

BGP Multi-Homing for VPLS Networks

In This Chapter

This section describes BGP Multi-Homing (BGP-MH) for VPLS network configurations.

Topics in this section include:

- [Applicability on page 172](#)
- [Summary on page 173](#)
- [Overview on page 175](#)
- [Configuration on page 177](#)
- [Conclusion on page 207](#)

Applicability

This section is applicable to all of the 7750 SR series (SR-7, SR-12, SR-c4 and SR-c12), 7710 SR, as well as 7450 ESS (ESS-7 and ESS-12) series in mixed mode. It was tested on release 9.0R1.

Summary

Beginning with SR OS 8.0R1, the SR/ESS portfolio supports the use of Border Gateway Protocol Multi-Homing for VPLS (hereafter called BGP-MH). BGP-MH is described in draft-ietf-l2vpn-vpls-multihoming, *BGP based Multi-homing in Virtual Private LAN Service*, and provides a network-based resiliency mechanism (no interaction from the PEs Provider Edge router — to MTU/CEs Multi-tenant Unit/Customer Equipment) that can be applied on access Service Access Points (SAPs) or network (pseudowires) topologies. The BGP-MH procedures will run between the PEs and will provide a loop-free topology from the network perspective (only one logical active path will be provided per VPLS among all the objects SAPs or pseudowires which are part of the same Multi-Homing site).

Each multi-homing site connected to two or more peers is represented by a site-id (2-bytes long) which is encoded in the BGP MH Network Layer Reachability Information (NLRI). The BGP peer holding the active path for a particular multi-homing site will be named as the Designated Forwarder (DF), whereas the rest of the BGP peers participating in the BGP MH process for that site will be named as non-DF and will block the traffic (in both directions) for all the objects belonging to that multi-homing site.

In SR OS 8.0, BGP MH uses the following rules to determine which PE is the DF for a particular multi-homing site:

1. A BGP MH NLRI with D flag = 0 (multi-homing object up) always takes precedence over a BGP MH NLRI with D flag = 1 (multi-homing object down). If there is a tie, then:
2. The BGP MH NLRI with the highest BGP LP (Local Preference) wins. If there is a tie, then:
3. The BGP MH NLRI issued from the PE with the lowest PE ID (system address) wins.

The main advantages of using BGP-MH as opposed to other resiliency mechanisms for VPLS are:

- Flexibility: BGP-MH uses a common mechanism for access and core resiliency. The designer has the flexibility of using BGP-MH to control the active/standby status of SAPs, spoke SDPs, Split Horizon Groups (SHGs) or even mesh SDP bindings.
- The standard protocol is based on BGP, a standard, scalable and well-known protocol.
- Specific benefits at the access:
 - It is network-based, independent of the customer CE and, as such, it does not need any customer interaction to determine the active path. Consequently the operator will spend less effort on provisioning and will minimize both operation costs and security risks (in particular, this removes the requirement for spanning tree interaction between the PE and CE).
 - Easy load balancing per service (no service fate-sharing) on physical links.

Summary

- Specific benefits in the core:
 - It is a network-based mechanism, independent of the MTU resiliency capabilities and it does not need MTU interaction, therefore operational advantages are achieved as a result of the use of BGP-MH: less provisioning is required and there will be minimal risks of loops. In addition, simpler MTUs can be used.
 - Easy load balancing per service (no service fate-sharing) on physical links.
 - Less control plane overhead: there is no need for an additional protocol running the pseudowire redundancy when BGP is already used in the core of the network. BGP-MH just adds a separate NLRI in the L2-VPN family (AFI=25, SAFI=65).

The objective of this section is to provide the required guidelines to configure and troubleshoot BGP-MH for VPLS based on SR OS 8.0R5.

Knowledge of the LDP/BGP VPLS (RFC 4762, *Virtual Private LAN Service (VPLS) Using Label Distribution Protocol (LDP) Signaling*, and RFC 4761, *Virtual Private LAN Service (VPLS) Using BGP for Auto-Discovery and Signaling*) architecture and functionality is assumed throughout this document. For further information refer to the relevant Alcatel-Lucent documentation.

Overview

The following network setup will be used throughout the rest of the section.

Note:

- IGP — ISIS, Level 2 on all routers; area 49.0001
- RSVP-TE for transport tunnels
- FRR in the core
- No protection at the access.




Figure 31: Network Topology

The setup consists of a three core 7x50 SR/ESS (PE-1, PE-2 and PE-3) and three Multi-Tenant Unit (MTU) nodes connected to the core. The MTU nodes can also be 7x50 service routers. All the nodes are running SR OS 8.0R5.

The VPLS service 500 is configured on all the six nodes with the following characteristics:

- The core VPLS instances are connected by a full mesh of BGP-signalled pseudowires (that is, pseudowires among PE-1, PE-2 and PE-3 will be signalled by BGP VPLS).
- As depicted in the [Figure 31](#), the MTUs are connected to the BGP VPLS core by TLDP pseudowires. While MTU-3 is connected to PE-3 by a single pseudowire, MTU-1 and

MTU-2 are dual-homed to PE-1 and PE-2. The following resiliency mechanisms are used on the dual-homed MTUs:

- MTU-1 is dual-connected to PE-1 and PE-2 by an active/standby pseudowire (A/S pseudowire hereafter).
- MTU-2 is dual-connected to PE-1 and PE-2 by two active pseudowires, being one of them blocked by BGP MH running between PE-1 and PE-2. The PE-1 and PE-2 pseudowires, set up from MTU-2, will be part of the BGP MH site MH-site-2.
- MTU-1 and MTU-2 are running BGP MH, being SHG site-1 and sap 1/1/4:500 on MTU-2 part of the same BGP MH site, MH-site-1.
- The CEs are connected to the network in the following way:
 - CE-1, CE-3 and CE-4 are single-connected to the network
 - CE-2 is dual connected to MTU-1 and MTU-2.
 - CE-1 and CE-2 are part of the split-horizon-group (SHG) site-1(saps 1/1/4:500 and 1/1/3:500 on MTU-1). Assume that CE-1 and CE-2 have a backdoor link between them so that when MTU-2 is elected as DF, CE1 does not get isolated. This configuration high-lights the use of a SHG within a site configuration.

For each BGP MH site, MH-site-1 and MH-site-2, the BGP MH process will elect a DF, blocking the site objects for the non-DF nodes. In other words, based on the specific configuration explained throughout the section:

- For MH-site-1, MTU-1 will be elected as the DF. The non-DF-MTU-2 will block the sap 1/1/4:500.
- For MH-site-2, PE-1 will be elected as the DF. The non-DF PE-1 will block the spoke-SDP to MTU-2.

Configuration

This section describes all the relevant configuration tasks for the setup shown in figure 1. Note that the appropriate associated IP/MPLS configuration is out of the scope of this section. In this particular example the following protocols will be configured beforehand:

- ISIS-TE as IGP with all the interfaces being level-2 (OSPF-TE could have been used instead).
- RSVP-TE as the MPLS protocol to signal the transport tunnels (LDP could have been used instead, but without Fast Re-Route).
- LSPs between core PEs will be Fast Re-Route protected (facility bypass tunnels) whereas LSP tunnels between MTUs and PEs will not be protected¹.

Once the IP/MPLS infrastructure is up and running, the specific service configuration including the support for BGP MH can begin.

