+ peers. When BGP+ peers advertise route information, they use Update messages to advertise route information. These RTM routes contain next-hop addresses from the BGP peer that the route was learned from. The static routes correlate the next-hop addresses represented by the IPv4–mapped IPv6 address to a specific outgoing interface.

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

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

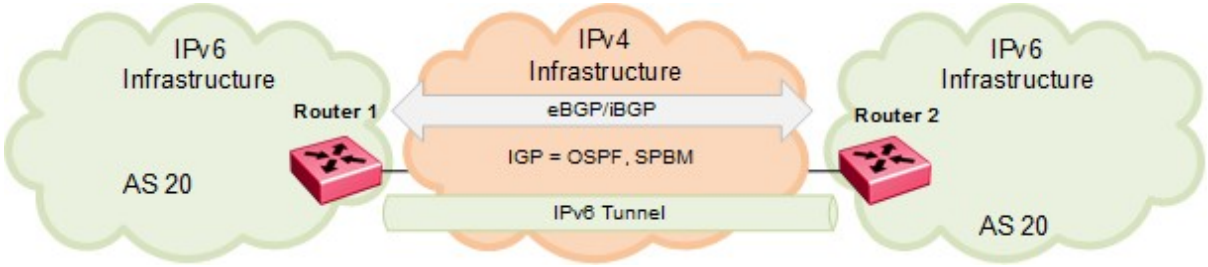

18.1 Configuration Example: iBGP+

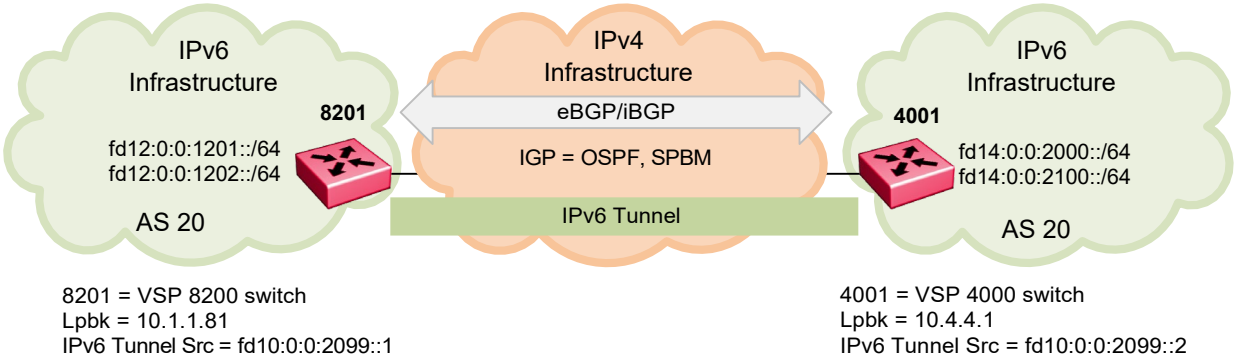

For this example, we will create an iBGP connection between switches 8201 and 4001 and use BGP+ to redistribute the local IPv6 interfaces shown in the above diagram.

18.1.1 BGP+ Configuration

Enable BGP

8201:

```
8201:1(config)#interface loopback 1
8201:1(config-if)#ip address 1 10.1.1.81/32
8201:1(config-if)#exit
8201:1(config)#router bgp
8201:1(router-bgp)#no auto-summary
8201:1(router-bgp)#no synchronization
8201:1(router-bgp)#router-id 10.1.1.81
8201:1(router-bgp)#neighbor 10.4.4.1
8201:1(router-bgp)#neighbor 10.4.4.1 remote-as 20
8201:1(router-bgp)#neighbor 10.4.4.1 next-hop-self
8201:1(router-bgp)#neighbor 10.4.4.1 update-source 10.1.1.81
8201:1(router-bgp)#neighbor 10.4.4.1 address-family ipv6
8201:1(router-bgp)#neighbor 10.4.4.1 enable
8201:1(router-bgp)#exit
8201:1(config)#router bgp 20 enable
```

4001:

```
4001:1(config)#interface loopback 1
4001:1(config-if)#ip address 1 10.4.4.1/32
4001:1(config-if)#exit
4001:1(config)#router bgp
```

```
4001:1(router-bgp)#no auto-summary
4001:1(router-bgp)#no synchronization
4001:1(router-bgp)#neighbor 10.1.1.81
4001:1(router-bgp)#neighbor 10.1.1.81 remote-as 20
4001:1(router-bgp)#neighbor 10.1.1.81 next—hop-self
4001:1(router-bgp)#neighbor 10.1.1.81 update-source 10.4.4.1
4001:1(router-bgp)#neighbor 10.1.1.81 address-family ipv6
4001:1(router-bgp)#neighbor 10.1.1.81 enable
4001:1(router-bgp)#exit
4001:1(config)#router bgp 20 enable
```

Enable IPv6 Forwarding

8201 & 4001: Same configuration on both switches

```
8201:1(config)#ipv6 forwarding
```

Enable IPv6 Tunnel

8201:

```
8201:1(config)#ipv6 tunnel 100 source 10.1.1.81 address fd10:0:0:2099::1/128
destination 10.4.4.1
```

4001:

```
4001:1(config)#ipv6 tunnel 100 source 10.4.4.1 address fd10:0:0:2099::2/64 destination
10.1.1.81
```

Enable IPv6 Static Route Configuration

8201:

```
8201:1(config)#ipv6 route ::ffff:10.4.4.1/128 cost 1 tunnel 100
```

4001:

```
4001:1(config)#ipv6 route ::ffff:10.1.1.81/128 cost 1 tunnel 100
```

The configuration used above is one way to express a static route for the IPv6 configured tunnel using the format *::ffff<ipv4 address>* or *0:0:0:0:ffff:<ipv4 address>.*

Enable redistribution of IPv6 direct interfaces

8201:

```
8201:1(config)#router bgp
8201:1(router-bgp)#redistribute ipv6-direct
8201:1(router-bgp)#redistribute ipv6-direct enable
8201:1(router-bgp)#exit
```

4001:

```
4001:1(config)#router bgp
```

```
4001:1(router-bgp)#redistribute ipv6-direct
4001:1(router-bgp)#redistribute ipv6-direct enable
4001:1(router-bgp)#exit
```

18.1.2 Verification

To verify BGP neighbor:

```
8201:1#show ip bgp neighbors 10.4.4.1
================================================================================
                        BGP Neighbor Info - GlobalRouter
================================================================================

BGP neighbor is 10.4.4.1 remote AS 20, Internal Peer, MP-BGP-capable,
BGP state [Established]
remote router ID 10.4.4.1


                                  vrf instance - 0
                                   admin-state - BGP ON
                         connect-retry-interval - 120
                                  ebgp-multihop - disable
                                     hold-time - 180
                                  keepalive-time - 60
                               hold-time-configured - 180
                          keepalive-time-configured - 60
                                     max-prefix - 12000
                                max-prefix-ipv6 - 8000
                                   nexthop-self - enable
                             originate-def-route - disable
                          originate-v6-def-route - disable

                               MD5-authentication - disable
                                  neighbor-debug - none
                               remove-private-as - disable
                     route-advertisement-interval - 5
                           route-reflector-client - disable
                                   send-community - disable
                        soft-reconfiguration-in - enable
                          updt-source-interface - 10.1.1.81
                                         weight - 100
                                Route Policy In -
                               Route Policy Out -
                             IPv6Route Policy In -
                            IPv6Route Policy Out -
                             address-family ipv6 - enable
                            address-family vpnv4 - disable
                                  route-refresh - disable
                    negotiated-session-capabilites - IPv6

Total bgp neighbors: 1
```

To verify IPv6 tunnel:

```
8201:1(config)#show ipv6 tunnel

================================================================================
                            Tunnel If Information
================================================================================
ID          LOCAL ADDRESS    REMOTE ADDRESS   TYPE   TTL

--------------------------------------------------------------------------------
100         10.1.1.81        10.4.4.1         manual 255

1 out of 1 Total number of entries displayed.
```

To verify IPv6 redistributed routes where LCL = local

```
8201:1#show bgp ipv6 redistributed-routes
================================================================================
                     BGPv6 Redistribute List - GlobalRouter
================================================================================

SRC-VRF           SRC   MET ENABLE RPOLICY
--------------------------------------------------------------------------------
GlobalRouter       LCL 0    TRUE
```

