

---

## In This Chapter

This section describes advanced BGP VPLS configurations.

Topics in this section include:

- [Applicability on page 210](#)
- [Summary on page 211](#)
- [Overview on page 212](#)
- [Configuration on page 214](#)
- [Conclusion on page 248](#)

## **Applicability**

This example is applicable to all of the 7450-ESS, 7750-SR and 7710-SR series and was tested on Release 12.0.R1. There are no pre-requisites for this configuration.

## Summary

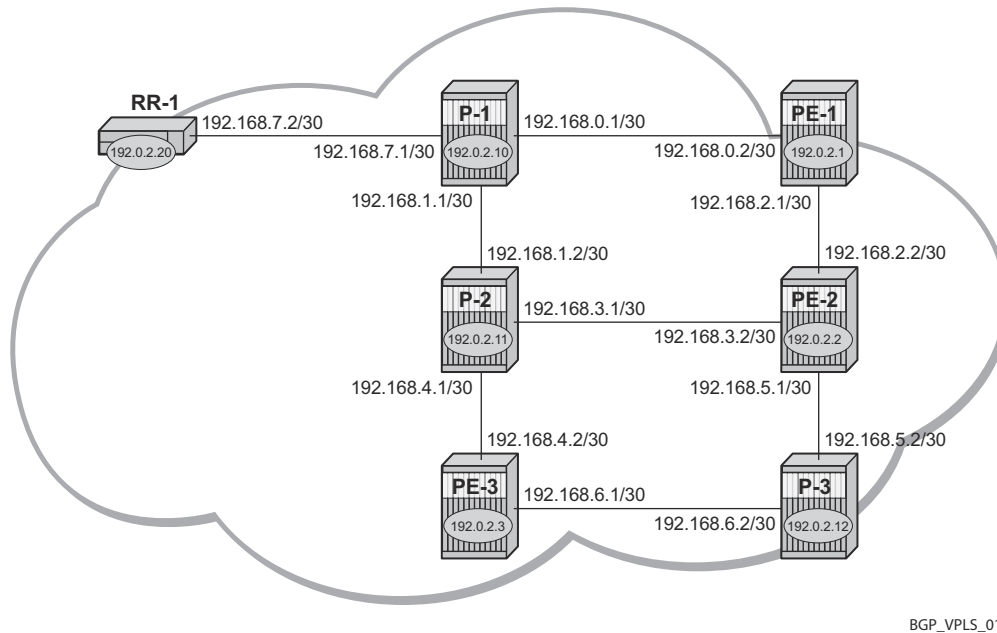
There are currently two IETF standards for the provisioning of Virtual Private LAN Services (VPLS). RFC 4762, *Virtual Private LAN Service (VPLS) Using Label Distribution Protocol (LDP) Signaling*, describes Label Distribution Protocol (LDP) VPLS, where VPLS pseudowires are signaled using LDP between VPLS Provider Edge (PE) routers, either configured manually or auto-discovered using BGP.

RFC 4761, *Virtual Private LAN Service (VPLS) Using BGP for Auto-Discovery and Signaling*, describes the use of Border Gateway Protocol (BGP) for both the auto-discovery of VPLS PEs and signaling of pseudowires between such PEs.

The purpose of this section is to describe the configuration and troubleshooting for BGP-VPLS.

Knowledge of BGP-VPLS RFC 4761 architecture and functionality is assumed throughout this section, as well as knowledge of Multi-Protocol BGP.

## Overview



**Figure 36: Network Topology**

The network topology is displayed in [Figure 36](#). The configuration uses seven 7750/7450/7710 Service Router (SR) nodes located in the same Autonomous System (AS). There are three Provider Edge (PE) routers, and RR-1 will act as a Route Reflector (RR) for the AS. The PE routers are all VPLS-aware, the Provider (P) routers are VPLS unaware and do not take part in the BGP process.

The following configuration tasks should be completed as a pre-requisite:

- ISIS or OSPF on each of the network interfaces between the PE/P routers and RR.
- MPLS should be configured on all interfaces between PE routers and P routers. MPLS is not required between P-1 and RR-1.
- LDP should be configured on interfaces between PE and P routers. It is not required between P-1 and the RR-1.
- The RSVP protocol should be enabled.

## BGP VPLS

In this architecture a VPLS instance is a collection of local VPLS instances present on a number of PEs in a provider network. In this context, any VPLS-aware PE is also known as a VPLS Edge (VE) device.

The PEs communicate with each other at the control plane level by means of BGP updates containing BGP-VPLS Network Layer Reachability Information (NLRI). Each update contains enough information for a PE to determine the presence of other local VPLS instances on peering PEs and to set-up pseudowire connectivity for data flow between peers containing a local VPLS within the same VPLS instance. Therefore, auto-discovery and pseudowire signaling are achieved using a single BGP update message.

Each PE within a VPLS instance is identified by a VPLS Edge identifier (ve-id) and the presence of a VPLS instance is determined using the exchange of standard BGP extended community route targets between PEs.

Each PE will advertise, via the route reflectors, the presence of each VPLS instance to all other PEs, along with a block of multiplexer labels that can be used to communicate between such instances plus a BGP next hop that determines a labelled transport tunnel between PEs.

Each VPLS instance is configured with import and export route target extended communities for topology control, along with VE identification.

## Configuration

The first step is to configure an MP-iBGP session between each of the PEs and the RR.

The configuration for PE-1 is as follows:

```
configure router bgp
  group internal
    family l2vpn
    type internal
    peer-as 65536
    neighbor 192.0.2.20
    exit
  exit
  no shutdown
exit
```

The configuration for the other PE nodes is very similar. The IP addresses can be derived from [Figure 36](#).

The configuration for RR-1 is as follows:

```
configure router bgp
  cluster 1.1.1.1
  group rr-internal
    family l2-vpn
    type internal
    peer-as 65536
    neighbor 192.0.2.1
    exit
    neighbor 192.0.2.2
    exit
    neighbor 192.0.2.3
    exit
  exit
  no shutdown
exit
```