Global BGP Configuration

BGP is used in this configuration guide for these purposes:

- a. Auto-discovery and signaling of the pseudowires in the core, as per RFC 4761.
- b. Exchange of multi-homing site NLRIs and redundancy handling from MTU-2 to the core.
- c. Exchange of multi-homing site NLRIs and redundancy handling at the access for CE-1/CE-2.

A BGP route reflector (RR), PE-3, is used for the reflection of BGP updates corresponding to the above uses **a** and **b**.

A direct peering is established between MTU-1 and MTU-2 for use **c**. Note that the same RR could have been used for the three cases, however, like in this example, the designer may choose to have a direct BGP peering between access devices. The reasons for this are:

- By having a direct BGP peering between MTU-1 and MTU-2, the BGP updates do not have to travel back and forth.

1. Note that the designer can choose whether to protect access link failures by means of MPLS FRR or A/S pseudowire or BGP MH. While FRR provides a faster convergence (around 50ms) and stability (it does not impact on the service layer, hence, link failures do not trigger MAC flush and flooding) some interim inefficiencies can be introduced compared to A/S pseudowire or BGP MH.

Global BGP Configuration

- On MTU-1 and MTU-2, BGP is exclusively used for multi-homing, therefore there will not be more BGP peers for either MTUs and a RR adds nothing in terms of control plane scalability.

As an example, the following CLI output shows the global BGP configuration required on MTU-1. Note that the 192.0.2.21 address will be replaced by the corresponding peer or the RR system address for PE-1 and PE-2.

```
A:MTU-1>config>router# info
-----
...
autonomous-system 65000
router-id 192.0.2.11
...
bgp
  rapid-withdrawal
  rapid-update l2-vpn
  family l2-vpn
  group "Multi-homing"
    neighbor 192.0.2.21
      type internal
      peer-as 65000
    exit
  exit
exit
```

In this example, PE-3 is the BGP RR, therefore its BGP configuration will contain a cluster with all its peers included (PE-1 and PE-2):

```
A:PE-3>config>router# info
-----
...
bgp
  rapid-withdrawal
  rapid-update l2-vpn
  family l2-vpn
  group "internal"
    cluster 1.1.1.1
    neighbor 192.0.2.1
      type internal
      peer-as 65000
    exit
    neighbor 192.0.2.2
      type internal
      peer-as 65000
    exit
  exit
exit
...
```


The relevant BGP commands for BGP-MH are in bold. Some considerations about those:

- It is required to specify **family l2-vpn** in the BGP configuration. That statement will allow the BGP peers to agree on the support for the family AFI=25 (Layer 2 VPN), SAFI=65 (VPLS). This family is used for BGP VPLS as well as for BGP MH and BGP AD.
- The **rapid-update l2-vpn** statement allows BGP MH to send BGP updates immediately after detecting link failures, without having to wait for the MRAI to send the updates in batches. This statement is required to guarantee a fast convergence for BGP MH.
- Optionally, rapid-withdrawal can also be added. Note that, in the context of BGP MH, this command is only useful if a particular multi-homing site is cleared². In that case, a BGP withdrawal is sent immediately without having to wait for the MRAI.

2. This means removing the BGP-MH site or even to remove the whole VPLS service.

Service Level Configuration

Once the IP/MPLS infrastructure is configured, including BGP, this section shows the configuration required at service level (VPLS 500). The focus is on the nodes involved on BGP MH, that is, MTU-1, MTU-2, PE-1 and PE-2. These nodes are highlighted in [Figure 32](#).

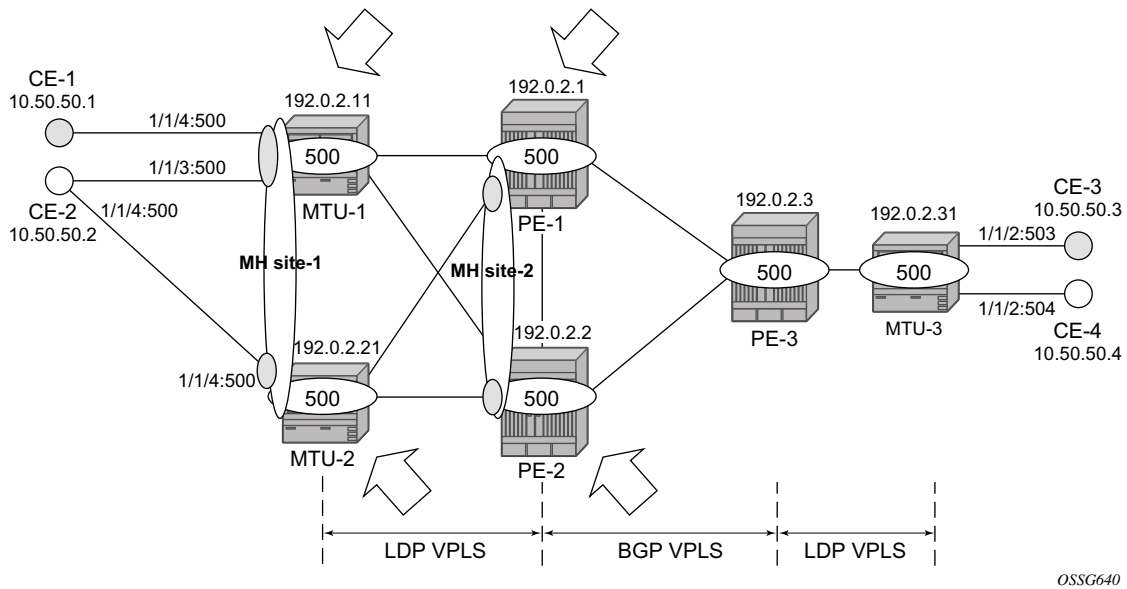


Figure 32: Nodes Involved in BGP MH

Core PE Service Configuration

The following CLI excerpt shows the service level configuration on PE-1.

```
A:PE-1>config>service# info
-----
...
sdp 111 mpls create
  far-end 192.0.2.11
  lsp "LSP-PE-1-MTU-1"
  path-mtu 8000
  keep-alive
  shutdown
  exit
no shutdown
exit
sdp 121 mpls create
  far-end 192.0.2.21
  lsp "LSP-PE-1-MTU-2"
  path-mtu 8000
  keep-alive
  shutdown
  exit
no shutdown
exit
sdp 1212 mpls create
  far-end 192.0.2.2
  lsp "LSP-PE-1-PE-2"
  signaling bgp # this sdp will transport BGP-signaled pseudowires
  path-mtu 8000
  keep-alive
  shutdown
  exit
no shutdown
exit
sdp 1313 mpls create
  far-end 192.0.2.3
  lsp LSP-PE-1-PE-3
  signaling bgp # this sdp will transport BGP-signaled pseudowires
  path-mtu 8000
  keep-alive
  shutdown
  exit
no shutdown
exit
...
pw-template 500 use-provisioned-sdp create
exit
...
vpls 500 customer 1 create
  split-horizon-group "CORE" create
  exit
  bgp # BGP context with common parameters for BGP VPLS and BGP MH
  route-distinguisher 65000:501
  vsi-export "vsi500_export"
  vsi-import "vsi500_import"
  pw-template-bind 500 split-horizon-group "CORE"
```

Core PE Service Configuration

```
exit
bgp-vpls # Specific context for BGP VPLS configuration
  max-ve-id 65535
  ve-name 501
    ve-id 501
  exit
  no shutdown
exit
stp
  shutdown
exit
site "MH-site-2" create # BGP MH specific configuration
  site-id 2
  spoke-sdp 121:500
  no shutdown
exit
spoke-sdp 111:500 create
exit
spoke-sdp 121:500 create
exit
no shutdown
exit
```

