

# BGP Virtual Private Wire Services

---

## In This Chapter

This section describes BGP Virtual Private Wire Service (VPWS) configurations.

Topics in this section include:

- [Applicability on page 250](#)
- [Overview on page 251](#)
- [Configuration on page 255](#)
- [Conclusion on page 280](#)

## Applicability

This example is applicable to all of the 7710, 7450, 7750 SR and 7950 XRS series and was tested on Release 12.0.R1. There are no prerequisites for this configuration.

---


## Introduction

There are currently two IETF standards for the provisioning of Virtual Private Wire Services (VPWS). RFC 4447, *Pseudowire Setup and Maintenance Using the Label Distribution Protocol (LDP)*, describes Label Distribution Protocol (LDP) VPWS, where VPWS pseudowires are signaled using LDP between Provider Edge (PE) Routers.

RFC 6624, *Layer 2 Virtual Private Networks Using BGP for Auto-Discovery and Signaling*, describes the use of Border Gateway Protocol (BGP) for signaling of pseudo-wires between such PEs.

The purpose of this example is to describe the configuration and troubleshooting for BGP VPWS.

## Overview



al\_0265

**Figure 39: Network Topology**

The network topology is displayed in [Figure 39](#). The setup uses five Service Router (SR) nodes located in the same Autonomous System (AS). There are three PE routers connected to a single P router and a Route Reflector (RR-1) for the AS. The Provider Edge routers are all BGP VPWS aware. The Provider (P) router is BGP VPWS unaware and also does not take part in the BGP process.

The following configuration tasks should be completed as a prerequisite:

- IS-IS or OSPF should be configured on each of the network interfaces between the PE/P routers and route reflector.
- MPLS should be configured on all interfaces between PE routers and P routers. It is not required between P-4 and RR-1.
- LDP should be configured on interfaces between PE and P routers. It is not required between P-4 and the RR-1.
- RSVP protocol should be configured on interfaces between PE and P routers. It is not required between P-4 and the RR-1.

## BGP VPWS

In this architecture, a VPWS is a collection of two (or three in case of redundancy) BGP VPWS service instances present on different PEs in a provider network.

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

Each PE with a BGP VPWS instance is identified by a VPWS Edge Identifier (VE-ID) and the presence of other BGP VPWS instances is determined using the exchange of standard BGP extended community route targets between PEs.

Each PE will advertise, via the route reflector, the presence of its BGP VPWS instance to all other PEs, along with a block of multiplexer labels (for BGP VPWS there is just one label per block) that can be used to communicate between each instance, plus a BGP next-hop that determines a labeled transport tunnel to be used between PEs.

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

The following examples show the configuration of four BGP VPWS scenarios.

- Single homed BGP VPWS
  - using auto-provisioned SDPs
  - using pre-provisioned SDPs
- Dual homed BGP VPWS
  - with single pseudowire
  - with active/standby pseudowire

## Configuration Tasks

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

The configuration for PE-1 is as follows:

```
configure router
  autonomous-system 65536
  bgp
    group internal
      family l2-vpn
      type internal
      peer-as 65536
      local-as 65536
      neighbor 192.0.2.5
      exit
    exit
  no shutdown
  exit
exit
```

The configuration for the other PE nodes is very similar. The IP addresses can be derived from [Figure 39](#).

The configuration for the Route Reflector (RR-1) is:

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

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

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

```
=====
```

## Configuration Tasks

```

BGP Router ID:192.0.2.5      AS:65536      Local AS:65536
=====
BGP Admin State      : Up          BGP Oper State      : Up
Total Peer Groups    : 1           Total Peers          : 3
Total BGP Paths      : 20          Total Path Memory    : 3840
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 : 16          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      58   0 00h21m05s 6/0/16 (L2VPN)
      79       0
192.0.2.2
      65536      50   0 00h20m06s 4/0/16 (L2VPN)
      63       0
192.0.2.3
      65536      60   0 00h21m33s 6/0/16 (L2VPN)
      84       0
-----
A:RR-1#

```

## Configuration

---

### Pseudowire Templates

BGP VPWS utilizes pseudowire (PW) templates 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, such as `vc-type`, `vlan-vc-tag`, `hash-label`, filters, etc. The following parameters are recognized by BGP VPWS; the remainder is ignored.