On PE-1, verify that the BGP session with RR-1 is established with address family l2-vpn capability negotiated:

```
*A:PE-1# show router bgp neighbor 192.0.2.20
=====
BGP Neighbor
=====
-----
Peer   : 192.0.2.20
Group  : internal
-----
Peer AS      : 65536           Peer Port      : 179
Peer Address : 192.0.2.20
Local AS     : 65536           Local Port     : 50622
Local Address: 192.0.2.1
```

```

Peer Type           : Internal
State               : Established      Last State           : Active
Last Event          : recvKeepAlive
Last Error          : Unrecognized Error
Local Family        : L2-VPN
Remote Family       : L2-VPN
Hold Time           : 90                Keep Alive           : 30
Min Hold Time       : 0
Active Hold Time    : 90                Active Keep Alive    : 30
Cluster Id          : None
Preference          : 170                Num of Update Flaps : 17
Recd. Paths         : 27
IPv4 Recd. Prefixes : 0                IPv4 Active Prefixes : 0
IPv4 Suppressed Pfxs : 0                VPN-IPv4 Suppr. Pfxs : 0
VPN-IPv4 Recd. Pfxs : 0                VPN-IPv4 Active Pfxs : 0
Mc IPv4 Recd. Pfxs. : 0                Mc IPv4 Active Pfxs. : 0
Mc IPv4 Suppr. Pfxs : 0                IPv6 Suppressed Pfxs : 0
IPv6 Recd. Prefixes : 0                IPv6 Active Prefixes : 0
VPN-IPv6 Recd. Pfxs : 0                VPN-IPv6 Active Pfxs : 0
VPN-IPv6 Suppr. Pfxs : 0
Mc IPv6 Recd. Pfxs. : 0                Mc IPv6 Active Pfxs. : 0
Mc IPv6 Suppr. Pfxs : 0                L2-VPN Suppr. Pfxs  : 0
L2-VPN Recd. Pfxs  : 21                L2-VPN Active Pfxs  : 14
MVPN-IPv4 Suppr. Pfxs : 0                MVPN-IPv4 Recd. Pfxs : 0
MVPN-IPv4 Active Pfxs : 0                MDT-SAFI Suppr. Pfxs : 0
MDT-SAFI Recd. Pfxs  : 0                MDT-SAFI Active Pfxs : 0
Flow-IPv4 Suppr. Pfxs : 0                Flow-IPv4 Recd. Pfxs : 0
Flow-IPv4 Active Pfxs : 0                Rte-Tgt Suppr. Pfxs  : 0
Rte-Tgt Recd. Pfxs  : 0                Rte-Tgt Active Pfxs  : 0
Backup IPv4 Pfxs    : 0                Backup IPv6 Pfxs     : 0
Mc Vpn Ipv4 Recd. Pf* : 0                Mc Vpn Ipv4 Active P* : 0
Backup Vpn IPv4 Pfxs : 0                Backup Vpn IPv6 Pfxs : 0
Input Queue         : 0                Output Queue         : 0
i/p Messages        : 142                o/p Messages         : 104
i/p Octets          : 7224                o/p Octets           : 2731
i/p Updates         : 49                o/p Updates          : 11
MVPN-IPv6 Suppr. Pfxs : 0                MVPN-IPv6 Recd. Pfxs : 0
MVPN-IPv6 Active Pfxs : 0
Flow-IPv6 Suppr. Pfxs : 0                Flow-IPv6 Recd. Pfxs : 0
Flow-IPv6 Active Pfxs : 0
Evpn Suppr. Pfxs   : 0                Evpn Recd. Pfxs     : 0
Evpn Active Pfxs   : 0
TTL Security        : Disabled                Min TTL Value        : n/a
Graceful Restart    : Disabled                Stale Routes Time    : n/a
Restart Time        : n/a
Advertise Inactive  : Disabled                Peer Tracking        : Disabled
Advertise Label     : None
Auth key chain      : n/a
Disable Cap Nego    : Disabled                Bfd Enabled          : Disabled
Flowspec Validate   : Disabled                Default Route Tgt    : Disabled
Aigp Metric         : Disabled                Split Horizon        : Disabled
Damp Peer Oscillatio* : Disabled                Update Errors        : 0
GR Notification     : Disabled                Fault Tolerance      : Disabled
Rem Idle Hold Time  : 00h00m00s
Next-Hop Unchanged  : None
L2 VPN Cisco Interop : Disabled
Local Capability    : RtRefresh MPBGP 4byte ASN
Remote Capability   : RtRefresh MPBGP 4byte ASN
Local AddPath Capabi* : Disabled

```

## Configuration

```
Remote AddPath Capab*: Send - None
                        : Receive - None
Import Policy         : None Specified / Inherited
Export Policy        : None Specified / Inherited
```

```
-----
Neighbors : 1
=====
```

```
* indicates that the corresponding row element may have been truncated.
*A:PE-1#
```

On RR-1, show that BGP sessions with each PE are established, and have a negotiated the l2-vpn address family capability.

```
A:RR-1# show router bgp summary
```

```
=====
BGP Router ID:192.0.2.20      AS:65536      Local AS:65536
=====
BGP Admin State      : Up          BGP Oper State      : Up
Total Peer Groups    : 1           Total Peers         : 3
Total BGP Paths      : 28         Total Path Memory   : 5376
Total IPv4 Remote Rts : 0         Total IPv4 Rem. Active Rts : 0
Total McIPv4 Remote Rts : 0       Total McIPv4 Rem. Active Rts : 0
Total McIPv6 Remote Rts : 0       Total McIPv6 Rem. Active Rts : 0
Total IPv6 Remote Rts : 0         Total IPv6 Rem. Active Rts : 0
Total IPv4 Backup Rts : 0         Total IPv6 Backup Rts  : 0

Total Supressed Rts  : 0           Total Hist. Rts     : 0
Total Decay Rts      : 0

Total VPN Peer Groups : 0         Total VPN Peers     : 0
Total VPN Local Rts   : 0
Total VPN-IPv4 Rem. Rts : 0       Total VPN-IPv4 Rem. Act. Rts : 0
Total VPN-IPv6 Rem. Rts : 0       Total VPN-IPv6 Rem. Act. Rts : 0
Total VPN-IPv4 Bkup Rts : 0       Total VPN-IPv6 Bkup Rts  : 0

Total VPN Supp. Rts   : 0         Total VPN Hist. Rts  : 0
Total VPN Decay Rts   : 0

Total L2-VPN Rem. Rts : 24         Total L2VPN Rem. Act. Rts : 0
Total MVPN-IPv4 Rem Rts : 0       Total MVPN-IPv4 Rem Act Rts : 0
Total MDT-SAFI Rem Rts : 0         Total MDT-SAFI Rem Act Rts : 0
Total MSPW Rem Rts     : 0         Total MSPW Rem Act Rts     : 0
Total RouteTgt Rem Rts : 0         Total RouteTgt Rem Act Rts : 0
Total McVpnIPv4 Rem Rts : 0       Total McVpnIPv4 Rem Act Rts : 0
Total MVPN-IPv6 Rem Rts : 0       Total MVPN-IPv6 Rem Act Rts : 0
Total EVPN Rem Rts     : 0         Total EVPN Rem Act Rts     : 0
Total FlowIpv4 Rem Rts : 0         Total FlowIpv4 Rem Act Rts : 0
Total FlowIpv6 Rem Rts : 0         Total FlowIpv6 Rem Act Rts : 0
=====
```

```
BGP Summary
```

```
=====
Neighbor
          AS PktRcvd InQ Up/Down State|Rcv/Act/Sent (Addr Family)
          PktSent OutQ
-----
192.0.2.1
          65536 169 0 00h45m51s 7/0/24 (L2VPN)
```



```

192.0.2.2          148    0
                  65536   167   0 00h25m28s 10/0/24 (L2VPN)
                  86     0
192.0.2.3          65536   167   0 01h16m16s 7/0/24 (L2VPN)
                  213    0

```

```
-----
A:RR-1#
```

Configure a full mesh of RSVP-TE LSPs between PE routers.

The MPLS interface and LSP configuration for PE-1 are:

```

A:PE-1# configure router mpls
      interface "system"
      exit
      interface "int-PE-1-P-1"
      exit
      interface "int-PE-1-PE-2"
      exit
      path "loose"
        no shutdown
      exit
      lsp "LSP-PE-1-PE-2"
        to 192.0.2.2
        primary "loose"
      exit
      no shutdown
    exit
      lsp "LSP-PE-1-PE-3"
        to 192.0.2.3
        primary "loose"
      exit
      no shutdown
    exit
  no shutdown

```

The MPLS and LSP configuration for PE-2 and PE-3 are similar to that of PE-1 with the appropriate interfaces and LSP names configured.

## BGP VPLS PE Configuration

---

### Pseudowire Templates

Pseudowire templates are used by BGP to dynamically instantiate SDP bindings, for a given service, to signal the egress service de-multiplexer labels used by remote PEs to reach the local PE.