To verify IPv6 route table:

```
8201:1#show bgp ipv6 route
================================================================================
                          BGPv6 Routes - GlobalRouter
================================================================================
The total number of routes is 3

NETWORK/MASK                 PEER REM ADDR      NEXTHOP ADDRESS                ORG LOC PREF
--------------------------------------------------------------------------------------------------------------------------------
fd10:0:0:2099:0:0:0:0/64     10.4.4.1           0:0:0:0:0:ffff:10.4.4.1        INC 100
        AS_PATH: path-is-empty
fd14:0:0:2000:0:0:0:0/64     10.4.4.1           0:0:0:0:0:ffff:10.4.4.1        INC 100
        AS_PATH: path-is-empty
fd14:0:0:2100:0:0:0:0/64     10.4.4.1           0:0:0:0:0:ffff:10.4.4.1        INC 100
        AS_PATH: path-is-empty
```

```
8201:1#% show ipv6 route
================================================================================================
                                    IPv6 Routing Table Information
================================================================================================
Destination Address/PrefixLen          NEXT HOP                  VID/BID/TID  PROTO    COST  AGE
------------------------------------------------------------------------------------------------
0:0:0:0:0:ffff:10.4.4.1/128            0:0:0:0:0:0:0:0           T-100        STATIC   1     0
fd10:0:0:2099:0:0:0:0/64               0:0:0:0:0:0:0:0           T-100        LOCAL    1     0
fd10:0:0:2099:0:0:0:1/128              0:0:0:0:0:0:0:0           T-100        LOCAL    1     0
fd12:0:0:1201:0:0:0:0/64               0:0:0:0:0:0:0:0           V-400        LOCAL    1     0
fd12:0:0:1201:0:0:0:1/128              0:0:0:0:0:0:0:0           V-400        LOCAL    1     0
fd12:0:0:1202:0:0:0:0/64               0:0:0:0:0:0:0:0           V-401        LOCAL    1     0
fd12:0:0:1202:0:0:0:1/128              0:0:0:0:0:0:0:0           V-401        LOCAL    1     0
fd14:0:0:2000:0:0:0:0/64               0:0:0:0:0:0:0:0           T-100        BGP      0     0
fd14:0:0:2100:0:0:0:0/64               0:0:0:0:0:0:0:0           T-100        BGP      0     0
```

Test connectivity:

```
8201:1#ping fd14:0:0:2000::1
fd14:0:0:2000:0:0:0:1 is Alive

8201:1#ping fd14:0:0:2100::1
fd14:0:0:2100:0:0:0:1 is Alive
```

19. Appendix A

19.1 Translating Cisco to Extreme Equivalents

This appendix shows you how to translate Cisco commands and functions into their Extreme Ethernet Routing Switch equivalents.

Configuration Command Equivalents

Table 3 lists the Extreme CLI and Enterprise Device Manager equivalents for Cisco router configuration commands. In this table, **Bold text** indicates variables that the user supplies. The items in the list following the table describe the functions that the correspondingly numbered row configures.

Table 3 Translating Cisco to Extreme Equivalents

Item	Cisco Configuration	Extreme CLI Command	Enterprise Device Manager Logical Steps
1	router bgp **333** neighbor **1.1.1.2** remote-as **444**	router bgp **333 enable** router bgp neighbor **1.1.1.2** neighbor 1.1.1.2 remote-as **444** neighbor 1.1.1.2 **enable** ** If changing local-as, disable BGP first – no router bgp enable	***IP->BGP->Generals*** LocalAS: **333** AdminStatus: Enable ***IP->BGP->Peers->Insert*** RemoteAddr: **1.1.1.2** RemoteAs: **444** **Enable** Insert
2	network **10.1.1.0** mask **255.255.255.0**	router bgp network **10.1.1.0/24**	***IP->BGP->Network->Insert*** NetworkAfAddr: **10.1.1.0** NetworkAfPrefixLen: **24**

Item	Cisco Configuration	Extreme CLI Command	Enterprise Device Manager Logical Steps
3	neighbor **1.1.1.1** distribute-list **5 out** …access list **5** deny **128.1.0.0 0.0.255.255** …access list **5** permit **0.0.0.0 255.255.255.255**	ip prefix-list "**128.1.0.0**" **128.1.0.0/16** route-map "**distribute**" 1 no permit enable match network "**128.1.0.0**" route-map "**distribute**" 2 enable exit router bgp neighbor **1.1.1.1** neighbor 1.1.1.1 remote-as **444** neighbor 1.1.1.1 out-route-map "**distribute**" neighbor 1.1.1.1 enable exit	***IP->Policy->Prefix List->Insert*** Id: **1** Prefix: **128.1.0.0** PrefixMaskLen: **24** Name: **128.1.0.0** ***IP->Policy->Route Policy->Insert*** Id: **1** SequenceNumber: **1** Name: **distribute** **Enable** Mode: **deny** MatchNetwork: **128.1.0.0** **Insert** ***IP->Policy->Route Policy->Insert*** Id: **1** SequenceNumber: **2** Name: **distribute** **Enable** Mode: **permit** MatchNetwork: **128.1.0.0** **Insert** ***IP->BGP->Peers->Insert*** RemoteAddr: **1.1.1.1** RemoteAs: **444** **Enable** RoutePolicyOut: **distribute** **Insert**

Item	Cisco Configuration	Extreme CLI Command	Enterprise Device Manager Logical Steps
4	neighbor **1.1.1.1** route-map **IncomingMap in** …route-map **IncomingMap** permit **10** match as-path **5** set local-preference **125** …ip as-path access-list permit **333_444**	ip as-list **1** memberid **1** permit as-path "**333 444**" route-map "**IncomingMap**" **1** **enable** match as-path **1** set local-preference **125** route-map "IncomingMap" **2** **enable** exit router bgp neighbor "**1.1.1.1**" neighbor 1.1.1.1 1remote-as **444** neighbor 1.1.1.1 in-route-map "**IncomingMap**" neighbor 1.1.1.1 **enable** exit	***IP->Policy->AS Path List->Insert*** Id: **1** MemberId: **1** Mode: **Permit** AsRegularExpression: **333 444** ***IP->Policy->Route Policy->Insert*** Id: **1** SequenceNumber: **1** Name: **IncomingMap** **Enable** Mode: **Permit** SetLocalPref: **125** **Insert** ***IP->Policy->Route Policy->Insert*** Id: **1** SequenceNumber: **2** Name: **IncomingMap** **Enable** Mode: **permit** **Insert** ***IP->BGP->Peers->Insert*** RemoteAddr: **1.1.1.1** RemoteAs: **444** **Enable** RoutePolicyIn: **IncomingMap** **Insert**

Item	Cisco Configuration	Extreme CLI Command	Enterprise Device Manager Logical Steps
5	neighbor **1.1.1.1** route-map **setASPath out** …route-map **setASPath** permit **10** set as-path prepend **123** **123**	ip prefix-list "**200.1.40.0**" **200.1.40.0/24** ip as-list **1** memberid **20** permit as-path "**123 123**" route-map "**setASPath**" **1** **enable** match network "**200.1.40.0**" set as-path **1** exit router bgp neighbor "**1.1.1.1**" neighbor 1.1.1.1 remote-as **444** neighbor 1.1.1.1 out-route-map "**setASPath**" neighbor 1.1.1.1 **enable** exit	***IP->Policy->Prefix List->Insert*** Id: **1** Prefix: **200.1.40.0** PrefixMaskLen: **24** Name: **200.1.40.0** ***IP->Policy->Route Policy->Insert*** Id: **1** SequenceNumber: **1** Name: **setASPath** **Enable** Mode: **permit** MatchNetwork: **200.1.40.0** **Insert** ***IP->BGP->Peers->Insert*** RemoteAddr: **1.1.1.1** RemoteAs: **444** **Enable** RoutePolicyOut: **setASPath** **Insert**

Item	Cisco Configuration	Extreme CLI Command	Enterprise Device Manager Logical Steps
6	neighbor **1.1.1.1** route-map **AdvertiseMap out** …route-map **AdvertiseMap** permit **10** match ip address **1** set metric **100** route-map **AdvertiseMap** permit **20** set metric **50** …access-list **1** permit **192.10.20.0 0.0.0.255**	ip prefix-list "**192.10.20.0**" **192.10.20.0/24** route-map "**AdvertiseMap**" **1** **enable** match network "**192.10.20.0**" set metric 1**00** route-map "AdvertiseMap" **2** **enable** set metric **50** router bgp neighbor "1.1.1.1" neighbor 1.1.1.1 remote-as **444** neighbor 1.1.1.1 out-route-map "**AdvertiseMap**" neighbor 1.1.1.1 **enable** exit	***IP->Policy->Prefix List->Insert*** Id: **1** Prefix: **192.10.20.0** PrefixMaskLen: **24** Name: **192.168.20.0** ***IP->Policy->Route Policy->Insert*** Id: **1** SequenceNumber: **1** Name: **AdvertiseMap** **Enable** Mode: **permit** MatchNetwork: **192.10.20.0** SetMetric: **100** **Insert** ***IP->Policy->Route Policy->Insert*** Id: **1** SequenceNumber: **2** Name: **AdvertiseMap** **Enable** Mode: **permit** SetMetric: **50** ***IP->BGP->Peers->Insert*** RemoteAddr: **1.1.1.1** RemoteAs: **444** **Enable** RoutePolicyOut: **AdvertiseMap** **Insert**