The following commands are supported parameters:

```

config
  service
    [no] pw-template policy-id [use-provisioned-sdp] [create]
      accounting-policy acct-policy-id
      no accounting-policy
      [no] collect-stats
    egress
      filter ipv6 ipv6-filter-id
      filter ip ip-filter-id
      filter mac mac-filter-id
      no filter [ip ip-filter-id] [mac mac-filter-id] [ipv6 ipv6-filter-id]
      qos network-policy-id port-redirect-group queue-group-name [instance instance-id]
      no qos
    [no] force-vlan-vc-forwarding
    hash-label [signal-capability]
    no hash-label
    ingress
      filter ipv6 ipv6-filter-id
      filter ip ip-filter-id
      filter mac mac-filter-id
      no filter [ip ip-filter-id] [mac mac-filter-id] [ipv6 ipv6-filter-id]
      qos network-policy-id fp-redirect-group queue-group-name instance instance-id
      no qos
    [no] sdp-exclude group-name
    [no] sdp-include group-name
    vc-type {ether | vlan}
    vlan-vc-tag 0..4094
    no vlan-vc-tag

```

Note that:

- The encapsulation type in the Layer-2 extended community is either 4 (Ethernet VLAN tagged mode) or 5 (Ethernet raw mode), depending on the `vc-type` parameter.
- The **force-vlan-vc-forwarding** function will add a tag (equivalent to `vc-type vlan`) and will allow for customer QoS transparency (dot1p+DE bits).

## Pseudowire Templates for Auto-SDP Creation using LDP

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

LDP-based SDPs can be automatically instantiated or pre-provisioned. RSVP-TE-based SDPs have to be pre-provisioned. If pre-provisioned pseudowires should be used, the pw-template must be created with the **use-provisioned-sdp** parameter.

```
*A:PE-1# configure service pw-template
- [no] pw-template <policy-id> [use-provisioned-sdp] [create]
```

---

## 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. For example, on PE-1:

```
A:PE-1# configure router ldp
      interface-parameters
        interface "PE-1-P-4"
      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. The example below is created using the default values:

```
A:PE-1# configure service
      pw-template 1 create
      exit
```

---

## Pseudowire Templates for Provisioned SDPs using RSVP-TE

RSVP-TE LSPs need to be created between the PE routers on which provisioned SDPs will be used as prerequisite.

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

```
A:PE-1# configure router mpls
      interface "system"
        no shutdown
      exit
      interface "PE-1-P-4"
        no shutdown
```



```

exit
path "dyn"
  no shutdown
exit
lsp "LSP-PE-1-PE-2"
  to 192.0.2.2
  primary "dyn"
  exit
  no shutdown
exit
lsp "LSP-PE-1-PE-3"
  to 192.0.2.3
  primary "dyn"
  exit
  no shutdown
exit
no shutdown
exit

```

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

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

The SDP creation on PE-1 towards PE-2 is shown below; similar SDPs are required on each PE to the remote PEs in the service where provisioned SDPs are to be used (only on PE-2 in the example below).

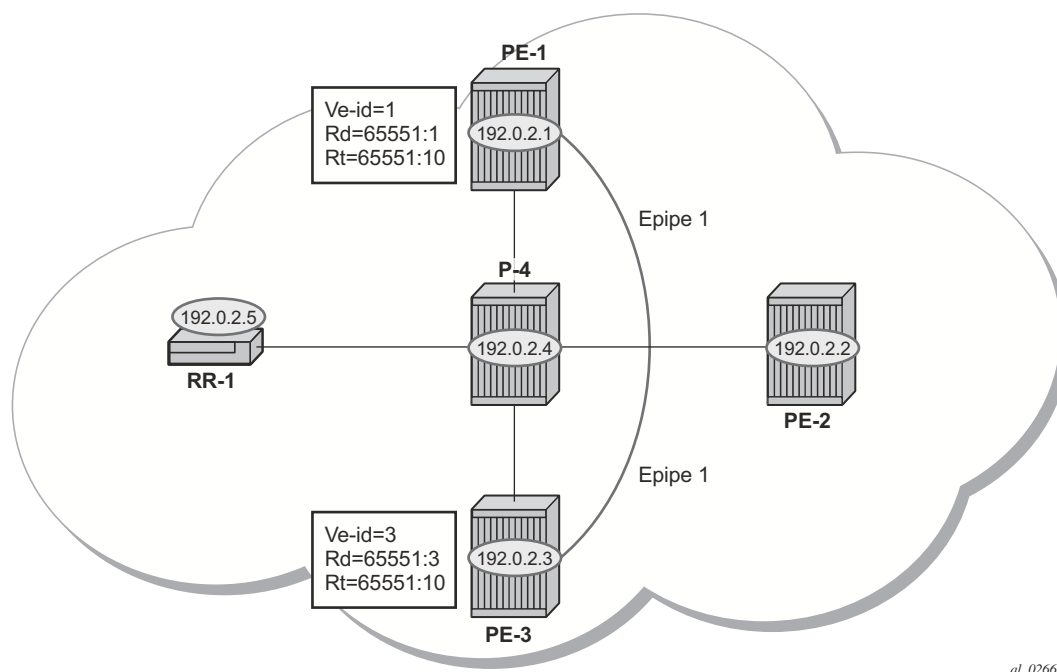
```