The following are general comments about the configuration of service 500:

- As it can be seen in the CLI output above for PE-1, there are four provisioned SDPs that the service VPLS 500 will use in this example. SDP 111 and SDP 121 are tunnels over which the TLDP FEC128 pseudowires for service 500 will be carried (according to RFC 4762), whereas SDP 1212 and SDP 1313 are the tunnels for the core BGP pseudowires (based on RFC 4761).
- The BGP context provides the general service BGP configuration that will be used by BGP VPLS and BGP MH:
 - Route distinguisher (notation chosen is based on <AS_number:500 + node_id>)
 - VSI export policies are used to add the export route-targets included in all the BGP updates sent to the BGP peers.
 - VSI import policies are used to control the NLRIs accepted in the RIB, normally based on the route-targets.
 - Both, VSI-export and VSI-import policies can be used to modify attributes like the Local-Preference (LP) that will be used to influence the BGP MH Designated Forwarder (DF) election (LP is the second rule in the BGP MH election process, as previously discussed). The use of these policies will be described later in the section.
 - The **pw-template-bind** command maps the previously defined pw-template 500 to the **split-horizon-group CORE**. In this way, all the BGP-signalled pseudowires will be part of this split-horizon-group. Although not shown in this example, the pw-template-bind command can also be used to instantiate pseudowires within different split-horizon-groups, based on different import route-targets³:

```
A:PE-1# configure service vpls 1 bgp pw-template-bind ?
- pw-template-bind <policy-id> [split-horizon-group <group-name>] [import-rt {ext-
community, ... (upto 5 max)}]
- no pw-template-bind <policy-id>
```

- The BGP-signalled pseudowires (from PE-1 to PE-2 and PE-3) are set up according to the configuration in the BGP context. Beside those pseudowires, the VPLS 500 also has two more pseudowires signalled by TLDP: spoke-sdp 111:500 (to MTU-1) and spoke-sdp 121:500 (to MTU-2).

3. Detailed BGP-VPLS configuration is out of the scope of this configuration guide. For more information please, refer to the Alcatel-Lucent SR OS documentation.

The general BGP MH configuration parameters for a particular multi-homing site are show in the following output:

```
*A:PE1> vpls <service-id> [customer <customer-id>][vpn <vpn-id>][m-vpls][b-vpls]
[no] site <site-name> [create] # 32 characters maximum
      [no] site-id <site-id-value> # Range: 1..65535

      [no] spoke-sdp sdp-id[:vc-id]
      [no] sap <sap-id>
      [no] split-horizon-group <group-name>
      [no] mesh-sdp-binding

      [no] failed-threshold <value|all> # Range:1-1000. Default: all

      [no] boot-timer <interval> # Range:0-600. Default: 10
      [no] site-activation-timer <interval> # Range:0-100. Default: 2

      [no] shutdown
```

Where:

- The **site-name** is defined by a string with less or equal to 32 characters.
- The **site-id** is an integer that identifies the multi-homing site and is encoded in the BGP MH NLRI. This ID must be the same one used on the peer node where the same multi-homing site is connected to. That is, MH-site-2 must use the same site-id in PE-1 and PE-2 (value = 2 in the PE-1 site configuration).
- Out of the four potential objects in a site, spoke SDP, SAP, SHG and mesh SDP binding only one can be used at the time on a particular site. To add more than just one sap/spoke-sdp to the same site, a split-horizon-group composed of the sap/spoke-sdp objects must be used in the site configuration. Otherwise, only one object (spoke SDP, SAP, SHG and mesh SDP binding) is allowed per site. A CLI log message warns the operator of such fact:

```
A:PE-1>config>service>vpls>site# mesh-sdp-binding
MINOR: SVCMGR #5855 only one object is allowed per site
```

- The **failed-threshold** command defines how many objects should be down for the site to be declared down. This command is obviously only valid for multi-object sites (split-horizon-groups and mesh-sdp-bindings). By default, all the objects in a site must be down for the site to be declared as operationally down.
- The **boot-timer** specifies for how long the service manager waits after a node reboot before running the MH Procedures. The boot-timer value should be configured to allow for the BGP sessions to come up and for the NLRI information to be refreshed/exchanged. In environments with the default BGP MRAI (30 seconds) it is highly recommended to increase this value (for instance, 120 seconds for a normal configuration). The **boot-timer** is only important when a node comes back up and would become the DF.

- The **site-activation-timer** command defines the amount of time the service manager will keep the local objects in standby (in the absence of BGP updates from remote PEs) before running the DF election algorithm to decide whether the site should be unblocked. The timer is started when one of the following events occurs only if the site is operationally up:
 - Manual site activation using the **no shutdown** command at the site-id level or at member object(s) level (SAP(s) or pseudowire(s))
 - Site activation after a failure

The BGP MH election procedures will be resumed upon expiration of this timer or the arrival of a BGP MH update for the multi-homing site.

- The **boot-timer** and **site-activation-timer** commands can be provisioned at service level or at global level. The service level settings have precedence and override the global configuration. The **no** form of the commands at global level, sets the value back to the default values, (10 and 2 seconds, respectively). The **no** form of the commands at service level, makes the timers inherit the global values.

```
*A:PE-1# configure redundancy bgp-multi-homing
- bgp-multi-homing

[no] boot-timer      - Configure BGP multi-homing boot-timer
[no] site-activatio* - Configure BGP multi-homing site activation timer
```

- The **shutdown** command controls the ADMIN state of the site. Note that each site has three possible states:
 - ADMIN state — controlled by the shutdown command.
 - OPERATIONAL state — controlled by the operational status of the individual site objects.
 - DESIGNATED-FORWARDER (DF) state — controlled by the BGP MH election algorithm.

The following CLI output shows the three states for a BGP MH site:

```
*A:MTU-2# show service id 500 site "MH-site-1"
=====
Site Information
=====
Site Name           : MH-site-1
-----
Site Id             : 1
Dest                : sap:1/1/4:500      Mesh-SDP Bind    : no
Admin Status       : Enabled           Oper Status      : up
Designated Fwdr    : No
DF UpTime          : 0d 00:00:00 secs   DF Chg Cnt      : 1
Failed Threshold    : default(all)
=====
```

Core PE Service Configuration

For this example, configure the site MH-site-2 in PE-1, where the site-id is 2 and the object in the site is spoke-sdp 121:500 (pseudowire established from PE-1 to MTU-2).