Item	Cisco Configuration	Extreme CLI Command	Enterprise Device Manager Logical Steps
7	neighbor **MyPeers** peer-group neighbor **MyPeers** remote-as **333** neighbor **MyPeers** route-map **AdvertiseMap** **out** neighbor **MyPeers** route-map **IncomingMap in** neighbor 1.1.1.1 peer-group **MyPeers** neighbor 2.2.2.2 peer-group **MyPeers**	router bgp neighbor **"2.2.2.2"** neighbor **"1.1.1.1"** neighbor "**MyPeers**" neighbor **1.1.1.1** peer-group "**MyPeers**" neighbor **2.2.2.2** peer-group "**MyPeers**" neighbor "**MyPeers**" remote-as **333** neighbor "**MyPeers**" in-route-map "**IncomingMap**" neighbor "**MyPeers**" out-route-map "**AdvertiseMap**" neighbor "**MyPeers**" **enable** exit	***IP->BGP->Peer Groups->Insert*** GroupName: **MyPeers** Enable: **enable** RemoteAs: **333** RoutePolicyIn: I**ncomingMap** RoutePolicyOut: **AdvertiseMap** **Insert** ***IP->BGP->Peers->Insert*** RemoteAddr: **1.1.1.1** GroupName: **MyPeers** **Enable** **Insert** ***IP->BGP->Peers->Insert*** RemoteAddr: **2.2.2.2** GroupName: **MyPeers** **Enable** **Insert**
8	aggregate-address **195.89.8.0 255.255.248.0**	router bgp aggregate-address **195.89.8.0/21**	***IP->BGP->Aggregates->Insert*** Address: **195.89.8.0** PrefixLen: **21** **Insert**
9	aggregate-address **172.1.1.0 255.255.255.0** **summary-only**	router bgp aggregate-address 172.1.1.0/24 summary-only	***IP->BGP->Aggregates->Insert*** Address: **172.1.1.0** PrefixLen: **24** **SummaryOnly** **Insert**

Item	Cisco Configuration	Extreme CLI Command	Enterprise Device Manager Logical Steps
10	router ospf **101** redistribute bgp **2000**	router ospf redistribute **bgp** redistribute **bgp enable** exit ip ospf apply redistribute bgp ** Prior to enabling ospf redistribution, make sure the switch is configured for ospf ASBR. router ospf as-boundary-router enable exit router ospf enable	***IP->OSPF->Redistribute->Insert*** RouteSource: **bgp** Enable: **enable** **Insert** ***IP->OSPF->General*** AdminStat: **enabled** ASBdrRtrStatus: **checked** **Apply**
11	router bgp **2000** redistribute ospf **101** redistribute static	router bgp redistribute **ospf** redistribute **ospf enable** redistribute **static** redistribute **static enable** exit ip bgp apply redistribute ospf ip bgp apply redistribute static	***IP->BGP->Redistribute>Insert*** RouteSource: **ospf** Enable: **enable** **Insert** ***IP->BGP->Redistribute>Insert*** RouteSource: **static** Enable: **enable** **Insert**
12	timers bgp **60 180**	router bgp neighbor **1.1.1.1** timers **60 180**	***IP->BGP->Peers->1.1.1.1*** HoldTimeConfigured: **180** KeepAliveConfigured: **60** **Apply**

Item	Cisco Configuration	Extreme CLI Command	Enterprise Device Manager Logical Steps
13	interface loopback0 ip address **1.1.1.1** **255.255.255.255**	interface loopback **1** ip address **1 1.1.1.1/255.255.255.255** exit ** To enable loopback for ospf and in turn BGP router-id, enter interface loopback **1** **ip ospf** exit router ospf router-id **1.1.1.1** exit	***IP->IP->Circuitless IP->Insert*** Interface: **1** Ip Address: **1.1.1.1** Net Mask: **255.255.255.255** OSPF: enable (click on tab)
14	ip subnet zero	The ERS or VSP switch has no parameter for zero subnet, already enabled.	The ERS or VSP switch has no parameter for zero subnet, already enabled.

Item	Cisco Configuration	Extreme CLI Command	Enterprise Device Manager Logical Steps
15	router bgp **4001** bgp confederation identifier **5** bgp confederation peers **4002 4003 4004** neighbor **1.2.3.4** remote-as **4002** neighbor **3.4.5.6** remote-as **510**	router bgp **4001** router bgp bgp confederation identifier **5** peers "**4002 4003 4004** " neighbor "**1.2.3.4"** neighbor **1.2.3.4** remote-as **4002** neighbor **1.2.3.4 enable** neighbor "**3.4.5.6**" neighbor **3.4.5.6** remote-as **510** neighbor **3.4.5.6** enable	***IP->BGP->Generals*** AdminStatus: **enable** LocalAS: **4001** ConfederationIdentifier: **5** ConfederationPeers: **4002,4003,4004** **Apply** ***IP->BGP->Peers->Insert*** RemoteAddress: **1.2.3.4** RemoteAs: **4002** **Enable** ***IP->BGP->Peers->Insert*** I RemoteAddress: **3.4.5.6** RemoteAs: **510** **Enable**
16	router bgp **1000** neighbor **132.245.10.2** password **bla4u00=2nkq**	router bgp neighbor "**132.245.10.2**" neighbor **132.245.10.2** MD5-authentication **enable** neighbor password 132.245.10.2 **bla4u00=2nkq**	CLI only

Item	Cisco Configuration	Extreme CLI Command	Enterprise Device Manager Logical Steps
17	neighbor **1.1.1.1** remote-as **100** neighbor **1.1.1.1** remote-as **100** route-reflector-client	router bgp **100** enable router bgp route-reflector **enable **** bgp client-to-client reflection ** exit ** Enabled by default router bgp neighbor "**1.1.1.1**" neighbor **1.1.1.1** remote-as **100** neighbor **1.1.1.1 route-reflector-client** neighbor **1.1.1.1 enable**	***IP->BGP->Generals*** AdminStatus: **enable** LocalAS: **100** RouteReflectionEnable: **enable **** ReflectorClientToClientReflection: **enable **** ** Enabled by default **Apply** ***IP- >BGP->Peers->Insert*** RemoteAddr: **1.1.1.1** RemoteAs: **100** **RouteReflectionClient** **Insert**
18	neighbor **5.5.5.5** remote-as **100** route-reflector-client bgp cluster-id **10**	router bgp **100** router bgp bgp cluster-id **0.0.0.10** neighbor **5.5.5.5** remote-as **100** neighbor **5.5.5.5 route-reflector-client** neighbor **5.5.5.5 enable**	***IP->BGP->Generals*** AdminStatus: enable LocalAS: **100** RouteReflectorClusterId: **0.0.0.10** **Apply** ***IP->BGP->Peers->Insert*** RemoteAddr: **5.5.5.5** RemoteAs: **100** **RouteReflectoinClient** **Insert**
19	router bgp **100** neighbor **10.1.1.42** default-originate	router bgp **100** router bgp neighbor **10.1.1.42** default-originate enable	***IP->BGP->Peers*** RemoteAddr: **10.1.1.42** DefaultOriginate: **true**

19.2 Interpreting the Cisco to Extreme BGP Translation Table

The numbers in the following list correspond to the item numbers in Table 2. Each numbered item in this list describes the function of the commands in the corresponding row of that table.

1. Enable the Border Gateway Protocol (BGP) routing process and identify the local router autonomous system (AS), 333. Activate a BGP session with peer router, IP address, 1.1.1.2 that belongs to AS 444. If the local and remote AS numbers are the same, the BGP session is internal otherwise it is an external BGP session.
2. Advertise network 10.1.1.0 mask 255.255.255.0 and originate it from my AS. Note that network 10.1.1.0 must be present in the IP routing table for BGP network command to advertise the route.
3. Deny incoming advertisement of network 128.1.0.0, mask 255.255.0.0 from peer IP address, 1.1.1.1, as specified by Cisco access list 5 or Extreme policy name distribute.
4. Accept incoming advertisements, from peer 1.1.1.1, match on AS-Path that contain either AS “333 444” or 345 and set Local Preference to 125, as specified by Cisco route-map and Extreme policy name IncomingMap.
5. Announce advertisements to peer 1.1.1.1 and append AS-Path <123 123> to all outgoing updates, as specified by Cisco route-map and Extreme policy name setASPath.
6. Announce advertisement of network 192.10.20.0 mask 255.255.255.0 to peer IP address 1.1.1.1, setting multi-exit discriminator (MED) to 100 as specified by Cisco route-map and Extreme policy name AdvertiseMap. In addition, advertise any other networks with MED set to 50.
7. Create a peer group named MyPeers with the following elements: peer router AS is 333, advertise networks as specified by route-map AdvertiseMap and accept incoming networks as specified by FilterMap. Assign peer routers 1.1.1.1 and 2.2.2.2 to peer group MyPeers
8. Advertise the aggregate address 195.89.8.0 mask 255.255.248.0 (195.89.8.0/21) as well as the more specific addresses i.e. 195.89.8.0 - 195.89.15.0.