A:PE-1# configure service sdp 12 mpls create
  description "from-PE1"
  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. BGP VPWS instances using BGP VPWS signaling are able to use these SDPs. 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 VPWS.

## Single Homed BGP VPWS using Auto-Provisioned SDPs



al\_0266

**Figure 40: Single Homed BGP VPWS using Auto-Provisioned SDPs**

Figure 40 shows a schematic of a single homed BGP VPWS between PE-1 and PE-2 where SDPs are auto-provisioned. In this case, the transport tunnels are LDP signaled.

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

```
A:PE-1# configure service
pw-template 1 create
  vc-type vlan
exit
epipe 1 customer 1 create
  bgp
    route-distinguisher 65551:1
    route-target export target:65551:10 import target:65551:10
    pw-template-binding 1
  exit
exit
bgp-vpws
  ve-name "PE1"
  ve-id 1
exit
  remote-ve-name "PE3"
  ve-id 3
```

```

        exit
        no shutdown
    exit
    sap 1/1/4:100 create
    exit
    no shutdown
exit
exit

```

The **bgp** context specifies parameters that are required for BGP VPWS.

Within the **bgp** context, parameters are configured that are used by the neighboring PEs to determine the membership of a given BGP VPWS; in other words, the auto-discovery of PEs in the same BGP VPWS, the route-distinguisher is configured, along with the route target extended communities. Route target communities are used to determine membership of a given BGP VPWS. Note that the import and export route targets at the BGP level are mandatory. The pw-template binding is then applied and its parameters are used for both the routes sent by this PE and the received routes matching the route target value.

Within the **bgp-vpws** context, the signaling parameters are also configured. These determine the service labels required for the data plane of the VPWS instance.

The VPWS Edge ID (VE-ID) is a numerical value assigned to each PE within a BGP VPWS. This value must be unique for a given BGP VPWS, with the exception of multi-homed scenarios, where two dual-homed PEs can have the same VE-ID and are distinguishable by the site preference (or by the tie breaking rules from the multi-homing draft RFC).

It is also worth noting that changes to the pseudowire template are not taken into account once the pseudowire has been set up (changes of route-target are refreshed though). PW-templates can be re-evaluated with the **tools perform eval-pw-template** command. The **eval-pw-template** checks if all of the bindings using this pw-template policy are still meant to be used this policy. If the template has changed and allow-service-impact is TRUE, then the old binding is removed and it is re-added using the new template.

## VE-ID and BGP Label Allocations

For a point-to-point VPWS, there are only two members within the BGP VPWS service, so only one label entry is required by each remote service. For dual-homed scenarios, there are two labels for the redundant site, one from each dual-homed PE.

Each PE allocates a label per BGP VPWS instance for the remote PEs, so it signals blocks with one label. 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 1 and cannot be changed.
- A VE Base Offset (VBO) corresponding to the first label in the label block.

## PE-3 Service Creation

On PE-3 create a BGP VPWS service using pseudowire template 1. PE-3 has been allocated a VE-ID of 3. For completeness, the pw-template is also shown.

```
A:PE-3# configure service
      pw-template 1 create
          vc-type vlan
      exit
      epipe 1 customer 1 create
          bgp
              route-distinguisher 65551:3
              route-target export target:65551:10 import target:65551:10
              pw-template-binding 1
          exit
      exit
      bgp-vpws
          ve-name "PE3"
          ve-id 3
      exit
      remote-ve-name "PE1"
          ve-id 1
      exit
      no shutdown
  exit
  sap 1/1/4:101 create
  exit
  no shutdown
exit
exit
```

## PE-1 Service Operation Verification

Verify that the BGP VPWS service is enabled on PE-1.

```
*A:PE-1# show service id 1 bgp-vpws
=====
BGP VPWS Information
=====
Admin State       : Enabled
VE Name           : PE1                VE Id           : 1
PW Tmpl used      : 1

Remote-Ve Information
-----
Remote VE Name    : PE3                Remote VE Id    : 3
=====
*A:PE-1
```