The following CLI output shows the service configuration for PE-2. Note that the site-id is 2, that is, the same value configured in PE-1. The object defined in PE-2's site is spoke-sdp 221:500 (pseudowire established from PE-2 to MTU-2).

```
A:PE-2>config>service# info
-----
...
    sdp 211 mpls create
        far-end 192.0.2.11
        lsp "LSP-PE-2-MTU-1"
        path-mtu 8000
        keep-alive
        shutdown
    exit
    no shutdown
exit
sdp 221 mpls create
    far-end 192.0.2.21
    lsp "LSP-PE-2-MTU-2"
    path-mtu 8000
    keep-alive
    shutdown
    exit
    no shutdown
exit
sdp 2121 mpls create
    far-end 192.0.2.1
    lsp "LSP-PE-2-PE-1"
    signaling bgp # this sdp will transport BGP-signaled pseudowires
    path-mtu 8000
    keep-alive
    shutdown
    exit
    no shutdown
exit
sdp 2323 mpls create
    far-end 192.0.2.3
    lsp "LSP-PE-2-PE-3"
    signaling bgp # this sdp will transport BGP-signaled pseudowires
    path-mtu 8000
    keep-alive
    shutdown
    exit
    no shutdown
exit
...
    pw-template 500 use-provisioned-sdp create
    exit
...
vpls 500 customer 1 create
    split-horizon-group "CORE" create
    exit
    bgp
```



```
        route-distinguisher 65000:502
        vsi-export "vsi500_export"
        vsi-import "vsi500_import"
        pw-template-bind 500 split-horizon-group "CORE"
    exit
    bgp-vpls
        max-ve-id 65535
        ve-name 502
        ve-id 502
    exit
    no shutdown
    exit
    stp
        shutdown
    exit
    site "MH-site-2" create
        site-id 2
        spoke-sdp 221:500
        no shutdown
    exit
    spoke-sdp 211:500 create
    exit
    spoke-sdp 221:500 create
    exit
    no shutdown
    exit
...

```

MTU Service Configuration

The following CLI output shows the service level configuration on MTU-1.

```
A:MTU-1>config>service# info
...
sdp 111 mpls create
  far-end 192.0.2.1
  lsp "LSP-MTU-1-PE-1"
  path-mtu 8000
  keep-alive
  shutdown
exit
no shutdown
exit
sdp 112 mpls create
  far-end 192.0.2.2
  lsp "LSP-MTU-1-PE-2"
  path-mtu 8000
  keep-alive
  shutdown
exit
no shutdown
exit
...
vpls 500 customer 1 create
  split-horizon-group "site-1" create
  exit
  bgp
    route-distinguisher 65000:511
    route-target export target:65000:500 import target:65000:500
  exit
  stp
    shutdown
  exit
  site "MH-site-1" create
    site-id 1
    split-horizon-group site-1
    no shutdown
  exit
  endpoint "CORE" create
    no suppress-standby-signaling
  exit
  sap 1/1/3:500 split-horizon-group "site-1" create
    eth-cfm
      mep 511 domain 1 association 1 direction down
      fault-propagation-enable use-if-tlv
      ccm-enable
      no shutdown
    exit
  exit
exit
exit
sap 1/1/4:500 split-horizon-group "site-1" create
exit
spoke-sdp 111:500 endpoint "CORE" create
  stp
```

```

        shutdown
    exit
    precedence primary
exit
spoke-sdp 112:500 endpoint "CORE" create
    stp
        shutdown
    exit
exit
no shutdown
exit

```

The MTU-1 node is configured with the following characteristics:

- The BGP context provides the general BGP parameters for service 500 in MTU-1. Note that the route-target command is now used instead of the vsi-import and vsi-export commands. The intent in this example is to configure only the export and import route-targets, and do not need to modify any other attribute. If the local preference are modified (to influence the DF election), a vsi-policy must be configured.
- An A/S pseudowire configuration is used to control the pseudowire redundancy towards the core.
- The multi-homing site, MH-site-1 has a site-id = 1 and a split-horizon-group as an object. The split-horizon-group site-1 is composed of sap 1/1/3:500 and sap 1/1/4:500. As previously discussed, the site will not be declared operationally down until the two SAPs belonging to the site are down. This behavior can be changed by the **failed-threshold** command (for instance, in order to bring the site down when only one object has failed even though the second sap is still up).
- Note that, as an example, a Y.1731 MEP with fault-propagation has been defined in sap 1/1/3:500. As discussed later in the section, this MEP will signal the status of the SAP (as a result of the BGP MH process) to CE-2.

MTU Service Configuration

The service configuration in MTU-2 is shown below.

```
A:MTU-2>config>service# info
...
sdp 211 mpls create
  far-end 192.0.2.1
  lsp "LSP-MTU-2-PE-1"
  path-mtu 8000
  keep-alive
  shutdown
  exit
  no shutdown
exit
sdp 212 mpls create
  far-end 192.0.2.2
  lsp "LSP-MTU-2-PE-2"
  path-mtu 8000
  keep-alive
  shutdown
  exit
  no shutdown
exit
...
vpls 500 customer 1 create
  bgp
    route-distinguisher 65000:521
    route-target export target:65000:500 import target:65000:500
  exit
  stp
    shutdown
  exit
  site "MH-site-1" create
    site-id 1
    sap 1/1/4:500
    no shutdown
  exit
  sap 1/1/4:500 create
  exit
  spoke-sdp 211:500 create
  exit
  spoke-sdp 212:500 create
  exit
  no shutdown
exit
```

Influencing the Designated Forwarder (DF) Decision

As previously explained, assuming that the sites on the two nodes taking part of the same multi-homing site are both up, the two tie-breakers for electing the DF are (in this order):

1. Highest LP
2. Lowest PE ID

The LP by default is 100 in all the routers. Under normal circumstances, if the LP in any router is not changed, MTU-1 will be elected the DF for MH-site-1, whereas PE-1 will be the DF for MH-site-2. Assume in this section that this behavior is changed for MH-site-2 to make PE-2 the DF. Since changing the system address (to make PE-2's ID the lower of the two IDs) is usually not an easy task to accomplish, a vsi-export policy in PE-2 is added to modify the LP with which the MH-site-2 NLRI is announced to PE-1. That policy will change the LP to 150 and as such, since 150 is greater than the default 100 in PE-1, PE-2 will be elected as the DF for MH-site-2. The configuration of the policy is outlined below:

```
A:PE-2# admin display-config
...
#-----
echo "Service Configuration"
#-----
...
    vpls 500 customer 1 create
        ...
            bgp
                route-distinguisher 65000:502
                vsi-export "vsi500_export"
                vsi-import "vsi500_import"
            ...
#-----
echo "Policy Configuration"
#-----
    policy-options
        begin
            community "comm_core" members "target:65000:500"
            policy-statement "vsi500_export"
                entry 10
                    action accept
                    community add "comm_core"
                    local-preference 150
                exit
            exit
        exit
    policy-statement "vsi500_import"
        entry 10
            from
                community "comm_core"
            exit
            action accept
            exit
        exit
```

Influencing the Designated Forwarder (DF) Decision

```
        default-action reject
    exit
```

In PE-1, simply import and export the same route-target without modifying the LP. The configuration is show below.

```
A:PE-1# admin display-config
...
#-----
echo "Service Configuration"
#-----
...
    vpls 500 customer 1 create
        ...
            bgp
                route-distinguisher 65000:501
                vsi-export "vsi500_export"
                vsi-import "vsi500_import"
            ...
#-----
echo "Policy Configuration"
#-----
    policy-options
        begin
            community "comm_core" members "target:65000:500"
            policy-statement "vsi500_export"
                entry 10
                    action accept
                    community add "comm_core"
                exit
            exit
        exit
    policy-statement "vsi500_import"
        entry 10
            from
                community "comm_core"
            exit
            action accept
        exit
    exit
    default-action reject
exit
```

Note that the policy is applied at service 500 level, which means that the LP changes for all the potential multi-homing sites configured under service 500. Therefore, load balancing can be achieved on a per-service basis, but not within the same service.

Another important remark is that these policies are applied on the VPLS 500 for all the potential BGP applications, that is, BGP VPLS, BGP MH and BGP AD. In the example, the LP for the PE-2 BGP updates for BGP MH and BGP VPLS will be set to 150 (note that this however has no impact on BGP VPLS since a given PE cannot receive two BGP VPLS NLRIs with the same VE-ID, such as a different VE-ID per PE within the same VPLS is required).