The template determines the signaling parameters of the pseudowire, control word presence, MAC-pinning, filters etc., plus other usage characteristics such as split horizon groups.

The MPLS transport tunnel between PEs can be signaled using LDP or RSVP-TE.

LDP based pseudowires can be automatically instantiated. RSVP-TE based SDPs have to be pre-provisioned.

---

### Pseudowire Templates for Auto-SDP Creation Using LDP

In order to use an LDP transport tunnel for data flow between PEs, it is necessary for link layer LDP to be configured between all PEs/Ps, so that a transport label for each PE's system interface is available.

```
A:PE-1# configure router ldp
      interface-parameters
        interface "int-PE-1-P-1"
        exit
        interface "int-PE-1-PE-2"
        exit
      exit
      targeted-session
      exit
      no shutdown
```

Using this mechanism SDPs can be auto-instantiated with SDP-ids starting at the higher end of the SDP numbering range, such as 17407. Any subsequent SDPs created use SDP-ids decrementing from this value.

A pseudowire template is required containing a split horizon group. Each SDP created with this template is contained within a split horizon group so that traffic cannot be forwarded between them.

```
A:PE-1# configure service
      pw-template 1 create
        split-horizon-group "vpls-shg"
      exit
```

```
exit
```

The pseudowire template also has the following options available when used for BGP-VPLS:

```
*A:PE-1# configure service pw-template
      [no] controlword
      force-vlan-vc-forwarding
      vc-type {ether | vlan}
```

- The control word will determine whether the C flag is set in the Layer 2 extended community and, therefore, if a control word is used in the pseudowire.
- The encap type in the Layer 2 extended community is always 19 (VPLS encap), therefore, the vc-type will always be **ether** regardless of the configured value on the vc-type.
- The **force-vlan-vc-forwarding** command will add a tag (equivalent to **vc-type vlan**) and will allow for customer QoS transparency (dot1p+DE bits).

## Pseudowire Templates for Provisioned SDPs using RSVP-TE


To use an RSVP-TE tunnel as transport between PEs, it is necessary to bind the RSVP-TE LSP between PEs to an SDP.

SDP creation from PE-1 to PE-2 is shown below:

```
A:PE-1# configure service sdp 33 mpls create
      description "from-192.0.2.1-id-33"
      far-end 192.0.2.2
      lsp "LSP-PE-1-PE-2"
      signaling bgp
      keep-alive
      shutdown
      exit
      no shutdown
```

Note that the **signaling bgp** parameter is required for BGP-VPLS to be able to use this SDP. Conversely, SDPs that are bound to RSVP-based LSPs with signaling set to the default value of **tldp** will not be used as SDPs within BGP-VPLS.

## BGP VPLS Using Auto-Provisioned SDPs



**Figure 37: BGP VPLS Using Auto-Provisioned SDPs**

Figure 37 shows a VPLS instance where SDPs are auto-provisioned. In this case, the transport tunnels are LDP signaled.

The following output shows the configuration required on PE-1 for a BGP-VPLS service using a pseudowire template configured for auto-provisioning of SDPs.

```
A:PE-1# configure service vpls 1 customer 1 create
      bgp
        route-distinguisher 65536:1
        route-target export target:65536:1 import target:65536:1
        pw-template-bind 1
      exit
      bgp-vpls
        max-ve-id 10
        ve-name PE-1
        ve-id 1
      exit
      no shutdown
    exit
  stp
    shutdown
  exit
  sap 1/1/4:1.0 create
  exit
  no shutdown
exit
```

The **bgp** context specifies parameters which are valid for all of the VPLS BGP applications, such as BGP-multi-homing, BGP-auto-discovery as well as BGP-VPLS.

Within the **bgp** context, parameters are configured that are used by neighboring PEs to determine membership of a given VPLS instance, such as the auto-discovery of PEs containing the same VPLS instance; the route-distinguisher is configured, along with the route target extended communities.

Route target communities are used to determine membership of a given VPLS instance. Note that the import route target at the BGP level is mandatory. The pseudowire template bind is then applied by the service manager on the received routes matching the route target value.

Also within the **bgp-vpls** context, the signaling parameters are configured. These determine the service labels required for the data plane of the VPLS instance.

The VPLS edge ID (ve-id) is a numerical value assigned to each PE within a VPLS instance. This value should be unique for a given VPLS instance, no two PEs within the same instance should have the same ve-id values.

Note that a more specific route target can be applied to a pseudowire template in order to define a specific pseudowire topology, rather than only a full mesh, using the command within the **bgp** context:

**pw-template** *template-id* [**split-horizon-group** *groupname*] [**import-rt** *import-rt-value* (up to 5 max)]

It is also worth noting that changes to the import policies are not taken once the pseudowire has been setup (changes on route-target are refreshed though). Pseudowire templates can be re-evaluated with the command **tools perform eval-pw-template**. The **eval-pw-template** command checks whether all the bindings using this pseudowire template policy are still meant to use this policy.

If the policy has changed and **allow-service-impact** is true, then the old binding is removed and it is re-added with the new template.

## VE-ID and BGP Label Allocations

The choice of `ve-id` is crucial in ensuring efficient allocation of de-multiplexer labels. The most efficient choice is for `ve-ids` to be allocated starting at 1 and incrementing for each PE as the following section explains.

The **max-ve-id** *value* determines the range of the `ve-id` value that can be configured. If a PE receives a BGP-VPLS update containing a `ve-id` with a greater value than the configured **max-ve-id**, then the update is dropped and no service labels are installed for this `ve-id`.

The **max-ve-id** command also checks the locally-configured `ve-id`, and prevents a higher value from being used.

Each PE allocates blocks of labels per VPLS instance to remote PEs, in increments of eight labels. It achieves this by advertising three parameters in a BGP update message,

- A label base (LB) which is the lowest label in the block
- A VE Block size (VBS) which is always eight labels, and cannot be changed
- A VE base offset (VBO).

This defines a block of labels in the range (LB, LB+1, ..., LB+VBS-1).

As an example, if the label base (LB) = 131055, then the range for the block is 131055 to 131062, which is exactly eight labels, as per the block size.

(The last label in the block is calculated as  $131055+8-1 = 131062$ )

The label allocated by the PE to each remote PE within the VPLS is chosen from this block and is determined by its `ve-id`. In this way, each remote PE has a unique de-multiplexer label for that VPLS.

To reduce label wastage, contiguous `ve-ids` in the range (N..N+7) per VPLS should be chosen, where  $N > 0$ .

Assuming a collection of PEs with contiguous `ve-ids`, the following labels will be chosen by PEs from the label block allocated by PE-1 which has a `ve-id` = 1.

**Table 2: VE-IDs and Labels**

VE-ID	Label
2	131056
3	131057
4	131058

**Table 2: VE-IDs and Labels (Continued)**

VE-ID	Label
5	131059
6	131060
7	131061
8	131062

This shows that the label allocated to a given PE is (LB+veid-1). The “1” is the VE block offset (VBO).

This means that the label allocated to a PE router within the VPLS can now be written as (LB + veid - VBO), which means that (ve-id - VBO) calculation must always be at least zero and be less than the block size, which is always 8.

For  $ve-id \leq 8$ , a label will be allocated from this block.

For the next block of 8 ve-ids (ve-id 9 to ve-id 16) a new block of 8 labels must be allocated, so a new BGP update is sent, with a new label base, and a block offset of 9.