Verify the BGP information used by the BGP VPWS service on PE-1.

```
*A:PE-1# show service id 1 bgp
=====
BGP Information
=====
Route Dist        : 65551:1
Rte-Target Import : 65551:10                Rte-Target Export: 65551:10

PW-Template Id    : 1
Import Rte-Tgt   : None
=====
*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     : Epipe
Name             : (Not Specified)
Description      : (Not Specified)
Customer Id      : 1                Creation Origin  : manual
Last Status Change: 03/10/2014 10:08:06
Last Mgmt Change : 03/10/2014 10:08:06
Admin State      : Up                Oper State       : Up
MTU              : 1514
Vc Switching     : False
SAP Count        : 1                SDP Bind Count   : 1
Per Svc Hashing  : Disabled
Force QTag Fwd   : Disabled
```

## PE-1 Service Operation Verification

```

-----
Service Access & Destination Points
-----
Identifier                               Type           AdmMTU  OprMTU  Adm  Opr
-----
sap:1/1/4:100                           q-tag         1578    1578    Up   Up
sdp:17407:4294967295 SB(192.0.2.3)      BgpVpws      0       1552    Up   Up
=====

```

\*A:PE-1#

The SAP and SDP are all operationally up. Note that the indication “SB” next to the SDP-id signify “Spoke” and “BGP”.

Further verification can be seen below where the ingress label for PE-3, that is, the labels used by PE-1 are shown.

```

*A:PE-1# show service id 1 sdp
=====
Services: Service Destination Points
=====
SdpId           Type Far End addr  Adm   Opr   I.Lbl   E.Lbl
-----
17407:4294967295 BVws 192.0.2.3    Up    Up    262141  262142
-----

```

Number of SDPs : 1

\*A:PE-1#

The following debug output from PE-1 shows the BGP VPWS NLRI update for Epipe 1 sent by PE-1 to the route reflector (192.0.2.5). This update will then be received by the other PEs.

```

170 2014/02/19 12:49:16.82 UTC MINOR: DEBUG #2001 Base Peer 1: 192.0.2.5
"Peer 1: 192.0.2.5: UPDATE
Peer 1: 192.0.2.5 - Send BGP UPDATE:
  Withdrawn Length = 0
  Total Path Attr Length = 76
  Flag: 0x90 Type: 14 Len: 32 Multiprotocol Reachable NLRI:
    Address Family L2VPN
    NextHop len 4 NextHop 192.0.2.1
    [VPLS/VPWS] veid: 1, vbo: 3, vbs: 1, label-base: 262131, RD 65551:1, csv
: 0x0
  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:65551:10
    12-vpn/vrf-imp:Encap=4: Flags=none: MTU=1514: PREF=0
"

```

Note the presence of the control flags within the extended community which indicate the status of the BGP VPWS instance.

The control flags are described below:

```

0 1 2 3 4 5 6 7
+-----+-----+
|D|A|F|Z|Z|Z|C|S| (Z = MUST Be Zero)
+-----+-----+

```

- D: access circuit down indicator. D is 1 if all access circuits are down, otherwise D is 0.
- A: automatic site id allocation, which is not supported. This is ignored on receipt and set to 0 on sending.
- F: MAC flush indicator, this relates to VPLS. This is set to 0 and ignored on receipt.
- C: presence of a control word. Control word usage is not supported. This is set to 0 on sending (control word not present) and if a non-zero value is received (indicating a control word is required) the pseudowire will not be created.
- S: sequenced delivery. Sequenced delivery is not supported. This is set to 0 on sending (no sequenced delivery) and if a non-zero value is received (indicating sequenced delivery required) the pseudowire will not be created.

The BGP VPWS NLRI is based on that defined for BGP VPLS but is extended with a circuit status vector. The circuit status vector is used to indicate the status of both the SAP and the spoke-SDP within the local service. As the VE block size used is 1, the most significant bit in the circuit status vector TLV value will be set to 1 if either the SAP or spoke-SDP is down; otherwise, it will be set to 0.

```

193 2014/02/19 13:22:10.61 UTC MINOR: DEBUG #2001 Base Peer 1: 192.0.2.5
"Peer 1: 192.0.2.5: UPDATE
Peer 1: 192.0.2.5 - Received BGP UPDATE:
  Withdrawn Length = 0
  Total Path Attr Length = 90
  Flag: 0x90 Type: 14 Len: 32 Multiprotocol Reachable NLRI:
    Address Family L2VPN
    NextHop len 4 NextHop 192.0.2.1
    [VPLS/VPWS] veid: 1, vbo: 3, vbs: 1, label-base: 262131, RD 65551:1, csv
: 0x80
  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: 0x80 Type: 9 Len: 4 Originator ID: 192.0.2.1
  Flag: 0x80 Type: 10 Len: 4 Cluster ID:
    1.1.1.1
  Flag: 0xc0 Type: 16 Len: 16 Extended Community:
    target:65551:10
    12-vpn/vrf-imp:Encap=4: Flags=D: MTU=1514: PREF=0
"