Black-hole Avoidance

The 7x50 supports the appropriate MAC flush mechanisms for BGP MH, regardless of the protocol being used for the pseudowire signaling:

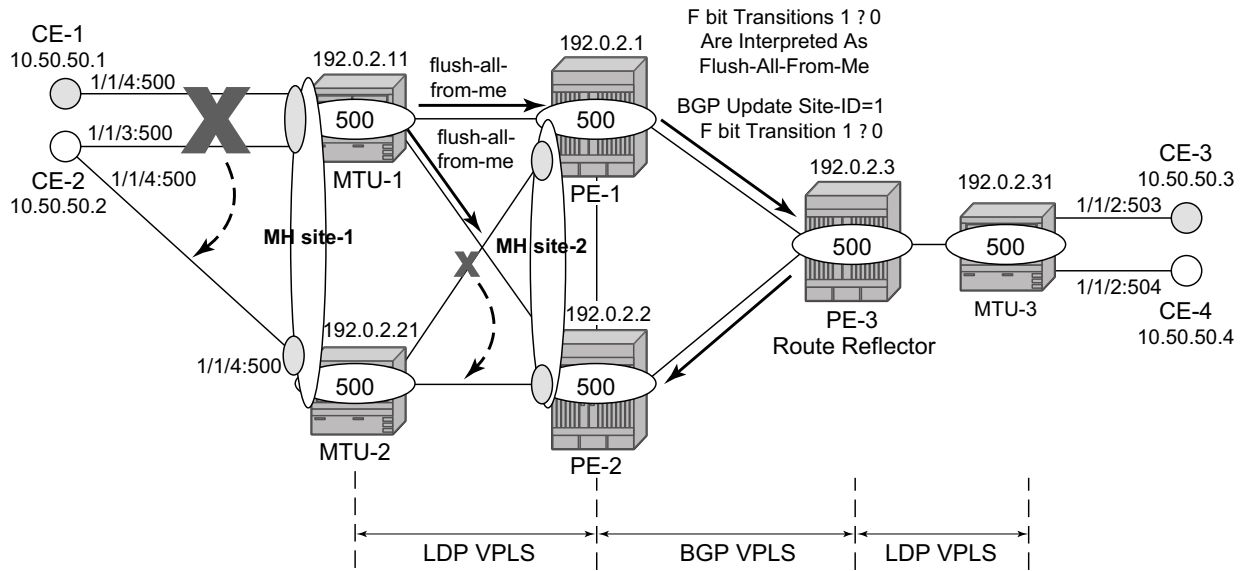
- LDP VPLS — The PE that contains the old DF site (the site that just experienced a DF->non-DF transition) always sends a LDP MAC **flush-all-from-me** to all LDP pseudowires in the VPLS, including the LDP pseudowires associated with the new DF site. No specific configuration is required.
- BGP VPLS — The remote BGP VPLS PEs interpret the F bit transitions from 1 to ZERO as an implicit MAC flush-all-from-me indication. If a BGP update with the flag F=0 is received from the previous DF PE, the remote PEs perform MAC flush-all-from-me, flushing all the MACs associated with the pseudowire to the old DF PE. No specific configuration is required.

Double flushing will not happen as it is expected that between any pair of PEs there will exist only one type of pseudowires, either BGP or LDP pseudowire, but not both types.

In the example, assuming MTU-1 and PE-1 are the DF nodes:

- When MH-site-1 is brought operationally down on MTU-1 (so by default, the two saps must go down unless the **failed-threshold** parameter is changed so that the site is down when only one sap is brought down), MTU-1 will issue a **flush-all-from-me** message.
- When MH-site-2 is brought operationally down on PE-2, a BGP update with F=0 and D=1 is issued by PE-1. PE-2 and PE-3 will receive the update and will flush the MACs learned on the pseudowire to PE-1.

Black-hole Avoidance



OSSG641

Figure 33: MAC Flush for BGP MH

Node failures implicitly trigger a MAC flush on the remote nodes, since the TLDP/BGP session to the failed node goes down.

Access CE/PE Signaling

BGP MH works at service level, therefore no physical ports are torn down on the non-DF but rather the objects are brought down operationally, while the physical port will stay up and used for any other services existing on that port. Due to this reason, there is a need for signaling the standby status of an object to the remote PE or CE. In SR OS 8.0R1 and higher:

- Access PEs running BGP MH on spoke-sdps and elected non-DF, will signal pseudowire standby status (0x20) to the other end. If no pseudowire status is supported on the remote MTU, a label withdrawal is performed⁴. If there is more than one spoke SDP on the site (part of the same SHG), the signaling is sent for all the pseudowires of the site.
- Multi-homed CEs connected through SAPs to the PEs running BGP MH, are signalled by the PEs using Y.1731 CFM, either by stopping the transmission of CCMs or by sending CCMs with ISDOWN (interface status down encoding in the interface status TLV).

In this example, down MEPs on MTU-1 sap 1/1/3:500, sap 1/1/4:500 and MTU-2 sap 1/1/4:500 are configured. Figure 34 shows only the MEPs on MTU-1 sap 1/1/3:500 and CE-2. Upon failure on the MTU-1 site MH-site-1 the MEP 1 will start sending CCMs with interface status down.

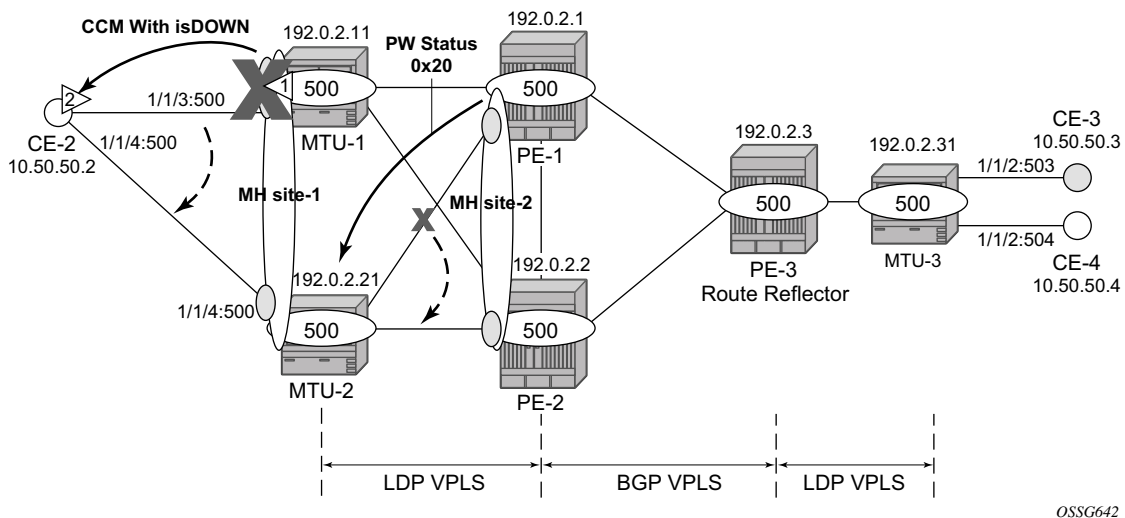


Figure 34: Access PE/CE Signaling

4. Note that the **configure service vpls x spoke-sdp y:z no pw-status-signaling** parameter allows to send a TLDP label-withdrawal instead of pseudowire status bits, even though the peer supports pseudowire status.