9. Advertise the aggregate address 195.89.8.0 mask 255.255.248.0 (195.89.8.0/21) only.

10. To redistribute BGP routes into OSPF.
11. To redistribute OSPF and static routes into BGP.
12. Keepalive timer is used between BGP peers as a periodic check of the TCP connection between them. Holddown timer is the amount of elapsed time before the BGP peering session is declared dead. RFC 1771 suggests values of 30 and 90 seconds respectively. Holddown timer is suggested to be three times the amount of the keepalive timer.
13. Cisco’s loopback interface and Extreme circuitless IP interface is useful in BGP environments to use as peer interfaces. It is highly recommended using loopback interfaces for BGP as it eliminates the dependency that would otherwise occur when you use the IP address of a physical interface.
14. Enable the use of subnet zero for interface addresses and routing updates.
15. Enable Confederations for IBGP full mesh reduction. In this example, the outside world sees this as a single AS, number 5, but within the AS it is divided into autonomous systems 4001, 4002, 4003 and 4004. This router's confederation ID is 4001. It has a peer 1.2.3.4 within its routing confederation domain and another peer 3.4.5.6 outside.
16. Enables MD5 authentication on the TCP connection between the two BGP peers (132.245.10.1 and 132.245.10.2). In this example, the MD5 key is **bla4u00=2nkq.**

17. Enable Route Reflectors for IBGP full mesh reduction. The ERS 8000 is also configured to allow router reflector client to client route distribution.
18. Enable Route Reflectors with two route reflectors for redundancy. A cluster id must be configured when there are two or more router reflectors in a cluster.
19. Send default route to BGP peer 10.1.1.42.

19.3 Comparing Cisco and Extreme BGP Operational Commands

Table 4 compares the corresponding Cisco and Extreme operational commands. The itemized list following this table describes the function of the commands in the corresponding row of this table.

Table 4: Cisco and Extreme BGP Operational Commands

Item	Cisco	Extreme
1	no synchronization	no synchronization
2	route reflector	route-reflector enable
3	bgp damping	flap-damping enable
4	Confederation	bgp confederation
BGP Monitoring Commands		
5	show ip route bgp	show ip bgp route
6	show ip bgp neighbors	show ip bgp sum
7	show ip bgp neighbors 1.1.1.2	show ip bgp neighbor 1.1.1.2
8	show ip bgp neighbors 1.1.1.2	show ip bgp neighbor 1.1.1.2 stats
9	show ip bgp neighbors 1.1.1.2	show ip bgp neighbor 1.1.1.2 routes
10	clear ip bgp neighbor-ip-address	ip bgp restart-bgp neighbor <ip address> ip bgp restart-bgp neighbor <ip address> soft-reconfiguration in ip bgp restart-bgp neighbor <ip address> soft-reconfiguration out
nei11	show ip route	show ip route
12	trace 1.1.1.1	traceroute 1.1.1.1
13	debug ip bgp	Through the local console port shows various debug commands that can configured for displaying bgp state, events, and more. Use the following command for bgp global debug:

		router bgp **global-debug mask <value>** List of mask values include: none, all, error, packet, event, trace, warning, state, init, filter, update Use the following command for bgp neighbor debug: **neighbor-debug-all mask <value>** List of mask values include: none, all, error, packet, event, trace, warning, state, init, filter, update **debug-screen on** This will output the debug information to the console. **NOTE:** excessive messages to the console will affect CPU performance.

19.4 Interpreting the Cisco and Extreme BGP Operational Table

The following list describes the function of the Cisco and Extreme operational commands in the corresponding row of Table 4.

1. Do not synchronize between BGP and IGP; this enables a router to advertise a BGP network to an external peer without having that network exist in the IP routing table.
2. Route reflection is a method to alleviate the need for “full mesh” IBGP by allowing an internal BGP speaker to reflect (or re-advertise) routes learned through an IBGP connection to another IBGP peer.
3. Minimize the instability caused by route flapping.
4. Confederations are used to reduce the number of peers in an AS by breaking the network into multiple (smaller) ASs.
5. Show BGP routing table.
6. Show status of BGP peers.
7. Show the router’s BGP timers. Within Cisco’s show ip bgp neighbor command the keepalive, hold-down and external advertisement timers are displayed.
8. Display the router’s statistics.
9. Cisco’s show ip bgp neighbor command displays the router’s incoming and outgoing route filters. The Extreme show ip bgp neighbor route command display incoming routes from peer 1.1.1.2.
10. Reset a neighbor’s BGP connection.
11. Display the IP routing table.
12. Discover the routes the router’s packets take when traveling to destination 1.1.1.1.
13. Display BGP updates/changes/events as they occur.

19.5 Interpreting the Cisco and Extreme BGP Operational Table

Table 5 compares the Cisco and Extreme route preference.

Table 5: Cisco and Extreme Route Preference Comparison

Route Type	Cisco –Pref. value	Extreme – Pref.value
Directly connected	0	0
Static	1	5
EBGP	20	45
OSPF Intra	110	20
OSPF Inter		25
BGP	20	30
RIP	120	100
OSPF External 1		120
OSPF External 2		125
IBGP	200	175

```
show ip route preference
================================================================================
                        IP Route Preference - GlobalRouter
================================================================================
PROTOCOL         DEFAULT    CONFIG
--------------------------------------------------------------------------------
LOCAL            0          0
STATIC           5          5
OSPF_INTRA       20         20
OSPF_INTER       25         25
EBGP             45         45
RIP              100        100
OSPF_E1          120        120
OSPF_E2          125        125
IBGP             175        175
STATICv6         5          5
OSPFv3_INTRA     20         20
OSPFv3_INTER     25         25
OSPFv3_E1        120        120
OSPFv3_E2        125        125
SPBM_L1          7          130
```

20. Appendix B

20.1 Translating Juniper to Extreme Equivalents

This appendix shows you how to translate Juniper commands and functions into their Extreme Switch 8000 equivalents.

Configuration Command Equivalents

Table 6 lists the Extreme CLI and Device Manager equivalents for Juniper router configuration commands. In this table, **Bold text** indicates variables that the user supplies. The items in the list following the table describe the functions that the correspondingly numbered row configures.

Table 6: Translating Juniper to ERS 8000 Equivalents

Item	Juniper Configuration	Extreme CLI Command	Enterprise Device Manager Logical Steps
1	set routing-options autonomous-system **333** edit protocols bgp group **ebgp** <enter> set type external set peer-as **444** set local-as **333** set neighbor **1.1.1.2**	router bgp **333 enable** router bgp neighbor **1.1.1.2** neighbor 1.1.1.2 remote-as **444** neighbor 1.1.1.2 **enable** ** If changing local-as, disable BGP first – no router bgp enable	***IP->BGP->Generals*** LocalAS: **333** AdminStatus: Enable ***IP->BGP->Peers->Insert*** RemoteAddr: **1.1.1.2** RemoteAs: **444** **Enable** Insert

2	protocols {
	bgp {
	export **direct**;
	
	policy-options {
	policy-statement **direct** {
	term **dir_export** {
	from protocol direct;
	then accept;
		router bgp
		
		network **10.1.1.0/24**
			IP->BGP->Network->Insert
			
			NetworkAfAddr: **1.1.1.0**
			
			NetworkAfPrefixLen: **24**

Item	Juniper Configuration	Extreme CLI Command	Enterprise Device Manager Logical Steps

3	protocols {
	bgp {
	group **ebgp** {
	type external;
	export **drop**;
	peer-as **300**;
	neighbor **1.1.1.1**;
	
	policy-options {
	policy-statement **drop** {
	term **list** {
	from {
	protocol bgp;
	route-filter **128.1.0.0/16** exact reject;
	}
	then accept;
		ip prefix-list "**128.1.0.0**" **128.1.0.0/16**
		
		route-map "**distribute**" 1
		no permit
		enable
		match network "**128.1.0.0**"
		route-map "**distribute**" 2
		enable
		exit
		
		router bgp
		neighbor **1.1.1.1**
		neighbor 1.1.1.1 remote-as **444**
		neighbor 1.1.1.1 out-route-map "**distribute**"
		neighbor 1.1.1.1 enable
		exit
			IP->Policy->Prefix List->Insert
			Id: **1**
			Prefix: **128.1.0.0**
			PrefixMaskLen: **24**
			Name: **128.1.0.0**
			IP->Policy->Route Policy->Insert
			Id: **1**
			SequenceNumber: **1**
			Name: **distribute**
			Enable
			Mode: **deny**
			MatchNetwork: **128.1.0.0**
			Insert
			IP->Policy->Route Policy->Insert
			Id: **1**
			SequenceNumber: **2**
			Name: **distribute**
			Enable
			Mode: **permit**
			MatchNetwork: **128.1.0.0**
			Insert
			IP->BGP->Peers->Insert
			RemoteAddr: **1.1.1.1**
			RemoteAs: **444**
			Enable
			RoutePolicyOut: **distribute**
			Insert

Item	Juniper Configuration	Extreme CLI Command	Enterprise Device Manager Logical Steps

4	policy-options {
	
	policy-statement **IncomingMap** {
	
	term **as** {
	
	from {
	
	neighbor **1.1.1.1**;
	
	as-path **aslist**;
	
	}
	
	then {
	
	local-preference **125**;
	