```

## PE-3 Service Operation Verification

After shutting down the local SAP, the CSV has the most-significant bit set to 1 (0x80).

The BGP VPWS update can be shown using the following command (note that the SDP-id has changed from that seen above as the spoke SDP has been re-signaled since the above output):

```
*A:PE-1# show service l2-route-table bgp-vpws detail
=====
Services: L2 Bgp-Vpws Route Information - Summary
=====

Svc Id       : 1
VeId        : 3
PW Temp Id   : 1
RD          : *65551:3
Next Hop     : 192.0.2.3
State (D-Bit) : up(0)
Path MTU     : 1514
Control Word : 0
Seq Delivery : 0
Status       : active
Tx Status    : active
CSV         : 0
Preference   : 0
Sdp Bind Id  : 17407:4294967295

<snipped>

*A:PE-1#
```

---


## PE-3 Service Operation Verification

Similar to PE-1, the service operation should be validated on PE-3.



## Single Homed BGP VPWS using Pre-Provisioned SDP

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



**Figure 41: Single Homed BGP VPWS using Pre-Provisioned SDP**

Figure 41 shows a schematic of a BGP VPWS where SDPs are pre-provisioned with RSVP-TE signaled transport tunnels.

### SDP on PE-1

```
configure service
  sdp 12 mpls create
    description "from-192.0.2.1"
    signaling bgp
    far-end 192.0.2.2
    lsp "LSP-PE-1-PE-2"
    keep-alive
      shutdown
    exit
  no shutdown
  exit
exit
```

## Single Homed BGP VPWS using Pre-Provisioned SDP

### SDP on PE-2

```
configure service
  sdp 21 mpls create
    description "from-192.0.2.2"
    signaling bgp
    far-end 192.0.2.1
    lsp "LSP-PE-2-PE-1"
    keep-alive
      shutdown
    exit
  no shutdown
  exit
exit
```

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

The pw-template is provisioned on both PEs as follows:

```
A:PE-1# configure service
  pw-template 2 use-provisioned-sdp create
  vc-type ether
  exit
```

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

```
A:PE-1# configure service epipe 2 customer 1 create
  bgp
    route-distinguisher 65551:2211
    route-target export target:65551:22 import target:65551:22
    pw-template-binding 2
  exit
  exit
  bgp-vpws
    ve-name "PE-1"
    ve-id 1
  exit
    remote-ve-name "PE-2"
    ve-id 2
  exit
  no shutdown
  exit
  sap 1/1/4:2200 create
  exit
  no shutdown
  exit
exit
```

The route distinguisher and route target extended community values for Epipe 2 are different from that in Epipe 1. This is to differentiate between the two as their visibility is global within the BGP

domain. The VE-ID values can be reused in each Epipe instance, as long as they are unique within the instance.

Similarly, on PE-2 the configuration is as seen below, where the VE-ID is 2:

```
A:PE-2# configure service epipe 2 customer 1 create
      bgp
        route-distinguisher 65551:2222
        route-target export target:65551:22 import target:65551:22
        pw-template-binding 2
        exit
      exit
      bgp-vpws
        ve-name "PE2"
        ve-id 2
        exit
        remote-ve-name "PE-1"
        ve-id 1
        exit
        no shutdown
      exit
      sap 1/1/4:200 create
      exit
      no shutdown
    exit
```

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

```
*A:PE-1# show service id 2 base
=====
Service Basic Information
=====
Service Id       : 2                Vpn Id          : 0
Service Type    : Epipe
Name            : (Not Specified)
Description     : (Not Specified)
Customer Id     : 1                Creation Origin  : manual
Last Status Change: 03/10/2014 10:09:04
Last Mgmt Change  : 03/10/2014 10:09:04
Admin State     : Up                Oper State      : Up
MTU             : 1514
Vc Switching   : False
SAP Count      : 1                SDP Bind Count  : 1
Per Svc Hashing : Disabled
Force QTag Fwd  : Disabled
-----
Service Access & Destination Points
-----
Identifier                                     Type           AdmMTU  OprMTU  Adm  Opr
-----
sap:1/1/4:2200                                q-tag         1578    1578    Up   Up
sdp:12:4294967292 S(192.0.2.2)                BgpVpws       0       1552    Up   Up
=====
*A:PE-1#
```

## Single Homed BGP VPWS using Pre-Provisioned SDP

Note that the SDP-id is the pre-provisioned SDP 12.