The CFM configuration required at sap 1/1/3:500 is outlined below⁵. Down MEPs will be configured on CE-2 and MTU-2 saps in the same way. Note that **fault-propagation-enable use-if-tlv** must be added. In case the CE does not understand the CCM interface status TLV, the **fault-propagation-enable suspend-ccm** option can be enabled instead. This will stop the transmission of CCMs upon site failures.

```
A:MTU-1>config>service# info
...
vpls 500 customer 1 create
  split-horizon-group "site-1" create
  exit
...
  site "MH-site-1" create
    site-id 1
    split-horizon-group site-1
    no shutdown
  exit
...
  sap 1/1/3:500 split-horizon-group "site-1" create
    eth-cfm
      mep 2 domain 1 association 1 direction down
        fault-propagation-enable use-if-tlv
        ccm-enable
        no shutdown
      exit
    exit
  exit
```

If CE-2 is a service router, upon receiving a CCM with ISDOWN, an alarm will be triggered and the SAP will be brought down:

```
56 2010/02/06 15:18:00.62 UTC MINOR: ETH_CFM #2001 Base
"MEP 1/1/2 highest defect is now defMACstatus"

57 2010/02/06 15:18:00.63 UTC MINOR: SVCAGR #2108 vprn502
"Status of interface int-CE2-MTU-1 in service 500 (customer 1) changed to admin=up
oper=down"

58 2010/02/06 15:18:00.63 UTC MINOR: SVCAGR #2203 vprn502
"Status of SAP 1/1/3:500 in service 500 (customer 1) changed to admin=up oper=down
flags=OamDownMEPFault"

59 2010/02/06 15:18:00.63 UTC WARNING: SNMP #2004 vprn500 int-CE2-MTU-1
"Interface int-CE2-MTU-1 is not operational"

*A:CE-2# show service id 500 sap 1/1/3:500
=====
Service Access Points(SAP)
-----
```

5. Detailed configuration guidelines for Y.1731 are out of the scope of this configuration guide.

```

=====
Service Id      : 500
SAP             : 1/1/3:500          Encap           : q-tag
Description    : (Not Specified)
Admin State    : Up                Oper State      : Down
Flags          : OamDownMEPFault
Multi Svc Site : None
Last Status Change : 02/06/2010 15:18:01
Last Mgmt Change  : 02/05/2010 12:21:26
=====

```

As also depicted in [Figure 34](#), PE-1 will signal pseudowire status standby (code 0x20) when PE-1 goes to non-DF state for MH-site-2 MTU-2 will receive that signaling and, based on the **ignore-standby-signaling** parameter, will decide whether to send the broadcast, unknown unicast and multicast (BUM) traffic to PE-1. In case MTU-2 uses in its configuration **ignore-standby-signaling**, it will be sending BUM traffic on both pseudowires at the same time (which is not normally desired), ignoring the pseudowire status bits. The following output shows the MTU-2 spoke **sdp** receiving the pseudowire status signaling. Note that although the spoke SDP stays operationally up, the peer Pw Bits field shows **pwFwdingStandby** and MTU-2 will not send any traffic if the **ignore-standby-signaling** parameter is disabled.

```

*A:MTU-2# show service id 500 sdp 211:500 detail
=====
Service Destination Point (Sdp Id : 211:500) Details
=====
-----
Sdp Id 212:500  -(192.0.2.1)
-----
Description      : (Not Specified)
SDP Id           : 211:500          Type            : Spoke
Spoke Descr     : (Not Specified)
Split Horiz Grp : (Not Specified)
VC Type         : Ether            VC Tag          : n/a
Admin Path MTU  : 8000             Oper Path MTU   : 8000
Far End         : 192.0.2.1        Delivery        : MPLS
Hash Label      : Disabled

Admin State     : Up                Oper State      : Up

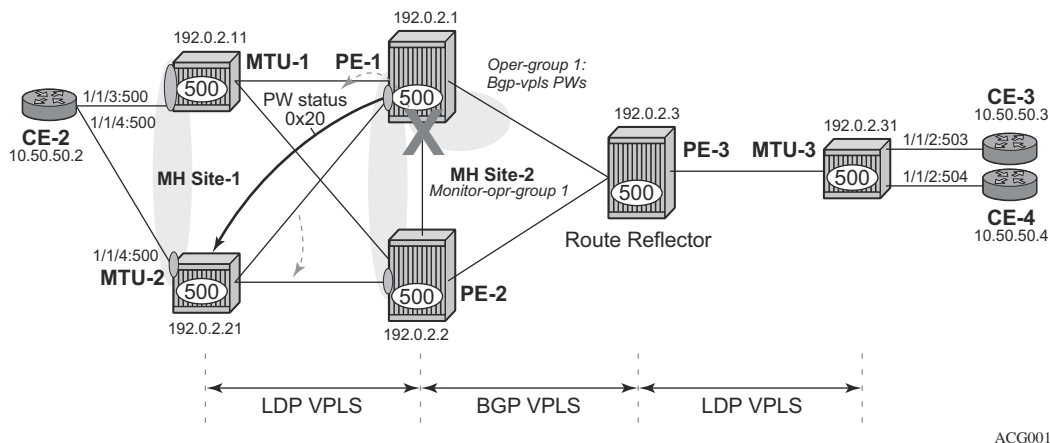
Endpoint       : N/A                Precedence     : 4
Class Fwding State : Down
Flags          : None
Time to RetryReset : never          Retries Left   : 3
Mac Move       : Blockable         Blockable Level : Tertiary
Peer Pw Bits   : pwFwdingStandby

```

Operational Groups for BGP-MH

Operational groups (“oper-groups” in the CLI) introduce the capability of grouping objects into a generic group object and associating its status to other service endpoints (pseudowires, SAPs, IP interfaces) located in the same or in different service instances. The operational group status is derived from the status of the individual components using certain rules specific to the application using the concept. A number of other service entities, the monitoring objects, can be configured to monitor the operational group status and to drive their own status based on the oper-group status. In other words, if the operational group goes down, the monitoring objects will be brought down. When one of the objects included in the oper-group comes up, the entire group will also come up, and therefore so will the monitoring objects.

This concept can be used to enhance the BGP-MH solution by avoiding black-holes on the PE selected as the Designated Forwarder (DF), if the rest of the VPLS endpoints fail (pseudowire spoke(s)/pseudowire mesh and/or SAP(s)). **Figure 35** illustrates the use of oper-groups together with BGP-MH. On PE-1 (and PE-2) all of the BGP-VPLS pseudowires in the core are configured under the same **oper-group group-1**. MH-site-2 is configured as a monitoring object. When the two BGP-VPLS pseudowires go down, **oper-group group-1** will be brought down, therefore MH-site-2 on PE-1 will go down as well (PE-2 will become DF and PE-1 will signal standby to MTU-2).



ACG0016

Figure 35: Oper-Groups and BGP-MH

In the example above, this feature provides a solution to avoid a black-hole when PE-1 loses its connectivity to the core.

Operational groups are configured in two steps:

1. Identify a set of objects whose forwarding state should be considered as a whole group then group them under an operational group (in this case **oper-group group-1**, which is configured in the **bgp pw-template-bind** context).
2. Associate other existing objects (clients) with the oper-group using the **monitor-group** command (configured, in this case, in the **site MH-site-2**).