	}
	
	}
	
	}
	
	as-path **aslist 333-444**;
	
	protocols {
	
	bgp {
	
	import **IncomingMap**;
		ip as-list **1** memberid **1** permit as-path "**333 444**"
		
		route-map "**IncomingMap**" **1**
		
		enable
		
		match as-path **1**
		
		set local-preference **125**
		
		route-map "IncomingMap" **2**
		
		enable
		
		exit
		
		router bgp
		
		neighbor "**1.1.1.1**"
		
		neighbor 1.1.1.1 1remote-as **444**
		
		neighbor 1.1.1.1 in-route-map "**IncomingMap**"
		
		neighbor 1.1.1.1 **enable**
		
		exit
			IP->Policy->AS Path List->Insert
			
			Id: **1**
			
			MemberId: **1**
			
			Mode: **Permit**
			
			AsRegularExpression: **333 444**
			
			IP->Policy->Route Policy->Insert
			
			Id: **1**
			
			SequenceNumber: **1**
			
			Name: **IncomingMap**
			
			Enable
			
			Mode: **Permit**
			
			SetLocalPref: **125**
			
			Insert
			
			IP->Policy->Route Policy->Insert
			
			Id: **1**
			
			SequenceNumber: **2**
			
			Name: **IncomingMap**
			
			Enable
			
			Mode: **permit**
			
			Insert
			
			IP->BGP->Peers->Insert
			
			RemoteAddr: **1.1.1.1**
			
			RemoteAs: **444**
			
			Enable
			
			RoutePolicyIn: **IncomingMap**
			
			Insert

Item	Juniper Configuration	Extreme CLI Command	Enterprise Device Manager Logical Steps

5	policy-options {
	
	policy-statement **setASPath** {
	
	term **ASList** {
	
	from {
	
	route-filter **200.1.40.0**/24 exact;
	
	}
	
	then as-path-prepend "**123 123**";
	
	protocols {
	
	bgp {
	
	group **ebgp** {
	
	type external;
	
	export **setASPath**;
	
	peer-as **300**;
	
	neighbor **1.1.1.1**;
	
	}
		ip prefix-list "**200.1.40.0**" **200.1.40.0/24**
		
		ip as-list **1** memberid **20** permit as-path "**123 123**"
		
		route-map "**setASPath**" **1**
		
		enable
		
		match network "**200.1.40.0**"
		
		set as-path **1**
		
		exit
		
		router bgp
		
		neighbor "**1.1.1.1**"
		
		neighbor 1.1.1.1 remote-as **444**
		
		neighbor 1.1.1.1 out-route-map "**setASPath**"
		
		neighbor 1.1.1.1 **enable**
		
		exit
			IP->Policy->Prefix List->Insert
			
			Id: **1**
			
			Prefix: **200.1.40.0**
			
			PrefixMaskLen: **24**
			
			Name: **200.1.40.0**
			
			IP->Policy->Route Policy->Insert
			
			Id: **1**
			
			SequenceNumber: **1**
			
			Name: **setASPath**
			
			Enable
			
			Mode: **permit**
			
			MatchNetwork: **200.1.40.0**
			
			Insert
			
			IP->BGP->Peers->Insert
			
			RemoteAddr: **1.1.1.1**
			
			RemoteAs: **444**
			
			Enable
			
			RoutePolicyOut: **setASPath**
			
			Insert

Item	Juniper Configuration	Extreme CLI Command	Enterprise Device Manager Logical Steps

6	policy-options {
	policy-statement **AdvertiseMap** {
	term **seq1** {
	from {
	route-filter **192.10.20.0/24** exact;
	}
	then {
	metric **100**;
	accept;
	}
	}
	term **seq2** {
	from {
	route-filter **0.0.0.0/0** orlonger;
	}
	then {
	metric **50**;
	accept;
	
	protocols {
	bgp {
	group **ebgp** {
	type external;
	export **AdvertiseMap**;
	peer-as **300**;
	neighbor **1.1.1.1**;
		ip prefix-list "**192.10.20.0**" **192.10.20.0/24**
		
		route-map "**AdvertiseMap**" **1**
		enable
		match network "**192.10.20.0**"
		set metric 1**00**
		route-map "AdvertiseMap" **2**
		enable
		set metric **50**
		
		router bgp
		neighbor "1.1.1.1"
		neighbor 1.1.1.1 remote-as **444**
		neighbor 1.1.1.1 out-route-map "**AdvertiseMap**"
		neighbor 1.1.1.1 **enable**
		exit
			IP->Policy->Prefix List->Insert
			Id: **1**
			Prefix: **192.10.20.0**
			PrefixMaskLen: **24**
			Name: **192.168.20.0**
			IP->Policy->Route Policy->Insert
			Id: **1**
			SequenceNumber: **1**
			Name: **AdvertiseMap**
			Enable
			Mode: **permit**
			MatchNetwork: **192.10.20.0**
			SetMetric: **100**
			Insert
			IP->Policy->Route Policy->Insert
			Id: **1**
			SequenceNumber: **2**
			Name: **AdvertiseMap**
			Enable
			Mode: **permit**
			SetMetric: **50**
			IP->BGP->Peers->Insert
			RemoteAddr: **1.1.1.1**
			RemoteAs: **444**
			Enable
			RoutePolicyOut: **AdvertiseMap**
			Insert

Item	Juniper Configuration	Extreme CLI Command	Enterprise Device Manager Logical Steps

7	protocols {
	bgp {
	group **ibgp** {
	type internal;
	export **NHS**
	peer-as **333**;
	neighbor **1.1.1.1**;
	neighbor **2.2.2.2**;
	
	policy-options {
	policy-statement **NHS** {
	term **next-hop** {
	from {
	protocol bgp;
	}
	then {
	next-hop self;
		router bgp
		neighbor **"2.2.2.2"**
		neighbor **"1.1.1.1"**
		neighbor "**MyPeers**"
		neighbor **1.1.1.1** peer-group "**MyPeers**"
		neighbor **2.2.2.2** peer-group "**MyPeers**"
		neighbor "**MyPeers**" remote-as **333**
		neighbor "**MyPeers**" in-route-map "**IncomingMap**"
		neighbor "**MyPeers**" out-route-map "**AdvertiseMap**"
		neighbor "**MyPeers**" **enable**
		exit
			IP->BGP->Peer Groups->Insert
			GroupName: **MyPeers**
			Enable: **enable**
			RemoteAs: **333**
			RoutePolicyIn: I**ncomingMap**
			RoutePolicyOut: **AdvertiseMap**
			Insert
			IP->BGP->Peers->Insert
			RemoteAddr: **1.1.1.1**
			GroupName: **MyPeers**
			Enable
			Insert
			IP->BGP->Peers->Insert
			RemoteAddr: **2.2.2.2**
			GroupName: **MyPeers**
			Enable
			Insert

8	policy-options {
	policy-statement **agg-add** {
	term **agg** {
	from {
	route-filter **195.89.8.0/20** orlonger;
	}
	then accept;
		router bgp
		aggregate-address **195.89.8.0/21**
			IP->BGP->Aggregates->Insert
			Address: **195.89.8.0**
			PrefixLen: **21**
			Insert

Item	Juniper Configuration	Extreme CLI Command	Enterprise Device Manager Logical Steps

9	policy-options {
	policy-statement **agg-add** {
	term **agg** {
	from {
	route-filter **172.1.1.0/24** exact;
	}
	then accept;
		router bgp
		aggregate-address **172.1.1.0/24** summary-only
			IP->BGP->Aggregates->Insert
			Address: **172.1.1.0**
			PrefixLen: **24**
			SummaryOnly
			Insert

10	protocols {
	ospf {
	export **bgp_routes**;
	area **0.0.0.0** {
	interface **20.1.1.1** {
	metric **200**;
	
	policy-options {
	policy-statement **bgp_routes** {
	from protocol bgp;
	then accept;
		router ospf
		redistribute **bgp**
		redistribute **bgp enable**
		exit
		ip ospf apply redistribute bgp
		
		** Prior to enabling ospf redistribution, make sure the switch is configured for ospf ASBR.
		router ospf
		as-boundary-router enable
		exit
		router ospf enable
			IP->OSPF->Redistribute->Insert
			RouteSource: **bgp**
			Enable: **enable**
			Insert
			
			IP->OSPF->General
			AdminStat: **enabled**
			ASBdrRtrStatus: **checked**
			Apply

Item	Juniper Configuration	Extreme CLI Command	Enterprise Device Manager Logical Steps

11	protocols {
	bgp {
	export **ospf_into_bgp**;
	
	policy-options {
	policy-statement **ospf_into_bgp** {
	term **ospf-only** {
	from protocol ospf;
	then accept;
		router bgp
		redistribute **ospf**
		redistribute **ospf enable**
		exit
		ip bgp apply redistribute ospf
			IP->BGP->Redistribute>Insert
			RouteSource: **ospf**
			Enable: **enable**
			Insert

12	protocols {
	bgp {
	group **ebgp** {
	type external;
	hold-time **180**;
	peer-as **300**;
	local-as **100**;
	neighbor **1.1.1.1**;
	
	**JUNOS software defaults a Keepalive time of always one-third the HoldTime. In this example, the HoldTime is set for 180 so the KeepAlive will default to 60.
		router bgp
		neighbor **1.1.1.1** timers **60 180**
			IP->BGP->Peers->1.1.1.1
			HoldTimeConfigured: **180**
			KeepAliveConfigured: **60**
			Apply

Item	Juniper Configuration	Extreme CLI Command	Enterprise Device Manager Logical Steps