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    : Epipe
Name            : (Not Specified)
Description     : (Not Specified)
Customer Id     : 1                Creation Origin  : manual
Last Status Change: 03/10/2014 12:06:37
Last Mgmt Change  : 03/10/2014 12:06:37
Admin State     : Up              Oper State      : Up
MTU             : 1514
Vc Switching   : False
SAP Count      : 1                SDP Bind Count  : 1
Per Svc Hashing : Disabled
Force QTag Fwd : Disabled
-----
Service Access & Destination Points
-----
Identifier                               Type           AdmMTU  OprMTU  Adm  Opr
-----
sap:1/1/4:200                            q-tag         1578    1578    Up   Up
sdp:21:4294967293 S(192.0.2.1)          BgpVpws       0       1552    Up   Up
=====
A:PE-2#
```

Again, the SDP-id used is the pre-provisioned SDP 21.

## Dual Homed BGP VPWS with Single Pseudowire

For access redundancy, an Epipe using a BGP VPWS service can be configured as dual-homed, as described in *draft-ietf-l2vpn-vpls-multihoming-03*. It can be configured with a single pseudowire setup, where the redundant pseudowire is not created until the initially active pseudowire is removed.

The diagram below shows a setup where an Epipe is configured on each PE. Site B is dual-homed to PE-1 and PE-3 with the remote PE-2 connected to site A; each site connection uses a SAP. A single pseudowire using Ethernet Raw Mode encapsulation connects PE-2 to PE-1 or PE-3 (but not both at the same time). The pseudowire is signaled using BGP VPWS over a tunnel LSP between the PEs.

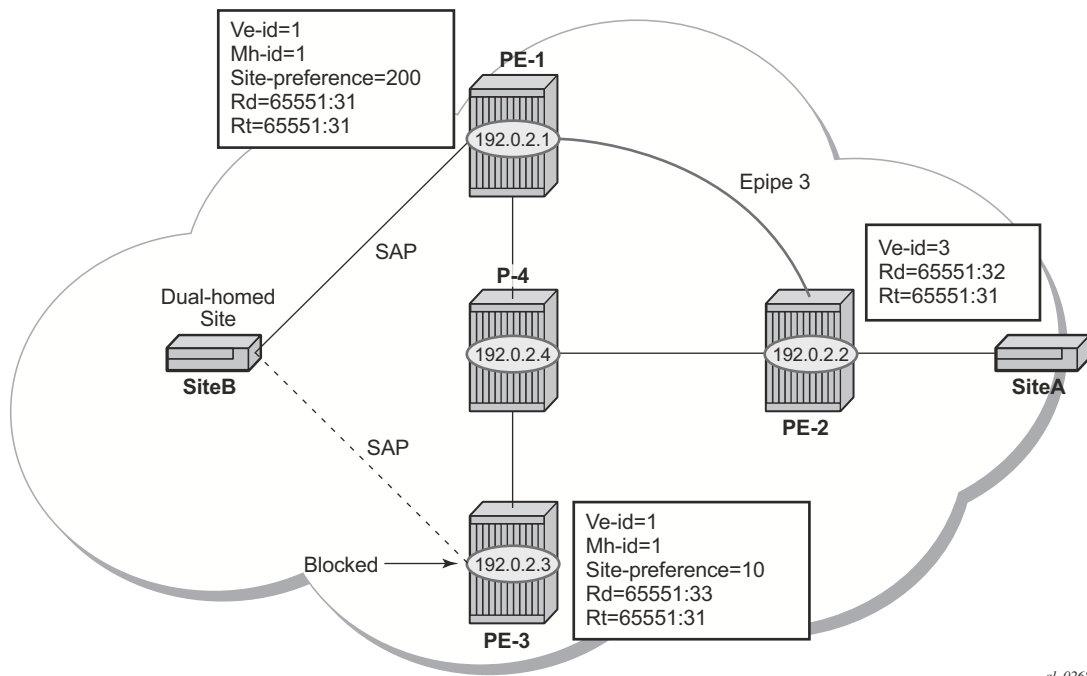


Figure 42: Dual Homed BGP VPWS with Single Pseudowire

BGP multi-homing is configured for the dual-homed site B using a site-id=1. The site-preference on PE-1 is set to 200 and to 10 on PE-3, this ensures that PE-1 will be the site's designated forwarder and the pseudowire from PE-2 will be created to PE-1 when PE-1 is fully operational (no pseudowire is created on PE-2 to PE-3). If PE-1 fails, or the multi-homing site fails over to PE-3, then the pseudowire from PE-2 to PE-1 will be removed and a new pseudowire will be created from PE-2 to PE-3.

## Dual Homed BGP VPWS with Single Pseudowire

### PE-1 Configuration