The following CLI excerpt shows the commands required (**oper-group**, **monitor-oper-group**).

```
A:PE-1>config>service# info
...
oper-group "group-1" create
...
vpls 500 customer 1 create
  split-horizon-group "CORE" create
  exit
  bgp
    route-distinguisher 65000:501
    vsi-export "vsi500_export"
    vsi-import "vsi500_import"
    pw-template-bind 500 split-horizon-group "CORE"
      oper-group "group-1"
    exit
  exit
  bgp-vpls
    max-ve-id 65535
    ve-name 501
      ve-id 501
    exit
    no shutdown
  exit
  stp
    shutdown
  exit
  site "MH-site-2" create
    site-id 2
    spoke-sdp 121:500
    monitor-oper-group "group-1"
    no shutdown
  exit
  spoke-sdp 111:500 create
  exit
  spoke-sdp 121:500 create
  exit
  no shutdown
exit
```

When all the BGP-VPLS pseudowires go down, **oper-group group-1** will go down and therefore the monitoring object, **site MH-site-2**, will also go down and PE-2 will then be elected as DF. The log 99 gives information about this sequence of events:

Operational Groups for BGP-MH

```
30 2012/02/02 17:46:51.58 UTC MINOR: SVCMGR #2542 Base
"Oper-group group-1 changed status to down"
```

```
31 2012/02/02 17:46:51.58 UTC MINOR: SVCMGR #2306 Base
"Status of SDP Bind 121:500 in service 500 (customer 1) changed to admin=up oper=do
wn flags="
```

```
33 2012/02/02 17:46:51.58 UTC WARNING: SVCMGR #2531 Base BGP-MH
"Service-id 500 site MH-site-2 is not the designated-forwarder"
```

Note that the process reverts when at least one BGP-VPLS pseudowire comes back up.

Show Commands and Debugging Options

The main command to find out the status of a given site is the **show service id x site** command. A **detail** modifier is available:

```
*A:MTU-2# show service id 500 site "MH-site-1"
=====
Site Information
=====
Site Name          : MH-site-1
-----
Site Id            : 1
Dest               : sap:1/1/4:500      Mesh-SDP Bind    : no
Admin Status      : Enabled              Oper Status      : up
Designated Fwdr   : No
DF UpTime         : 0d 00:00:00 secs    DF Chg Cnt      : 1
Failed Threshold  : default(all)
=====
*A:MTU-2#
```

```
*A:MTU-2# show service id 500 site "MH-site-1" detail
=====
Site Information
=====
Site Name          : MH-site-1
-----
Site Id            : 1
Dest               : sap:1/1/4:500      Mesh-SDP Bind    : no
Admin Status      : Enabled              Oper Status      : up
Designated Fwdr   : No
DF UpTime         : 0d 00:00:00 secs    DF Chg Cnt      : 1
Boot Timer        : 120 (Ovr)           Timer Remaining   : 0d 00:00:01
Site Activation Timer: 10 (Ovr)
Failed Threshold  : default(all)
=====
*A:MTU-2#
```

Note that the **detail** view of the command displays information about the BGP MH timers. The values are only shown if the global values are overridden by specific ones at service level (and will be tagged with **Ovr** if they have been configured at service level). The **Timer Remaining** field reflects the count down from the boot/site activation timers down to the moment when this router tries to become DF again. Again, this is only shown when the global timers have been overridden by the ones at service level.

It is important to note that the objects on the non-DF site will be brought down operationally and flagged with **StandByForMHPProtocol**.

```
*A:MTU-2# show service id 500 sap 1/1/4:500
=====
Service Access Points(SAP)
=====
Service Id        : 500
```

Show Commands and Debugging Options

```
SAP                : 1/1/4:500                Encap                : q-tag
Description        : (Not Specified)
Admin State        : Up                      Oper State           : Down
Flags              : StandByForMHPProtocol
Multi Svc Site     : None
Last Status Change : 02/06/2010 14:28:37
Last Mgmt Change   : 02/06/2010 14:21:30
```

```
=====
*A:PE-2# show service id 500 sdp 221:500 detail
```

```
=====
Service Destination Point (Sdp Id : 221:500) Details
=====
```

```
-----
Sdp Id 221:500  -(192.0.2.21)
-----
```

```
Description        : (Not Specified)
SDP Id              : 221:500                Type                 : Spoke
Admin State         : Up                      Oper State           : Down
Flags               : StandbyForMHPProtocol
```

The BGP MH routes in the RIB, RIB-IN and RIB-OUT can be shown by using the corresponding **show router bgp routes** and **show router bgp neighbor x.x.x.x received-routes|advertised-routes** commands. Note that the BGP MH routes are only shown when the operator uses the **l2-vpn family** modifier. Should the operator want to filter only the BGP MH routes out of the l2-vpn routes, the **multi-homing** filter has to be added to the **show router bgp routes** commands.

```
*A:PE-3# show router bgp routes l2-vpn multi-homing site id 1 detail
```

```
=====
BGP L2VPN-MULTIHOME Routes
=====
```

```
Route Type        : MultiHome
Route Dist.       : 65000:511
Site Id           : 1
Nexthop           : 192.0.2.11
From              : 192.0.2.11
Res. Nexthop      : n/a
Local Pref.       : 100                      Interface Name       : NotAvailable
Aggregator AS    : None                      Aggregator          : None
Atomic Aggr.     : Not Atomic                MED                 : 0
Community        : target:65000:500 l2-vpn:Encap=19: Flags=-DF: MTU=0
Cluster          : No Cluster Members
Originator Id    : None                      Peer Router Id      : 192.0.2.11
Flags             : Used Valid Best IGP
AS-Path          : No As-Path
```

```
Route Type        : MultiHome
Route Dist.       : 65000:521
Site Id           : 1
Nexthop           : 192.0.2.21
From              : 192.0.2.21
Res. Nexthop      : n/a
Local Pref.       : 100                      Interface Name       : NotAvailable
Aggregator AS    : None                      Aggregator          : None
```



```
Atomic Aggr.   : Not Atomic           MED           : 0
Community     : target:65000:500 l2-vpn:Encap=19: Flags=none: MTU=0
Cluster       : No Cluster Members
Originator Id : None                 Peer Router Id : 192.0.2.21
Flags         : Used Valid Best IGP
AS-Path       : No As-Path
```

A new command was added in to show the Layer 2 BGP routes.

```
*A:PE-1# show service l2-route-table
- l2-route-table [detail] [bgp-ad] [multi-homing] [bgp-vpls]
<detail>           : keyword - display detailed information

*A:PE-1# show service l2-route-table multi-homing
=====
Services: L2 Multi-Homing Route Information - Summary
=====
Svc Id      L2-Routes (RD-Prefix)      Next Hop      SiteId      State      DF
-----
500         65000:511                     192.0.2.11   1           up(0)     set
500         65000:521                     192.0.2.21   1           up(0)     clear
500         65000:502                     192.0.2.2    2           up(0)     clear
-----
No. of L2 Multi-Homing Route Entries: 3
=====
```

Finally, in terms of debugging, the recommendation would be to check the following CLI sources:

- **log-id 99** — Provides information about the site object changes and DF changes.
- **debug router bgp update** — Shows the BGP updates for BGP MH, including the sent and received BGP MH NLRIs and flags.
- **debug router ldp** commands — Provides information about the pseudowire status bits being signalled as well as the MAC flush messages.

```
*A:MTU-1# debug router ldp peer 192.0.2.1 packet init detail
*A:MTU-1# debug router ldp peer 192.0.2.1 packet label detail
```