Table 3 shows how the choice of ve-ids can affect the number of label blocks allocated, and hence the number of labels:

**Table 3: VE-IDs and Number of Labels**

VE-ID	Block Offset	Labels Allocated
1-8	1	8
9-16	9	8
17-24	17	8
25-32	25	8
33-40	33	8
41-48	41	8
49-56	49	8

## BGP VPLS PE Configuration

This shows that the most efficient use of labels occurs when the ve-ids for a set of PEs are chosen from the same block offset.

Note that if ve-ids are chosen that map to different block offsets, then each PE will have to send multiple BGP updates to signal service labels. Each PE sends label blocks in BGP updates to each of its BGP neighbors for all label blocks in which at least one ve-id has been seen by this PE (it does not advertise label blocks which do not contain an active ve-id, where active ve-id means the ve-id of this PE or any other PE in this VPLS).

The **max-ve-id** must be configured first, and determines the maximum value of the ve-id that can be configured within the PE. The ve-id value cannot be higher than this within the PE configuration,  $ve-id \leq max-ve-id$ . Similarly, if the ve-id within a received NLRI is higher than the **max-ve-id value**, it will not be accepted as valid consequently the max-ve-id configured on all PEs must be greater than or equal to any ve-id used in the VPLS.

Only one ve-id value can be configured. If the ve-id value is changed, BGP withdraws the NLRI and sends a route-refresh.

Note that if the same ve-id is used in different PEs for the same VPLS, a Designated Forwarder election takes place.

Executing the **shutdown** command triggers an MP-UNREACH-NLRI from the PE to all BGP peers

The **no shutdown** command triggers an MP-REACH-NLRI to the same peers.



## PE-2 Service Creation

On PE-2 create VPLS service using pseudowire template 1. In order to make the label allocation more efficient, PE-2 has been allocated a ve-id value of 2. For completeness, the pseudowire template is also shown.

```
A:PE-2# configure service
      pw-template 1 create
        split-horizon-group "vpls-shg"
        exit
      exit
    vpls 1 customer 1 create
      bgp
        route-distinguisher 65536:1
        route-target export target:65536:1 import target:65536:1
        pw-template-binding 1
        exit
      exit
    bgp-vpls
      max-ve-id 10
      ve-name PE-2
        ve-id 2
      exit
      no shutdown
    exit
  stp
    shutdown
  exit
  service-name "VPLS service-174 PE-2 (192.0.2.2)"
  sap 1/1/4:1.0 create
  exit
  no shutdown
exit
```

Note that the **max-ve-id** value is set to 10 to allow an increase in the number of PEs that could be a part of this VPLS instance.

## BGP VPLS PE Configuration

### PE-3 Service Creation

Create VPLS instance on PE-3, using a ve-id value of 3.

```
A:PE-3# configure service
pw-template 1 create
  split-horizon-group "vpls-shg"
  exit
exit
vpls 1 customer 1 create
  bgp
    route-distinguisher 65536:1
    route-target export target:65536:1 import target:65536:1
    pw-template-binding 1
    exit
  exit
  bgp-vpls
    max-ve-id 10
    ve-name PE-3
    ve-id 3
    exit
    no shutdown
  exit
  stp
    shutdown
  exit
  sap 1/1/4:1.0 create
  exit
  no shutdown
exit
```

## PE-1 Service Operation Verification

Verify that the BGP-VPLS site is enabled on PE-1

```
*A:PE-1# show service id 1 bgp-vpls
=====
BGP VPLS Information
=====
Max Ve Id       : 10           Admin State    : Enabled
VE Name         : PE-1         VE Id          : 1
PW Tmpl used    : 1
=====
*A:PE-1#
```

Verify that the service is operationally up on PE-1

```
*A:PE-1# show service id 1 base
=====
Service Basic Information
=====
Service Id      : 1           Vpn Id        : 0
Service Type    : VPLS
Name            : VPLS service-174 PE-1 (192.0.2.1)
Description     : (Not Specified)
Customer Id     : 1           Creation Origin : manual
Last Status Change: 06/03/2014 14:36:59
Last Mgmt Change : 06/03/2014 14:58:26
Admin State     : Up           Oper State     : Up
MTU             : 1514         Def. Mesh VC Id : 1
SAP Count       : 1           SDP Bind Count : 2
Snd Flush on Fail : Disabled   Host Conn Verify : Disabled
Propagate MacFlush: Disabled   Per Svc Hashing  : Disabled
Allow IP Intf Bind: Disabled
Def. Gateway IP : None
Def. Gateway MAC : None
Temp Flood Time : Disabled     Temp Flood      : Inactive
Temp Flood Chg Cnt: 0

-----
Service Access & Destination Points
-----
Identifier                                     Type           AdmMTU  OprMTU  Adm  Opr
-----
sap:1/1/4:1.0                                 qinq           1522    1522    Up   Up
sdp:17401:4294967274 SB(192.0.2.2)           BgpVpls       0        1556    Up   Up
sdp:17404:4294967286 SB(192.0.2.3)           BgpVpls       0        1556    Up   Up
=====
*A:PE-1#
```

The SAP and SDPs are all operationally up. Note that the SB flags signify Spoke and BGP.

Further verification can be seen below where the ingress labels for PE-2 and PE-3, the labels allocated by PE-1, can be seen.

## BGP VPLS PE Configuration

```
*A:PE-1# show service id 1 sdp
=====
Services: Service Destination Points
=====
SdpId          Type Far End addr   Adm   Opr   I.Lbl   E.Lbl
-----
17401:4294967274 Bgp* 192.0.2.2      Up    Up    131064  131063
17404:4294967286 Bgp* 192.0.2.3      Up    Up    131065  131063
-----
Number of SDPs : 2
-----
* indicates that the corresponding row element may have been truncated.
*A:PE-1#
```

As can be seen from the following output, a BGP-VPLS NLRI update is sent to the route reflector (192.0.2.20) and is received by each PE.

The following debug trace from PE-1 shows the BGP NLRI update for VPLS 1 sent by PE-1 to the route reflector.

```
DEBUG #2001 Base Peer 1: 192.0.2.20
"Peer 1: 192.0.2.20: UPDATE
Peer 1: 192.0.2.20 - 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:65536:1
    12-vpn:Encap=19: Flags=none: MTU=1514
  Flag: 0x90 Type: 14 Len: 28 Multiprotocol Reachable NLRI:
    Address Family L2VPN
    NextHop len 4 NextHop 192.0.2.1
    [VPLS] veid: 1, vbo: 1, vbs: 8, label-base: 131063, RD 65536:1
```

Note the presence of the control flags within the extended community which indicates the status of the VPLS instance.

New control flags have been introduced to allow support for BGP multi-homing, D indicates that all attachment circuits are Down, or the VPLS is shutdown. The flags are used in BGP Multi-homing when determining which PEs are designated forwarders.

When flags=none, then all attachment circuits are up. In the example above no flags are present, but should all SAPs become operationally down, then the following debug would be seen:

```
DEBUG #2001 Base Peer 1: 192.0.2.20
"Peer 1: 192.0.2.20: UPDATE
Peer 1: 192.0.2.20 - 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:65536:1
    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: 1, vbo: 1, vbs: 8, label-base: 131063, RD 65536:1

```

The BGP VPLS signaling parameters are also present, namely the ve-id of the PE within the VPLS instance, the VBO and VBS, and the label base. The target indicates the VPLS instance, which must be matched against the import route targets of the receiving PEs.