13	interfaces {
	lo0 {
	unit **0** {
	family inet {
	address **1.1.1.1/32**;
		interface loopback **1**
		ip address **1 1.1.1.1/255.255.255.255**
		exit
		** To enable loopback for ospf and in turn BGP router-id, enter
		interface loopback **1**
		ip ospf
		exit
		router ospf
		router-id **1.1.1.1**
		exit
			IP->IP->Circuitless IP->Insert
			Interface: **1**
			Ip Address: **1.1.1.1**
			Net Mask: **255.255.255.255**
			OSPF: enable (click on tab)

14	Synchronization Disabled. In JUNOS software, synchronization is disabled by default. There is no option to enable or disable synchronization.	router bgp **no** synchronization	***IP->BGP->Generals*** Synchronization: **disable**

Item	Juniper Configuration	Extreme CLI Command	Enterprise Device Manager Logical Steps

15	routing-options {
	autonomous-system **4001**;
	confederation **5** members [**4002 4003 4004**]
	
	protocols {
	bgp {
	group **1234** {
	type external;
	peer-as **4002**;
	neighbor **1.2.3.4**;
	}
	group **3456** {
	type external;
	peer-as **510**;
	neighbor **3.4.5.6**;
		router bgp **4001**
		router bgp
		bgp confederation identifier **5** peers "**4002 4003 4004** "
		neighbor "**1.2.3.4"**
		neighbor **1.2.3.4** remote-as **4002**
		neighbor **1.2.3.4 enable**
		neighbor "**3.4.5.6**"
		neighbor **3.4.5.6** remote-as **510**
		neighbor **3.4.5.6** enable
			IP->BGP->Generals
			AdminStatus: **enable**
			LocalAS: **4001**
			ConfederationIdentifier: **5**
			ConfederationPeers: **4002,4003,4004**
			Apply
			
			IP->BGP->Peers->Insert
			RemoteAddress: **1.2.3.4**
			RemoteAs: **4002**
			Enable
			IP->BGP->Peers->Insert
			I RemoteAddress: **3.4.5.6**
			RemoteAs: **510**
			Enable

16	set protocols bgp group **ebgp** authentication-key **bla4u00=2nkq**	router bgp neighbor "**132.245.10.2**" neighbor **132.245.10.2** MD5-authentication **enable** neighbor password 132.245.10.2 **bla4u00=2nkq**	CLI only

Item	Juniper Configuration	Extreme CLI Command	Enterprise Device Manager Logical Steps

17	interfaces {
	fe-1/1/0 {
	unit **0** {
	family inet {
	address **10.10.10.1/30**;
	}
	}
	}
	fe-1/1/1 {
	unit **0** {
	family inet {
	address **10.10.10.13/30**;
	}
	}
	}
	lo0 {
	unit **0** {
	family inet {
	address **1.1.1.1/32**;
	
	routing-options {
	static {
	route **1.1.1.2/32** next-hop [**10.10.10.2 10.10.10.14**];
	
	protocols {
	bgp {
	group **ebgp** {
	type external;
	multihop ttl **2**;
	local-address **1.1.1.1**; ©
	peer-as **300**;
		interface GigabitEthernet 1/15
		brouter vlan **2078** subnet **10.10.10.2/30**
		exit
		interface GigabitEthernet 1/17
		brouter vlan **2079** subnet **10.10.10.14/30**
		exit
		
		interface loopback **1**
		ip address 1 10.1.1.2/255.255.255.255
		ip ospf 1
		exit
		
		router bgp neighbor
		"**1.1.1.1**"
		neighbor 1.1.1.1 remote-as **100**
		neighbor 1.1.1.1 ebgp-multihop
		neighbor 1.1.1.1 update-source **1.1.1.2**
		neighbor 1.1.1.1 enable
		exit
		
		p route **1.1.1.1 255.255.255.255 10.10.10.1** weight **1**
		ip route **1.1.1.1 255.255.255.255 10.10.10.13** weight **10**
			Right-Click port 1/15->Edit IP->Insert
			Ip Address: **10.10.10.2/30**
			Net Mask: **255.255.255.252**
			Vlan Id: **2078**
			Right-Click port 1/17->Edit IP->Insert
			Ip Address: **10.10.10.14/30**
			Net Mask: **255.255.255.252**
			Vlan Id: **2079**
			
			IP->IP->Circuitless IP->Insert
			Interface: **1**
			Ip Address: **10.1.1.2**
			Net Mask: **255.255.255.0**
			Insert
			
			IP->BGP->Peers->Insert
			RemoteAddress: **1.1.1.1**
			RemoteAs: **100**
			Enable
			EbgpMultiHop
			Insert
			UpdateSourceInterface: **1.1.1.2**
			
			IP->IP->Static Routes->Insert
			Dest: **1.1.1.1**
			Mask: **255.255.255.255**
			NextHop: **10.10.10.1**
			Metic: **1**
			Insert
			IP->IP->Static Routes->Insert
			Dest: **1.1.1.1**
			Mask: **255.255.255.255**
			NextHop: **10.10.10.13**

Item	Juniper Configuration	Extreme CLI Command	Enterprise Device Manager Logical Steps

18	routing-options {
	autonomous-system **100**;
	
	protocols {
	bgp {
	group **rr-cluster1** {
	peer-as **100**
	local-address **5.5.5.4**;
	cluster **0.0.0.10**
	neighbor **5.5.5.5**;
	group **rr-cluster2** {
	peer-as **100**
	local-address **1.1.1.2**;
	cluster **0.0.0.10**
	neighbor **1.1.1.1**;
		router bgp **100**
		router bgp
		bgp cluster-id **0.0.0.10**
		neighbor **5.5.5.5** remote-as **100**
		neighbor **5.5.5.5 route-reflector-client**
		neighbor **5.5.5.5 enable**
			P->BGP->Generals
			AdminStatus: enable
			LocalAS: **100**
			RouteReflectorClusterId: **0.0.0.10**
			Apply
			IP->BGP->Peers->Insert
			RemoteAddr: **5.5.5.5**
			RemoteAs: **100**
			RouteReflectoinClient
			Insert

20.2 Interpreting the Juniper to Extreme BGP Translation Table

The numbers in the following list correspond to the item numbers in Table 5. Each numbered item in this list describes the function of the commands in the corresponding row of that table.

1. Enable the Border Gateway Protocol (BGP) routing process and identify the local router autonomous system (AS), 333. Activate a BGP session with peer router, IP address, 1.1.1.2 that belongs to AS 444. If the local and remote AS numbers are the same, the BGP session is internal, otherwise it is an external BGP session.
2. Advertise network 1.1.1.0 and 1.1.1.4 mask 255.255.255.252 that are direct interfaces on the ERS 8000 and originate it from my AS. Note that by default Juniper will advertise all learned routes and the BGP Network command is not used. A policy statement can be added, as shown in this configuration example, in order for the Juniper router to advertise its direct interfaces.
3. Deny incoming advertisement of network 128.1.0.0, mask 255.255.0.0 from peer IP address, 1.1.1.1, as specified by Juniper policy-statement drop or Extreme policy name distribute.
4. Accept incoming advertisements, from peer 1.1.1.1, match on AS-Path that contain either AS “333 444” or 345 and set Local Preference to 125, as specified by Juniper policy-statement IncomingMap and Extreme policy name IncomingMap.
5. Announce advertisements to peer 1.1.1.1 and append AS-Path <123 123> to all outgoing updates, as specified by Juniper policy-statement setASPath route-map and Extreme policy name setASPath.
6. Announce advertisement of network 192.10.20.0 mask 255.255.255.0 to peer IP address 1.1.1.1, setting multi-exit discriminator (MED) to 100 as specified by Juniper policy-statement AdvertiseMap and Extreme policy name AdvertiseMap. In addition, advertise any other networks with MED set to 50.
7. Accept incoming advertisements from peer 1.1.1.1, of AS-Path that contain either exactly AS 1000 or 5000 as specified by Juniper policy-statement AS_Filter and Extreme policy name AS_Filter.
8. Announce advertisements to peer 1.1.1.1 if the update includes an AS-Path that matches <350 400> and deny updates of AS-Path that contain <350 400 500> as specified by Juniper policy-statement Deny_AS and Extreme policy name Deny_AS.
9. Create a peer group named NHS with the following elements: nexthop-self enabled. Assign peer routers 1.1.1.1 and 2.2.2.2 to peer group MyPeers. Similar functionality is performed on Juniper by using the policy-statement NHS.
10. Advertise the aggregate address 195.89.8.0 mask 255.255.248.0 (195.89.8.0/21) as well as the more specific addresses i.e. 195.89.8.0 - 195.89.15.0.
11. Advertise the aggregate address 195.89.8.0 mask 255.255.248.0 (195.89.8.0/21) only.
12. To redistribute BGP routes into OSPF.
13. To redistribute OSPF into BGP.
14. Keepalive timer is used between BGP peers as a periodic check of the TCP connection between them. Holddown timer is the amount of elapsed time before the BGP peering session is declared dead. RFC 1771 suggests values of 30 and 90 seconds respectively. Holddown timer is suggested to be three times the amount of the keepalive timer.
15. Juniper’s loopback interface and Extreme circuitless IP interface is useful in BGP environments to use as peer interfaces. It is highly recommended using loopback interfaces for BGP as it

eliminates the dependency that would otherwise occur when you use the IP address of a physical interface.