```
A:PE-1# configure service
      pw-template 3 create
      exit
      epipe 3 customer 1 create
        bgp
          route-distinguisher 65551:31
          route-target export target:65551:31 import target:65551:31
          pw-template-binding 3
          exit
        exit
      bgp-vpws
        ve-name "PE-1"
          ve-id 1
          exit
        remote-ve-name "PE-2"
          ve-id 2
          exit
        no shutdown
      exit
      site "siteB" create
        site-id 1
        sap 1/1/4:99
        site-preference 200
        no shutdown
      exit
      sap 1/1/4:99 create
      exit
      no shutdown
    exit
  exit
```

### PE-3 Configuration

```
A:PE-3# configure service
      pw-template 3 create
      exit
      epipe 3 customer 1 create
        bgp
          route-distinguisher 65551:33
          route-target export target:65551:31 import target:65551:31
          pw-template-binding 3
          exit
        exit
      bgp-vpws
        ve-name "PE-3"
          ve-id 1
          exit
        remote-ve-name "PE-2"
          ve-id 2
          exit
        no shutdown
      exit
      site "siteB" create
        site-id 1
        sap 1/1/4:99
```

```

        site-preference 10
        no shutdown
    exit
    sap 1/1/4:99 create
    exit
    no shutdown
exit
exit

```

Note that in the above configurations, the remote-ve-name for PE-2 uses VE-ID 2 on both PE-1 and PE-3.

### PE-2 Configuration

```

A:PE-2# configure service
pw-template 3 create
exit
epipe 3 customer 1 create
    bgp
        route-distinguisher 65551:32
        route-target export target:65551:31 import target:65551:31
        pw-template-binding 3
    exit
exit
bgp-vpws
    ve-name "PE-2"
        ve-id 2
    exit
    remote-ve-name "PE-1orPE-3"
        ve-id 1
    exit
    no shutdown
exit
sap 1/1/4:99 create
exit
no shutdown
exit
exit

```

On PE-2, the remote-ve-name is configured as PE-1 or PE-3; this is because both of these PEs are configured with VE-ID 1.

As a result of this configuration, observe that on PE-2, there are multiple route entries for Route-Target 65551:31. In the BGP routing table, there are two entries per partner PE, one for the BGP-MH update (with site-id=1) and the other for the BGP-VPWS update (with VE-ID=1).

```

*A:PE-2# show router bgp routes l2-vpn rd 65551:31
=====
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

```

## Dual Homed BGP VPWS with Single Pseudowire

```

=====
BGP L2VPN Routes
=====
Flag  RouteType      Prefix          MED
      RD          SiteId         Label
      Nexthop     VeId           LocalPref
      As-Path     BaseOffset     BlockSize
                        vplsLabelBa
                        se
-----
u*>i  MultiHome      -              -
      65551:31    1              -
      192.0.2.1  -              200
      No As-Path -              -
u*>i  VPWS           -              -
      65551:31    -              -
      192.0.2.1  1              1
                        262139
      No As-Path  2              262139
-----
Routes : 2
=====
*A:PE-2#

```

```

*A:PE-2# show router bgp routes l2-vpn rd 65551:33
=====
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         Label
      Nexthop     VeId           LocalPref
      As-Path     BaseOffset     BlockSize
                        vplsLabelBa
                        se
-----
u*>i  MultiHome      -              -
      65551:33    1              -
      192.0.2.3  -              10
      No As-Path -              -
u*>i  VPWS           -              -
      65551:33    -              -
      192.0.2.3  1              1
                        262140
      No As-Path  2              262140
-----
Routes : 2
=====
*A:PE-2#

```

The route to PE-1 has the higher site preference, so it is selected as the target for the pseudowire.