The signaling parameters can be seen within the BGP update with following command:

```

*A:PE-1# show router bgp routes l2-vpn rd 65536:1 hunt
.. snipped ..
-----
RIB Out Entries
-----
Route Type      : VPLS
Route Dist.    : 65536:1
VeId           : 1
Block Size     : 8
Base Offset    : 1
Label Base     : 131063
NextHop        : 192.0.2.1
To             : 192.0.2.20
Res. NextHop   : n/a
Local Pref.    : 100
Interface Name : NotAvailable
Aggregator AS : None
Aggregator     : None
MED            : 0
Atomic Aggr.  : Not Atomic
AIGP Metric    : None
Connector      : None
Community      : target:65536:1
                l2-vpn/vrf-imp:Encap=19: Flags=none: MTU=1514: PREF=0
Cluster        : No Cluster Members
Originator Id  : None
Peer Router Id : 192.0.2.20
Origin         : IGP
AS-Path        : No As-Path
Neighbor-AS    : N/A
-----
Routes : 4
=====
*A:PE-1#

```

In this configuration example, PE-1 (192.0.2.1) with ve-id =1 has sent an update with base offset (VBO) =1, block size (VBS) = 8 and label base 131063. This means that labels 131063 (LB) to 131070 (LB+VBS-1) are available as de-multiplexer labels, egress labels to be used to reach PE-1 for VPLS 1.

## BGP VPLS PE Configuration

PE-2 receives this update from PE-1. This is seen as a valid VPLS BGP route from PE-1 through the route reflector with nexthop 192.0.2.1.

```
*A:PE-2# show router bgp routes l2-vpn rd 65536:1
=====
BGP Router ID:192.0.2.2      AS:65536      Local AS:65536
=====
Legend -
Status codes : u - used, s - suppressed, h - history, d - decayed, * - valid
Origin codes : i - IGP, e - EGP, ? - incomplete, > - best, b - backup
=====
BGP L2VPN Routes
=====
Flag  RouteType      Prefix      MED
RD    SiteId
Nexthop  VeId      BlockSize  LocalPref
As-Path  BaseOffset vplsLabelBase
-----
u*>i  VPLS          -           -           0
      65536:1     -           -           -
      192.0.2.1   1           8           100
      No As-Path  1           131063
i     VPLS          -           -           0
      65536:1     -           -           -
      192.0.2.2   2           8           100
      No As-Path  1           131063
u*>i  VPLS          -           -           0
      65536:1     -           -           -
      192.0.2.3   3           8           100
      No As-Path  1           131063
-----
Routes : 3
=====
*A:PE-2#
```

PE-2 uses this information in conjunction with its own ve-id to calculate the egress label towards PE-1, using the condition  $VBO \leq ve-id < (VBO+VBS)$ .

The ve-id of PE-2 is in the Label Block covered by  $VBO=1$ , thus,

Label calculation = label base + local ve-id - Base offset  
 $= 131063 + 2 - 1$

Egress label used = 131064

This is verified using the following command on PE-2 where the egress label toward PE-1 (192.0.2.1) is 131064.

```
*A:PE-2# show service id 1 sd
=====
Services: Service Destination Points
=====
SdpId      Type  Far End addr  Adm  Opr  I.Lbl  E.Lbl
-----
17406:4294967287 Bgp* 192.0.2.3    Up   Up   131065  131064
```

```
17407:4294967294 Bgp* 192.0.2.1      Up      Up      131063    131064
```

```
-----
Number of SDPs : 2
-----
```

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

```
*A:PE-2#
```

PE-3 also receives this update from PE-1 by the route reflector. This is seen as a valid VPLS BGP route from PE-1 with nexthop 192.0.2.1.

```
A:PE-3# show router bgp routes l2-vpn rd 65536:1
```

```
=====
BGP Router ID:192.0.2.3      AS:65536      Local AS:65536
=====
```

```
Legend -
Status codes : u - used, s - suppressed, h - history, d - decayed, * - valid
Origin codes : i - IGP, e - EGP, ? - incomplete, > - best, b - backup
```

```
=====
BGP L2VPN Routes
=====
```

Flag	RouteType	Prefix	MED
	RD	SiteId	Label
	Nexthop	VeId	LocalPref
	As-Path	BaseOffset	vplsLabelBase
u*>i	VPLS	-	0
	65536:1	-	-
	192.0.2.1	1	8
	No As-Path	1	131063
u*>i	VPLS	-	0
	65536:1	-	-
	192.0.2.2	2	8
	No As-Path	1	131063
i	VPLS	-	0
	65536:1	-	-
	192.0.2.3	3	8
	No As-Path	1	131063

```
-----
Routes : 3
=====
```

```
A:PE-3#
```

The ve-id of PE-3 is also in the label block covered by block offset VBO =1.

Label calculation = label base + local ve-id - VBO  
 = 131063 + 3 - 1

Egress label used = 131065

This is verified using the following command on PE-3 where egress label towards 192.0.2.1 is 131065.

## BGP VPLS PE Configuration

```
*A:PE-3# show service id 1 sdp
=====
Services: Service Destination Points
=====
SdpId           Type Far End addr   Adm   Opr       I.Lbl   E.Lbl
-----
17399:4294967268 Bgp* 192.0.2.2      Up    Up        131064  131065
17403:4294967281 Bgp* 192.0.2.1      Up    Up        131063  131065
-----
Number of SDPs : 2
-----
* indicates that the corresponding row element may have been truncated.
A:PE-3#
```



## PE-2 Service Operation Verification

For completeness, verify the service is operationally up on PE-2.

```
*A:PE-2# show service id 1 base
=====
Service Basic Information
=====
Service Id       : 1                Vpn Id           : 0
Service Type    : VPLS
Name            : VPLS service-174 PE-2 (192.0.2.2)
Description     : (Not Specified)
Customer Id     : 1                Creation Origin  : manual
Last Status Change: 06/03/2014 14:57:28
Last Mgmt Change  : 06/03/2014 14:58:20
Admin State     : Up                Oper State       : Up
MTU             : 1514              Def. Mesh VC Id  : 1
SAP Count       : 1                SDP Bind Count   : 2
Snd Flush on Fail : Disabled        Host Conn Verify : Disabled
Propagate MacFlush: Disabled        Per Svc Hashing  : Disabled
Allow IP Intf Bind: Disabled
Def. Gateway IP  : None
Def. Gateway MAC : None
Temp Flood Time  : Disabled        Temp Flood       : Inactive
Temp Flood Chg Cnt: 0
-----
Service Access & Destination Points
-----
Identifier                               Type           AdmMTU  OprMTU  Adm  Opr
-----
sap:1/1/4:1.0                             qinq           1522    1522    Up   Up
sdp:17406:4294967287 SB(192.0.2.3)         BgpVpls        0       1556    Up   Up
sdp:17407:4294967294 SB(192.0.2.1)         BgpVpls        0       1556    Up   Up
=====
*A:PE-2#
```

### PE-2 De-Multiplexer Label Calculation

In the same way that PE-1 allocates a label base (LB), block size (VBS), and base offset (VBO), PE-2 also allocates the same parameters for PE-1 and PE-3 to calculate the egress service label required to reach PE-2.

```
*A:PE-2# show router bgp routes 12-vpn rd 65536:1 hunt
.. snipped ..
-----
RIB Out Entries
-----
Route Type      : VPLS
Route Dist.     : 65536:1
VeId            : 2                               Block Size      : 8
Base Offset     : 1                               Label Base      : 131063
Nexthop        : 192.0.2.2
To              : 192.0.2.20
Res. Nexthop    : n/a
Local Pref.     : 100                             Interface Name  : NotAvailable
Aggregator AS  : None                             Aggregator     : None
Atomic Aggr.   : Not Atomic                       MED            : 0
AIGP Metric    : None
Connector       : None
Community       : target:65536:1
                  12-vpn/vrf-imp:Encap=19: Flags=none: MTU=1514: PREF=0
Cluster        : No Cluster Members
Originator Id  : None                             Peer Router Id  : 192.0.2.20
Origin         : IGP
AS-Path        : No As-Path
Neighbor-AS    : N/A
-----
Routes : 4
=====
*A:PE-2#
```