16. Disable synchronization on the ERS 8000. By default, synchronization is disabled on Juniper and there is no option to enable or disable this functionality.
17. Enable Confederations for IBGP full mesh reduction. In this example, the outside world sees this as a single AS, number 5, but within the AS it is divided into autonomous systems 4001, 4002, 4003 and 4004. This router's confederation ID is 4001. It has a peer 1.2.3.4 within its routing confederation domain and another peer 3.4.5.6 outside.
18. Enables MD5 authentication on the TCP connection between the two BGP peers (132.245.10.1 and 132.245.10.2). In this example, the MD5 key is **bla4u00=2nkq.**
19. Enable EBGP multihop load balancing. The EBGP peering is between the loopback interface on Juniper and the circuitless ip on Extreme. On each router, static routes to the remote peer's loopback address must be configured for each data link connection.
20. Enable Route Reflectors for IBGP full mesh reduction. A cluster id is always used by Juniper and must be configured on ERS 8000 when there are two or more router reflectors in a cluster.

20.3 Comparing Juniper and Extreme BGP Operational Commands

Table 7 compares the corresponding Juniper and Extreme operational commands. The itemized list following this table describes the function of the commands in the corresponding row of this table.

Table 7: Juniper and Extreme BGP Operational Commands

Item	Juniper	Extreme
1	no synchronization	no synchronization
2	Route reflector	route-reflector enable
3	Bgp damping	flap-damping enable
4	Confederation	bgp confederation
BGP Monitoring Commands		
5	show route protocol bgp	show ip bgp route
6	show bgp summary	show ip bgp sum
7	show bgp neighbor 1.1.1.2	show ip bgp neighbor 1.1.1.2
8	show bgp neighbor 1.1.1.2	show ip bgp neighbor 1.1.1.2 stats
9	show route advertising-protocol bgp 1.1.1.2	show ip bgp neighbor 1.1.1.2 route
10	clear bgp neighbor <*ip address*>	ip bgp restart-bgp neighbor <ip address> ip bgp restart-bgp neighbor <ip address> soft-reconfiguration in ip bgp restart-bgp neighbor <ip address> soft-reconfiguration out
11	show route	show ip route
12	traceroute 1.1.1.1	traceroute 1.1.1.1
13	a) show log messages	Through the local console port shows various debug commands that can configured for displaying bgp state, events, and more. Use the following command for bgp global debug:

	b) configure the following: [edit protocols bgp] set traceoptions file bgp-log size 1m files 10 then use the following command: show log bgp-log	**router bgp** **global-debug mask <value>** List of mask values include: none, all, error, packet, event, trace, warning, state, init, filter, update Use the following command for bgp neighbor debug: **neighbor-debug-all mask <value>** List of mask values include: none, all, error, packet, event, trace, warning, state, init, filter, update **debug-screen on** This will output the debug information to the console. **NOTE:** excessive messages to the console will affect CPU performance.

20.4 Interpreting the Juniper and Extreme BGP Operational Table

The following list describes the function of the Juniper and Extreme operational commands in the corresponding row of Table 7.

1. Do not synchronize between BGP and IGP; this enables a router to advertise a BGP network to an external peer without having that network exist in the IP routing table.
2. Route reflection is a method to alleviate the need for “full mesh” IBGP by allowing an internal BGP speaker to reflect (or re-advertise) routes learned through an IBGP connection to another IBGP peer.
3. Minimize the instability caused by route flapping.
4. Confederations are used to reduce the number of peers in an AS by breaking the network into multiple (smaller) ASs.
5. Show BGP routing table.
6. Show status of BGP peers.
7. Show the router’s BGP neighbor information.
8. Display the router’s statistics.
9. Juniper’s show route advertising-protocol bgp command displays the router’s incoming and outgoing routes. The Extreme show ip bgp neighbor route command display incoming routes from peer 1.1.1.2.
10. Reset a neighbor’s BGP connection.
11. Display the IP routing table.
12. Discover the routes the router’s packets take when traveling to destination 1.1.1.1.
13. Display BGP updates/changes/events as they occur.

20.5 Interpreting the Juniper and Extreme BGP Operational Table

Table 8 compares the Cisco and Extreme route preference.

Table 8: Route Preference Comparison

Route Type	Juniper –Pref. value	P8000 – Pref.value
Directly connected	0	0
Static	5	5
EBGP	170	12
OSPF Intra	10	15
OSPF Inter	150	17
BGP	170	30
RIP	100	100
OSPF External 1	150	120
OSPF External 2	150	125
IBGP	170	200

21. Appendix C – BGP Events

BGP State Transitions and Actions:

This Appendix discusses the transitions between states in response to BGP events. The following is the list of these states and events when the negotiated Hold Time value is non-zero.

```
BGP States:

          1 - Idle
          2 - Connect
          3 - Active
          4 - OpenSent
          5 - OpenConfirm
          6 - Established

BGP Events:

          1 - BGP Start
          2 - BGP Stop
          3 - BGP Transport connection open
          4 - BGP Transport connection closed
          5 - BGP Transport connection open failed
          6 - BGP Transport fatal error
          7 - ConnectRetry timer expired
          8 - Hold Timer expired
          9 - KeepAlive timer expired
         10 - Receive OPEN message
         11 - Receive KEEPALIVE message
         12 - Receive UPDATE messages
         13 - Receive NOTIFICATION message
```

Event	Actions	Message Sent	Next State
Idle (1)			
1	Initialize resources Start ConnectRetry timer Initiate a transport connection	none	2
others	none	none	1
Connect(2)			
1	none	none	2
3	Complete initialization Clear ConnectRetry timer	OPEN	4
5	Restart ConnectRetry timer	none	3
7	Restart ConnectRetry timer Initiate a transport connection	none	2
others	Release resources	none	1
Active (3)			
1	none	none	3
3	Complete initialization Clear ConnectRetry timer	OPEN	4
5	Close connection Restart ConnectRetry timer		3
7	Restart ConnectRetry timer Initiate a transport connection	none	2
others	Release resources	none	1
OpenSent(4)			
1	none	none	4
4	Close transport connection Restart ConnectRetry timer	none	3
6	Release resources	none	1
10	Process OPEN is OK	KEEPALIVE	5
	Process OPEN failed	NOTIFICATION	1
others	Close transport connection Release resources	NOTIFICATION	1
OpenConfirm (5)			
1	none	none	5
4	Release resources	none	1
6	Release resources	none	1
9	Restart KeepAlive timer	KEEPALIVE	5
11	Complete initialization Restart Hold Timer	none	6
13	Close transport connection Release resources		1
others	Close transport connection Release resources	NOTIFICATION	1
Established (6)			
1	none	none	6
4	Release resources	none	1
6	Release resources	none	1
9	Restart KeepAlive timer	KEEPALIVE	6
11	Restart Hold Timer	KEEPALIVE	6

```
    12              Process UPDATE is OK          UPDATE              6
                    Process UPDATE failed         NOTIFICATION        1
    13              Close transport connection                        1
                    Release resources
    others          Close transport connection    NOTIFICATION        1
                    Release resources

----------------------------------------
```

The following is a condensed version of the above state transition table.

Events	Idle (1)	Connect (2)	Active (3)	OpenSent (4)	OpenConfirm (5)	Estab (6)
1	2	2	3	4	5	6
2	1	1	1	1	1	1
3	1	4	4	1	1	1
4	1	1	1	3	1	1
5	1	3	3	1	1	1
6	1	1	1	1	1	1
7	1	2	2	1	1	1
8	1	1	1	1	1	1
9	1	1	1	1	5	6
10	1	1	1	1 or 5	1	1
11	1	1	1	1	6	6
12	1	1	1	1	1	1 or 6
13	1	1	1	1	1	1

BGP States:	**BGP Events:**
1 - Idle 2 - Connect 3 - Active 4 - OpenSent 5 - OpenConfirm 6 - Established	1 - BGP Start 2 - BGP Stop 3 - BGP Transport connection open 4 - BGP Transport connection closed 5 - BGP Transport connection open failed 6 - BGP Transport fatal error 7 - ConnectRetry timer expired 8 - Hold Timer expired 9 - KeepAlive timer expired 10 - Receive OPEN message 11 - Receive KEEPALIVE message 12 - Receive UPDATE messages 13 - Receive NOTIFICATION message

22. Appendix D – EDM BGP Command Options

Table 9 displays the various BGP configuration options available while Table 10 displays the various BGP peer configuration options available.

Table 9: EDM BGP Configuration Options

Field	Description
AdminStatus	Enables or disables BGP on the router. The default is disable. You cannot enable AdminStatus until you change the LocalAS value to a nonzero value.
4ByteAs	Enables or disables the switch from using 4 byte numbers for autonomous systems.
LocalAs	Sets the local autonomous system (AS) number. You cannot change this field when AdminStatus is set to enable. This field sets a 2-byte local AS number in the range from 1 to 65535. To set a 4-byte local AS number, click enable in the 4ByteAs field and enter a number in the NewLocalAs field. Attention: This parameter is not supported with BGP+.
AsDot	Enables or disables representing AS numbers in octects. The default is disable so the switch uses the plain notation format. If you enable this field and the 4ByteAs field, enter the AS number in the NewLocalAs field. Attention: This parameter is not supported with BGP+.
Aggregate	Enables or disables aggregation. The default is enable. You cannot change the value when BGP is enabled.
DefaultMetric	Sets the metric sent to BGP neighbors. The Default Metric determines the cost of a route a neighbor uses. Use this parameter in conjunction with the redistribute parameters so that BGP uses the same metric for all redistributed routes. The default is -1. The range is -1 to 2147483647. Note: A default metric value helps solve the problems associated with redistributing routes that have incompatible metrics. For example, whenever metrics do not convert, using a default metric provides a reasonable substitute and allows the redistribution to proceed.