As an example, log-id 99 and debug output is displayed below after shutting down MH-site-1 on MTU-1:

```
*A:MTU-1# configure service vpls 500 sap 1/1/4:500 shutdown
*A:MTU-1# configure service vpls 500 sap 1/1/3:500 shutdown

6047 2010/02/06 14:46:04.05 UTC MINOR: SVCMGR #2203 Base
"Status of SAP 1/1/3:500 in service 500 (customer 1) changed to admin=down oper=down
flags=SapAdminDown "
```

```
6046 2010/02/06 14:46:04.05 UTC WARNING: SVCMGR #2531 Base BGP-MH
"Service-id 500 site MH-site-1 is not the designated-forwarder"
```

Show Commands and Debugging Options

```
12 2010/02/06 14:46:04.04 UTC MINOR: DEBUG #2001 Base Peer 1: 192.0.2.3
"Peer 1: 192.0.2.3: UPDATE
Peer 1: 192.0.2.3 - Send BGP UPDATE:

    Flag: 0x40 Type: 5 Len: 4 Local Preference: 100
    Flag: 0xc0 Type: 16 Len: 16 Extended Community:
        target:65000:500
        12-vpn:Encap=19: Flags=D: MTU=0
    Flag: 0x90 Type: 14 Len: 28 Multiprotocol Reachable NLRI:
        Address Family L2VPN
        NextHop len 4 NextHop 192.0.2.11
        [MH] site-id: 1, RD 65000:511
"
13 2010/02/06 14:46:04.04 UTC MINOR: DEBUG #2001 Base LDP
"LDP: LDP
Send Address Withdraw packet (msgId 434) to 192.0.2.1:0
MAC Flush (All MACs learned from me)
Service FEC EpipePWE3: ENET(5)/500 Group ID = 0 cBit = 0
"
14 2010/02/06 14:46:04.05 UTC MINOR: DEBUG #2001 Base Peer 1: 192.0.2.3
"Peer 1: 192.0.2.3: UPDATE
Peer 1: 192.0.2.3 - Received BGP UPDATE:
    Withdrawn Length = 0
    Total Path Attr Length = 86

    Flag: 0x40 Type: 5 Len: 4 Local Preference: 100
    Flag: 0x80 Type: 9 Len: 4 Originator ID: 192.0.2.21
    Flag: 0x80 Type: 10 Len: 4 Cluster ID:
        1.1.1.1
    Flag: 0xc0 Type: 16 Len: 16 Extended Community:
        target:65000:500
        12-vpn:Encap=19: Flags=-DF: MTU=0
    Flag: 0x90 Type: 14 Len: 28 Multiprotocol Reachable NLRI:
        Address Family L2VPN
        NextHop len 4 NextHop 192.0.2.21
        [MH] site-id: 1, RD 65000:521
"
```

Note that assuming all the recommended tools are enabled, a DF to non-DF transition can be shown as well as the corresponding MAC flush messages and related BGP processing.

If MH-site-2 is torn down on PE-1, the **debug router bgp update** command would allow us to see two BGP updates from PE-1:

- A BGP MH update for site-id 2 with flag D set (since the site is down).
- A BGP VPLS update for veid=501 and flag D set. This is due to the fact that there are no more active objects on the VPLS, besides the BGP pseudowires.

```
*A:PE-1>config>service>vpls# spoke-sdp 121:500 shutdown
11 2010/02/20 11:21:55.49 UTC MINOR: DEBUG #2001 Base Peer 1: 192.0.2.3
"Peer 1: 192.0.2.3: UPDATE
Peer 1: 192.0.2.3 - Send BGP UPDATE:
    Withdrawn Length = 0
    Total Path Attr Length = 72
    Flag: 0x40 Type: 1 Len: 1 Origin: 0
```

```

Flag: 0x40 Type: 2 Len: 0 AS Path:
Flag: 0x80 Type: 4 Len: 4 MED: 0
Flag: 0x40 Type: 5 Len: 4 Local Preference: 100
Flag: 0xc0 Type: 16 Len: 16 Extended Community:
    target:65000:500
    l2-vpn:Encap=19: Flags=D: MTU=0
Flag: 0x90 Type: 14 Len: 28 Multiprotocol Reachable NLRI:
    Address Family L2VPN
    NextHop len 4 NextHop 192.0.2.1
    [MH] site-id: 2, RD 65000:501
12 2010/02/20 11:21:55.49 UTC MINOR: DEBUG #2001 Base Peer 1: 192.0.2.3
"Peer 1: 192.0.2.3: UPDATE
Peer 1: 192.0.2.3 - Send BGP UPDATE:
    Withdrawn Length = 0
    Total Path Attr Length = 72
    Flag: 0x40 Type: 1 Len: 1 Origin: 0
    Flag: 0x40 Type: 2 Len: 0 AS Path:
    Flag: 0x80 Type: 4 Len: 4 MED: 0
    Flag: 0x40 Type: 5 Len: 4 Local Preference: 100
    Flag: 0xc0 Type: 16 Len: 16 Extended Community:
        target:65000:500
        l2-vpn:Encap=19: Flags=D: MTU=1514
    Flag: 0x90 Type: 14 Len: 28 Multiprotocol Reachable NLRI:
        Address Family L2VPN
        NextHop len 4 NextHop 192.0.2.1
        [VPLS] veid: 501, vbo: 497, vbs: 8, label-base: 131032, RD 65000:501

```

The D flag, sent along with the BGP VPLS update for veid 501, would be seen on the remote core PEs as though it was a pseudowire status fault (although there is no TLDP running in the core).

```

*A:PE-2# show service id 500 all | match Flag
Flags          : None
Flags          : None
Flags          : PWPeerFaultStatusBits
Flags          : None

```

When oper-groups are configured (as previously shown), the following **show** command helps to find the operational dependencies between monitoring objects and group objects.

```

*A:PE-1# show service oper-group "group-1" detail
=====
Service Oper Group Information
=====
Oper Group          : group-1
Creation Origin     : manual           Oper Status        : up
Hold DownTime      : 0 secs           Hold UpTime        : 4 secs
Members            : 1                 Monitoring         : 1
=====
Member SDF-Binds for OperGroup: group-1
=====
SdpId              SvcId      Type IP address      Adm   Opr
-----
1212:4294967293   500        Bgp* 192.0.2.2      Up    Up
1313:4294967294   500        Bgp* 192.0.2.3      Up    Up

```

Show Commands and Debugging Options

```
-----  
SDP Entries found: 2  
=====
```

* indicates that the corresponding row element may have been truncated.

```
=====
```

Monitoring Sites for OperGroup: group-1

```
=====
```

SvcId	Site	Site-Id	Dest	Admin	Oper	Fwdr
500	MH-site-2	2	sdp:121:500	Enabled	up	Yes

```
-----  
Site Entries found: 1  
=====
```

Conclusion

The SR OS supports a wide range of service resiliency options as well as the best-of-breed system level HA and MPLS mechanisms for the access and the core. BGP MH for VPLS completes the service resiliency toolset by adding a mechanism that has some good advantages over the alternative solutions:

- BGP MH provides a common resiliency mechanism for attachment circuits (SAPs), pseudowires (spoke SDPs), split horizon groups and mesh bindings
- BGP MH is a network-based technique which does not need interaction to the CE or MTU to which it is providing redundancy to.

The examples used in this section illustrate the configuration of BGP MH for access CEs and MTUs. Show and debug commands have also been suggested so that the operator can verify and troubleshoot the BGP MH procedures.

Conclusion