This is verified using the following command on PE-1 to show the egress label towards PE-2 (192.0.2.2) where the egress label towards PE-2 = 131063 + 1 - 1 = 131063.

```
*A:PE-1# show service id 1 sdp
=====
Services: Service Destination Points
=====
SdpId          Type Far End addr   Adm   Opr       I.Lbl   E.Lbl
-----
17401:4294967274 Bgp* 192.0.2.2     Up    Up        131064  131063
17404:4294967286 Bgp* 192.0.2.3     Up    Up        131065  131063
-----
Number of SDPs : 2
=====
* indicates that the corresponding row element may have been truncated.
*A:PE-1#
```

This is also verified using the following command on PE-3 to show the egress label towards PE-2 (192.0.2.2) where the egress label towards PE - 2 = 131063 + 3 - 1 = 131065.

```
A:PE-3# show service id 1 sdp
=====
Services: Service Destination Points
=====
SdpId           Type Far End addr   Adm   Opr   I.Lbl   E.Lbl
-----
17399:4294967268 Bgp* 192.0.2.2       Up    Up    131064  131065
17403:4294967281 Bgp* 192.0.2.1       Up    Up    131063  131065
-----
Number of SDPs : 2
=====
* indicates that the corresponding row element may have been truncated.
A:PE-3#
```

## PE-3 Service Operation Verification

Verify service is operationally up on PE-3:

```
A:PE-3# show service id 1 base
=====
Service Basic Information
=====
Service Id       : 1                Vpn Id           : 0
Service Type    : VPLS
Name            : VPLS service-174 PE-3 (192.0.2.3)
Description     : (Not Specified)
Customer Id     : 1                Creation Origin  : manual
Last Status Change: 06/03/2014 14:06:18
Last Mgmt Change  : 06/03/2014 14:58:30
Admin State     : Up
MTU             : 1514
SAP Count       : 1
Snd Flush on Fail : Disabled
Propagate MacFlush: Disabled
Allow IP Intf Bind: Disabled
Def. Gateway IP  : None
Def. Gateway MAC : None
Temp Flood Time  : Disabled
Temp Flood Chg Cnt: 0
Oper State      : Up
Def. Mesh VC Id : 1
SDP Bind Count  : 2
Host Conn Verify : Disabled
Per Svc Hashing : Disabled
Temp Flood      : Inactive

-----
Service Access & Destination Points
-----
Identifier           Type           AdmMTU  OprMTU  Adm  Opr
-----
sap:1/1/4:1.0       qinq          1522    1522    Up   Up
sdp:17399:4294967268 SB(192.0.2.2) BgpVpls    0      1556    Up   Up
sdp:17403:4294967281 SB(192.0.2.1) BgpVpls    0      1556    Up   Up
=====
A:PE-3#
```

## BGP VPLS PE Configuration

```
A:PE-3# show service id 1 sdp
=====
Services: Service Destination Points
=====
SdpId          Type Far End addr   Adm   Opr    I.Lbl   E.Lbl
-----
17399:4294967268 Bgp* 192.0.2.2      Up    Up     131064  131065
17403:4294967281 Bgp* 192.0.2.1      Up    Up     131063  131065
-----
Number of SDPs : 2
-----
* indicates that the corresponding row element may have been truncated.
A:PE-3#
```

## PE-3 De-Multiplexer Label Verification

PE-3 also allocates the required parameters for PE-1 and PE-2 to calculate the egress service label required to reach PE-3.

This is verified using the following command on PE-1 to show the egress label towards PE-3 (192.0.2.3) where egress label towards PE-2 = 131063.

```
*A:PE-1# show service id 1 sdp
=====
Services: Service Destination Points
=====
SdpId          Type Far End addr   Adm   Opr    I.Lbl   E.Lbl
-----
17401:4294967274 Bgp* 192.0.2.2      Up    Up     131064  131063
17404:4294967286 Bgp* 192.0.2.3      Up    Up     131065  131063
-----
Number of SDPs : 2
-----
* indicates that the corresponding row element may have been truncated.
*A:PE-1#
```

This is also verified using the following command on PE-2 to show the egress label towards PE-3 (192.0.2.3) which is using auto-provisioned SDP 17406.

```
*A:PE-2# show service id 1 sdp
=====
Services: Service Destination Points
=====
SdpId          Type Far End addr   Adm   Opr    I.Lbl   E.Lbl
-----
17406:4294967287 Bgp* 192.0.2.3      Up    Up     131065  131064
17407:4294967294 Bgp* 192.0.2.1      Up    Up     131063  131064
-----
Number of SDPs : 2
-----
```

```
=====
* indicates that the corresponding row element may have been truncated.
*A:PE-2#
```


This example has shown that for VPLS instance with 3 PEs, not all labels allocated by a PE will be used by remote PEs as de-multiplexor service labels. There will be some wastage of label space, so there is a necessity to choose ve-ids that keep this waste to a minimum.

The next example will show an even more wasteful use of labels by using random choice of ve-ids.

## BGP VPLS Using Pre-Provisioned SDP

It is possible to configure BGP-VPLS instances that use RSVP-TE transport tunnels. In this case, the SDP must be created with the MPLS LSPs mapped and with signaling set to BGP, as the service labels are signaled using BGP. The pseudowire template configured within the BGP-VPLS instance must use the **use-provisioned-sdp** keyword.

This example also examines the effect of using ve-ids that are not all within the same contiguous block.



**Figure 38: BGP VPLS Using Pre-Provisioned SDP**

Figure 38 shows an example of a VPLS instance where SDPs are pre-provisioned with LDP signalled transport tunnels.

### SDPs on PE-1

```
configure service
  sdp 33 mpls create
    description "from-192.0.2.1-id-33"
    far-end 192.0.2.2
    lsp "LSP-PE-1-PE-2"
    signaling bgp
    keep-alive
      shutdown
    exit
    no shutdown
  exit
  sdp 42 mpls create
    description "from-192.0.2.1-id-42"
    far-end 192.0.2.3
```

```

        lsp "LSP-PE-1-PE-3"
        signaling bgp
        keep-alive
            shutdown
        exit
        no shutdown
    exit
exit

```

### SDPs on PE-2

```

configure service
    sdp 40 mpls create
        far-end 192.0.2.1
        lsp "LSP-PE-2-PE-1"
        signaling bgp
        keep-alive
            shutdown
        exit
        no shutdown
    exit
    sdp 41 mpls create
        far-end 192.0.2.3
        lsp "LSP-PE-2-PE-3"
        signaling bgp
        keep-alive
            shutdown
        exit
        no shutdown
    exit
exit

```

### SDPs on PE-3

```

configure service
    sdp 38 mpls create
        description "from-192.0.2.3-id-38"
        far-end 192.0.2.1
        lsp "LSP-PE-3-PE-1"
        signaling bgp
        keep-alive
            shutdown
        exit
        collect-stats
        no shutdown
    exit
    sdp 39 mpls create
        description "from-192.0.2.3-id-39"
        far-end 192.0.2.2
        lsp "LSP-PE-3-PE-2"
        signaling bgp
        keep-alive
            shutdown
        exit
        collect-stats
        no shutdown

```

## BGP VPLS PE Configuration

```
        exit
exit
```

Note that pre-provisioned BGP-SDPs can also be used with BGP-VPLS. For reference, they are configured as follows:

```
*A:PE-3>config service sdp 64 mpls create
    far-end 192.0.2.2
    signaling bgp
    keep-alive
        shutdown
    exit
    no shutdown
```

To create an SDP within a service that uses the RSVP transport tunnel, a pseudowire template is required that has the **use-provisioned-sdp** parameter set.