Field	Description
DefaultLocalPreference	Specifies the default value of the local preference attribute. The default value is 100. Specify an integer value in the range 0 to 2147483647 Note: You cannot change the default value when AdminStatus is set to enable.
DefaultInformationOriginate	Specifies the default local preference. The default is 100. The range is 0 to 2147483647.
DefaultInformationOriginateIPv6	Enables or disables the redistribution of a default IPv6 network into BGP. The default is disable.
AlwaysCompareMed	Enables or disables the comparison of the multi-exit discriminator (MED) parameter for paths from neighbors in different ASs. A path with a lower MED is preferred over a path with a higher MED. The default is disable.
DeterministicMed	Enables or disables deterministic MED. Deterministic MED compares the MEDs when routes advertised by different peers in the same AS are chosen. The default is disable.
AutoPeerRestart	Enables or disables the process that automatically restarts a connection to a BGP neighbor. The default is enable. Attention: Because the switch does not support IPv6 BGP peers, this parameter does not support IPv6 peers.
AutoSummary	Enables or disables automatic summarization. When enabled, BGP summarizes networks based on class limits (for example, Class A, B, or C networks). The default is enable.
NoMedPathIsWorst	Enables or disables NoMedPathIsWorst. When set to enable (default), BGP treats an update that is missing a MED attribute as the worst path.
BestPathMedConfed	Enables or disables the comparison of MED attributes within a confederation. The default is disable.

Field	Description
DebugMask	Displays the specified debug information for BGP global configurations. The default value is none. • none disables all debug messages. • all enables all debug messages. • error enables the display of debug error messages. • packet enables the display of debug packet messages. • event enables the display of debug event messages. • trace enables the display of debug trace messages. • warning enables the display of debug warning messages. • state enables display of debug state transition messages. • init enables the display of debug initialization messages. • filter enables the display of debug messages related to filtering. • update enables display of debug messages related to updates transmission and reception. • error enables the display of debug error messages.
IgnoreillegalRouterId	Enables BGP to overlook an illegal router ID. For example, it enables the acceptance of a connection from a peer that sends an open message using a router ID of 0. The default is enable.
Synchronization	Enables or disables the router to accept routes from iBGP peers without waiting for an update from the IGP. The default is enable.
MaxEqualCostRoutes	Sets the maximum number of equal-cost-paths that are available to a BGP router by limiting the number of equal-cost-paths that can be stored in the routing table. The default value is 1; the range is 1 to 8.
IbgpReportImportRoute	Configures BGP to report imported routes to an interior BGP (IBGP) peer. This command also enables or disables reporting of non-BGP imported routes to other IBGP neighbors. The default is enable.
FlapDampEnable	Enables or disables route suppression for routes that go up and down (flap). The default is disable. This parameter is not supported with BGP+.
QuickStart	Enables or disables the Quick Start feature, which forces the BGP speaker to begin establishing peers immediately, instead of waiting for

Field	Description
	the peer's autorestart timer to expire. The default is disable.
TrapEnable	Enables or disables BGP traps.
ConfederationIdentifier	Specifies a BGP confederation identifier in the range of 0 to 65535. The default is 0.
ConfederationPeers	Lists adjoining ASs that are part of the confederation in the format (5500,65535,0,10,...,...). The default is none.
RouteReflectionEnable	Enables or disables the reflection of routes from IBGP neighbors. The default value is disable.
RouteReflectorClusterId	Sets a reflector cluster ID IP address. This option applies only if RouteReflectionEnable is set to enable, and if multiple route reflectors are in a cluster. The default value is 0.0.0.0.
ReflectorClientToClientReflection	Enables or disables route reflection between two route reflector clients. This option is applicable only if the RouteReflectionEnable value is set to enable. The default value is disable.
RouteRefresh	Enables or disables IP VPN Route Refresh for BGP. The default is disable. If enabled, a route refresh request received by a BGP speaker causes the speaker to resend all route updates it contains in its database that are eligible for the peer that issues the request. This parameter is not supported with BGP+ and not supported on the VSP 9000.

Table 10: EDM BGP Peer Configuration Options

Field	Description
RemoteAddr	Specifies the IP address of this peer or subscriber group. The default is none. (Peer creation is not possible without a remote address.)
GroupName	Specifies the peer group name to which this peer belongs (optional). The default is none.
RemoteAs	Configures a remote AS number for the peer or peer-group in the range 0 to 65535. The default is 0.
Enable	Enables or disables the peer. Double-click in the field to access the pull-down menu.
RoutePolicyIn	Specifies the route policy that applies to all IPv4 networks learned from this peer. The default value is none. To remove a route policy, click under the RoutePolicyIn column, hit Ctrl and highlight the policy to remove, and click OK.
RoutePolicyOut	Specifies the route policy that applies to all outgoing IPv4 updates to this peer. The default value is none. To remove a route policy, click under the RoutePolicyOut column, hit Ctrl and highlight the policy to remove, and click OK.
DefaultOriginate	When enabled, allows the local router to send the default IPv4 route to the neighbor for use as a default route. The default value is disable.
DefaultOriginateIpv6	When enabled, allows the local router to send the default IPv6 route to the neighbor for use as a default route. The default value is disable.
ConnectRetryInterval	Specifies the time interval (in seconds) for the ConnectRetry timer. The suggested value for this timer is 120 seconds (the default).
HoldTimeConfigured	Specifies the time interval (in seconds) for the Hold Time for this BGP speaker with this peer. This value is placed in an OPENmessage sent to this peer by this BGP speaker, and is compared with the HoldTime in an OPEN message received from the peer when the switch determines the Hold Time with the peer. The HoldTime must be at least three seconds. if it is zero, the Hold Time is not to be established with the peer. The suggested value for this timer is 90 seconds. The default is 180 seconds.
KeepAliveConfigured	Specifies the time interval (in seconds) for the KeepAlivetimer configured for this BGP speaker with this peer. KeepAliveConfigured determines the KEEPALIVE messages frequency relative to HoldTimeConfigured; the actual time interval for

Field	Description
	the KEEPALIVE messages is indicated byKeepAlive. The recommended maximum value for this timer is one-third of HoldTimeConfigured. If KeepAliveConfigured is zero, no periodic KEEPALIVE messages are sent to the peer after the BGP connection is established. The suggested value for this timer is 30 seconds. The default is 60.
DebugMask	Displays the specified debug information for the BGP peer. The default value is none. • none disables all debug messages. • all enables all debug messages. • error enables the display of debug error messages. • packet enables the display of debug packet messages. • event enables the display of debug event messages. • trace enables the display of debug trace messages. • warning enables the display of debug warning messages. • state enables display of debug state transition messages. • init enables the display of debug initialization messages. • filter enables the display of debug messages related to filtering. • update enables display of debug messages related to updates transmission and reception.
AdvertisementInterval	Specifies the time interval (in seconds) that elapses between each transmission of an advertisement from a BGP neighbor. The default value is 30 seconds and the range is 5 to 120 seconds.
Weight	Specifies this peer's or peer groups' weight, or the priority of updates that can be received from this BGP peer. The default value is 100 and the range is 0 to 65535.
MaxPrefix	Sets a limit on the number of routes that are accepted from a neighbor. The default value is 12000 routes and the range is 0 to 2147483647. 0 means there is no limit to the number of routes that are accepted.
RouteReflectorClient	Specifies that this peer is a route reflector client. The default is false.
SoftReconfigurationIn	When enabled, the router relearns routes from the specified neighbor or group of neighbors without resetting the connection when the policy changes in the inbound direction. The default value is enable. Enabling SoftReconfigurationIn

Field	Description
	causes all BGP routes to be stored in local memory (even non-best routes).
SendCommunity	Enables or disables sending the update message's community attribute to the specified peer. The default value is disable.
Ipv6Cap	Specifies (when enabled) that IPv6capability can be enabled or disabled on the BGP neighbor peer. The default is disable.
RouteRefresh	Configures a route refresh for the BGP peer.
AsOverride	Specifies that the AS Override parameter can be enabled or disabled for the BGP peer. The default is disable.
AllowAsIn	Specifies the number of AS-in allowed for the BGP peer. The range is 0 to 10. The default is 0.
Ipv6RoutePolicyIn	Specifies the route policy that applies to all IPv6 networks learned from this peer. The default value is none. To remove a route policy, click under the Ipv6RoutePolicyIn column, hit Ctrl and highlight the policy to remove, and click OK.
Ipv6RoutePolicyOut	Specifies the route policy that applies to all outgoing IPv6 updates to this peer. The default value is none. To remove a route policy, click under the Ipv6RoutePolicyOut column, hit Ctrl and highlight the policy to remove, and click OK.