Once again, a split horizon group is included to prevent forwarding between pseudowires.

The pseudowire template must be provisioned on all PEs and looks like:

```
A:PE-1# configure service
    pw-template 3 use-provisioned-sdp create
        split-horizon-group "vpls-shg"
    exit
exit
```

The following output shows the configuration required for a BGP-VPLS service using a pseudowire template configured for using pre-provisioned RSVP-TE SDPs.

```
*A:PE-1# configure service vpls 2 customer 1 create
    bgp
        route-distinguisher 65536:2
        route-target export target:65536:2 import target:65536:2
        pw-template-binding 3
    exit
exit
    bgp-vpls
        max-ve-id 100
        ve-name PE-1
            ve-id 1
        exit
        no shutdown
    exit
    stp
        shutdown
    exit
    sap 1/1/4:2.0 create
    exit
    no shutdown
exit
```



Note that the route distinguisher and route target extended community values for VPLS 2 are different from that in VPLS 1. The ve-id value for PE-1 can be the same as that in VPLS 1, but these must be different within the same VPLS instance on the other PEs — PE-2 should not have ve-id = 1.

Similarly, on PE-2 the configuration example shows where the ve-id value is 20:

```
A:PE-2# configure service vpls 2
      bgp
        route-distinguisher 65536:2
        route-target export target:65536:2 import target:65536:2
        pw-template-binding 3
      exit
    exit
  bgp-vpls
    max-ve-id 100
    ve-name PE-2
    ve-id 20
  exit
  no shutdown
exit
stp
  shutdown
exit
sap 1/1/4:2.0 create
exit
no shutdown
```

and on PE-3:

```
A:PE-3# configure service vpls 2
      bgp
        route-distinguisher 65536:2
        route-target export target:65536:2 import target:65536:2
        pw-template-binding 3
      exit
    exit
  bgp-vpls
    max-ve-id 100
    ve-name PE-3
    ve-id 3
  exit
  no shutdown
exit
stp
  shutdown
exit
sap 1/1/4:2.0 create
exit
no shutdown
```

Verify that the service is operationally up on PE-1.

```
*A:PE-2# show service id 2 base
```

```
=====
```

## BGP VPLS PE Configuration

```
Service Basic Information
=====
Service Id       : 2                Vpn Id           : 0
Service Type    : VPLS
Name            : (Not Specified)
Description     : (Not Specified)
Customer Id     : 1                Creation Origin  : manual
Last Status Change: 06/03/2014 14:57:28
Last Mgmt Change : 06/03/2014 14:58:52
Admin State     : Up                Oper State       : Up
MTU             : 1514              Def. Mesh VC Id : 2
SAP Count       : 1                SDP Bind Count  : 2
Snd Flush on Fail : Disabled        Host Conn Verify : Disabled
Propagate MacFlush: Disabled        Per Svc Hashing  : Disabled
Allow IP Intf Bind: Disabled
Def. Gateway IP : None
Def. Gateway MAC : None
Temp Flood Time : Disabled          Temp Flood       : Inactive
Temp Flood Chg Cnt: 0

-----
Service Access & Destination Points
-----
Identifier                                     Type           AdmMTU  OprMTU  Adm  Opr
-----
sap:1/1/4:2.0                                 qinq           1522    1522    Up   Up
sdp:40:4294967285 S(192.0.2.1)                BgpVpls        0       1556    Up   Up
sdp:41:4294967284 S(192.0.2.3)                BgpVpls        0       1556    Up   Up
=====
*A:PE-2#
```

Note that the SDP-ids are the pre-provisioned SDPs.

For completeness, verify the service is operationally up on PE-2

```
*A:PE-2# show service id 2 base
=====
Service Basic Information
=====
Service Id       : 2                Vpn Id           : 0
Service Type    : VPLS
Name            : (Not Specified)
Description     : (Not Specified)
Customer Id     : 1                Creation Origin  : manual
Last Status Change: 06/03/2014 14:57:28
Last Mgmt Change : 06/03/2014 14:58:52
Admin State     : Up                Oper State       : Up
MTU             : 1514              Def. Mesh VC Id : 2
SAP Count       : 1                SDP Bind Count  : 2
Snd Flush on Fail : Disabled        Host Conn Verify : Disabled
Propagate MacFlush: Disabled        Per Svc Hashing  : Disabled
Allow IP Intf Bind: Disabled
Def. Gateway IP : None
Def. Gateway MAC : None
Temp Flood Time : Disabled          Temp Flood       : Inactive
Temp Flood Chg Cnt: 0

-----
```

## Service Access &amp; Destination Points

```

-----
Identifier                               Type           AdmMTU  OprMTU  Adm  Opr
-----
sap:1/1/4:2.0                           qinq           1522    1522    Up   Up
sdp:40:4294967285 S(192.0.2.1)          BgpVpls       0        1556    Up   Up
sdp:41:4294967284 S(192.0.2.3)          BgpVpls       0        1556    Up   Up
=====

```

\*A:PE-2#

## Verify service is operational on PE-3:

A:PE-3# show service id 2 base

## Service Basic Information

```

-----
Service Id           : 2                Vpn Id           : 0
Service Type        : VPLS
Name                 : (Not Specified)
Description          : (Not Specified)
Customer Id         : 1                Creation Origin   : manual
Last Status Change: 06/03/2014 14:06:18
Last Mgmt Change   : 06/03/2014 14:58:57
Admin State         : Up                Oper State        : Up
MTU                  : 1514             Def. Mesh VC Id  : 2
SAP Count           : 1                SDP Bind Count   : 2
Snd Flush on Fail   : Disabled         Host Conn Verify  : Disabled
Propagate MacFlush : Disabled         Per Svc Hashing  : Disabled
Allow IP Intf Bind : Disabled
Def. Gateway IP     : None
Def. Gateway MAC    : None
Temp Flood Time     : Disabled         Temp Flood        : Inactive
Temp Flood Chg Cnt : 0
-----

```

## Service Access &amp; Destination Points

```

-----
Identifier                               Type           AdmMTU  OprMTU  Adm  Opr
-----
sap:1/1/4:2.0                           qinq           1522    1522    Up   Up
sdp:38:4294967282 S(192.0.2.1)          BgpVpls       0        1556    Up   Up
sdp:39:4294967266 S(192.0.2.2)          BgpVpls       0        1556    Up   Up
=====

```

A:PE-3#

## PE-1 De-Multiplexor Label Calculation

In the case of VPLS 1, all ve-ids are in the range of a single label block. In the case of VPLS 2, the ve-ids are in different blocks, for example, the ve-id 20 is in a different block to ve-ids 1 and 2.

As the label allocation is block-dependent, multiple labels blocks must be advertised by each PE to encompass this.

Consider PE-1's BGP update NLRIs.

```
*A:PE-1# show router bgp routes l2-vpn rd 65536:2 hunt
=====
BGP Router ID:192.0.2.1      AS:65536      Local AS:65536
=====
Legend -
Status codes : u - used, s - suppressed, h - history, d - decayed, * - valid
Origin codes : i - IGP, e - EGP, ? - incomplete, > - best, b - backup
=====
BGP L2VPN Routes
=====
.. snipped ..
-----
RIB Out Entries
-----
Route Type      : VPLS
Route Dist.     : 65536:2
VeId            : 1                               Block Size      : 8
Base Offset     : 1                               Label Base      : 131055
Nexthop        : 192.0.2.1
To              : 192.0.2.20
Res. Nexthop    : n/a
Local Pref.     : 100                             Interface Name  : NotAvailable
Aggregator AS  : None                             Aggregator     : None
Atomic Aggr.   : Not Atomic                       MED            : 0
AIGP Metric     : None
Connector      : None
Community       : target:65536:2
                  12-vpn/vrf-imp:Encap=19: Flags=none: MTU=1514: PREF=0
Cluster        : No Cluster Members
Originator Id  : None                             Peer Router Id  : 192.0.2.20
Origin         : IGP
AS-Path        : No As-Path
Neighbor-AS    : N/A

Route Type      : VPLS
Route Dist.     : 65536:2
VeId            : 1                               Block Size      : 8
Base Offset     : 17                              Label Base      : 131028
Nexthop        : 192.0.2.1
To              : 192.0.2.20
Res. Nexthop    : n/a
Local Pref.     : 100                             Interface Name  : NotAvailable
Aggregator AS  : None                             Aggregator     : None
Atomic Aggr.   : Not Atomic                       MED            : 0
AIGP Metric     : None
Connector      : None
Community       : target:65536:2
```

```

l2-vpn/vrf-imp:Encap=19: Flags=none: MTU=1514: PREF=0
Cluster      : No Cluster Members
Originator Id : None                Peer Router Id : 192.0.2.20
Origin       : IGP
AS-Path      : No As-Path
Neighbor-AS  : N/A

```

```

-----
Routes : 8
=====

```

```
*A:PE-1#
```

Two NLRIs updates are sent to the route reflector, with the following label parameters

1. LB = 131055, VBS = 8, VBO = 1
2. LB = 131028, VBS = 8, VBO = 17

PE-2 has a ve-id of 20. Applying the condition  $VBO \leq \text{ve-id} < (VBO+VBS)$

```

Update 1: LB = 131055, VBS = 8, VBO = 1
VBO ≤ ve-id for ve-id = 20 is TRUE
ve-id < (VBO+VBS) for ve-id = 20 is FALSE.
PE-2 cannot choose a label from this block.

```

```

Update 2: LB = 131028, VBS = 8, VBO = 17
VBO ≤ ve-id for ve-id = 20 is TRUE
ve-id < (VBO+VBS) for ve-id = 20 is TRUE.
PE-2 chooses label 131028 + 20 - 17 = 131031 (LB + veid - VBO)

```

The egress label chosen is verified by examining the egress label towards PE-1 (192.0.2.1) on PE-2.

```
*A:PE-2# show service id 2 sdp
```

```

=====
Services: Service Destination Points
=====

```

SdpId	Type	Far End addr	Adm	Opr	I.Lbl	E.Lbl
40:4294967285	Bgp*	192.0.2.1	Up	Up	131020	131031
41:4294967284	Bgp*	192.0.2.3	Up	Up	131022	131031

```

-----
Number of SDPs : 2
=====

```

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

```
*A:PE-2#
```

PE-3 has a ve-id of 3. Applying the condition  $VBO \leq \text{ve-id} < (VBO+VBS)$

```

Update 1: LB = 131055, VBS = 8, VBO = 1
VBO ≤ ve-id for ve-id = 3 is TRUE

```

## BGP VPLS PE Configuration

ve-id < (VBO+VBS) for ve-id = 3 is TRUE.  
PE-3 chooses label 131055 + 3 - 1 = 131057 (LB + veid - VBO)

Update 2: LB = 131028, VBS = 8, VBO = 17  
VBO ≤ ve-id for ve-id = 3 is FALSE  
ve-id < (VBO+VBS) for ve-id = 3 is FALSE.  
PE-3 cannot choose a label from this block.

The egress label chosen is verified by examining the egress label towards PE-1 (192.0.2.1) on PE-3.

```
A:PE-3# show service id 2 sdp
=====
Services: Service Destination Points
=====
SdpId          Type Far End addr   Adm   Opr       I.Lbl   E.Lbl
-----
38:4294967282  Bgp* 192.0.2.1      Up    Up        131055  131057
39:4294967266  Bgp* 192.0.2.2      Up    Up        131031  131022
-----
Number of SDPs : 2
=====
* indicates that the corresponding row element may have been truncated.
A:PE-3#
```

To illustrate the allocation of label blocks by a PE, against the actual use of the same labels, consider the following. When BGP updates from each PE signal the multiplexer labels in blocks of eight, the allocated label values are added to the in-use pool. The label pool of PE-1 can be verified as per the following output which shows labels used along with the associated protocol:

```
*A:PE-1# show router mpls label 32 131071 in-use
=====
MPLS Labels from 32 to 131071 (In-use)
=====
Label          Label Type      Label Owner
-----
131024         dynamic        TLDP
131025         dynamic        TLDP
131026         dynamic        TLDP
131027         dynamic        TLDP
131028         dynamic        BGP
131029         dynamic        BGP
131030         dynamic        BGP
131031         dynamic        BGP
131032         dynamic        BGP
131033         dynamic        BGP
131034         dynamic        BGP
131035         dynamic        BGP
131036         dynamic        TLDP
131037         dynamic        TLDP
131038         dynamic        TLDP
131039         dynamic        RSVP
131040         dynamic        RSVP
131041         dynamic        TLDP
```

131042	dynamic	ILDP
131043	dynamic	ILDP
131044	dynamic	ILDP
131045	dynamic	ILDP
131046	dynamic	ILDP
131047	dynamic	BGP
131048	dynamic	BGP
131049	dynamic	BGP
131050	dynamic	BGP
131051	dynamic	BGP
131052	dynamic	BGP
131053	dynamic	BGP
131054	dynamic	BGP
131055	dynamic	BGP
131056	dynamic	BGP
131057	dynamic	BGP
131058	dynamic	BGP
131059	dynamic	BGP
131060	dynamic	BGP
131061	dynamic	BGP
131062	dynamic	BGP
131063	dynamic	BGP
131064	dynamic	BGP
131065	dynamic	BGP
131066	dynamic	BGP
131067	dynamic	BGP
131068	dynamic	BGP
131069	dynamic	BGP
131070	dynamic	BGP
131071	dynamic	ILDP

```
-----
In-use labels (Owner: All) in specified range : 48
In-use labels in entire range : 48
=====
```

```
*A:PE-1#
```

This shows that 32 labels have been allocated for use by BGP. Of this number, 16 labels have been allocated for use by PEs within VPLS 2 to communicate with PE-1, the blocks with label base 131028 and with label base 131055.

There are only two neighboring PEs within this VPLS instance, so only two labels will ever be used in the data plane for traffic destined to PE-1. These are 131031 and 131057. The remaining labels have no PE with the associated ve-id that can use them.

Once again, this case emphasizes that to reduce label wastage, contiguous ve-ids in the range (N..N+7) per VPLS should be chosen, where N>0.

## Conclusion

BGP-VPLS allows the delivery of Layer 2 VPN services to customers where BGP is commonly used. The examples presented in this section show the configuration of BGP-VPLS together with the associated show outputs which can be used for verification and troubleshooting.