```

A:PE-2# show service l2-route-table bgp-vpws detail
=====
Services: L2 Bgp-Vpws Route Information - Summary
=====

```



```

<snipped>

Svc Id      : 3
VeId       : 1
PW Temp Id  : 3
RD         : *65551:31
Next Hop   : 192.0.2.1
State (D-Bit) : up(0)
Path MTU   : 1514
Control Word : 0
Seq Delivery : 0
Status     : active
Tx Status  : active
CSV        : 0
Preference : 200
Sdp Bind Id : 17407:4294967295
<snipped>

```

```

=====
*A:PE-2#

```

After disabling the SAP in the service on PE-1, BGP UPDATE messages are received. The VPLS/VPWS message received on PE-2 from PE-1 shows in the CSV that the access circuit is down (the CSV has the most-significant bit set to 1 (0x80)), so PE-2 selects the update from PE-3 to create the pseudowire. The BGP-MH update received by PE-2 from PE-1 also shows that the local site is down as indicated by the flags=D.

Note in the debug output below,

- BGP MH (multi-homing) entry uses encap-type=19.
- BGP VPWS entry uses encap-type=5 (Ethernet raw mode).

```

240 2014/02/19 15:08:15.65 UTC MINOR: DEBUG #2001 Base Peer 1: 192.0.2.5
"Peer 1: 192.0.2.5: UPDATE
Peer 1: 192.0.2.5 - Received BGP UPDATE:
  Withdrawn Length = 0
  Total Path Attr Length = 86
  Flag: 0x90 Type: 14 Len: 28 Multiprotocol Reachable NLRI:
    Address Family L2VPN
    NextHop len 4 NextHop 192.0.2.1
    [MH] site-id: 1, RD 65551:31
  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: 0
  Flag: 0x80 Type: 9 Len: 4 Originator ID: 192.0.2.1
  Flag: 0x80 Type: 10 Len: 4 Cluster ID:
    1.1.1.1
  Flag: 0xc0 Type: 16 Len: 16 Extended Community:
    target:65551:31
    l2-vpn/vrf-imp:Encap=19: Flags=D: MTU=0: PREF=200
"

```

## Dual Homed BGP VPWS with Single Pseudowire

```
241 2014/02/19 15:08:15.64 UTC MINOR: DEBUG #2001 Base Peer 1: 192.0.2.5
"Peer 1: 192.0.2.5: UPDATE
Peer 1: 192.0.2.5 - Received BGP UPDATE:
  Withdrawn Length = 0
  Total Path Attr Length = 90
  Flag: 0x90 Type: 14 Len: 32 Multiprotocol Reachable NLRI:
    Address Family L2VPN
    NextHop len 4 NextHop 192.0.2.1
    [VPLS/VPWS] veid: 1, vbo: 2, vbs: 1, label-base: 262139, RD 65551:31, cs v:
0x80
  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: 0
  Flag: 0x80 Type: 9 Len: 4 Originator ID: 192.0.2.1
  Flag: 0x80 Type: 10 Len: 4 Cluster ID:
    1.1.1.1
  Flag: 0xc0 Type: 16 Len: 16 Extended Community:
    target:65551:31
    l2-vpn/vrf-imp:Encap=5: Flags=D: MTU=1514: PREF=200
""
```

The result can be shown on PE-2 as now the spoke SDP is up (active) to PE-3.

```
A:PE-2# show service l2-route-table bgp-vpws detail
=====
Services: L2 Bgp-Vpws Route Information - Summary
=====

Svc Id       : 3
VeId        : 1
PW Temp Id   : 3
RD          : *65551:33
Next Hop    : 192.0.2.3
State (D-Bit) : up(0)
Path MTU    : 1514
Control Word : 0
Seq Delivery : 0
Status      : active
Tx Status    : active
CSV         : 0
Preference  : 200
Sdp Bind Id : 17407:4294967295
=====
A:PE-2#
```



## Dual Homed BGP VPWS with Active/Standby Pseudowire

being created to PE-3. If PE-1 fails, or the multi-homing site fails over to PE-3, then the pseudowire from PE-2 to PE-3 will become active (used as the data path between site A and B).

### PE-1 configuration:

```
A:PE-1# configure service
pw-template 3 create
exit
epipe 4 customer 1 create
  bgp
    route-distinguisher 65551:41
    route-target export target:65551:41 import target:65551:41
    pw-template-binding 3
  exit
exit
bgp-vpws
  ve-name "PE-1"
  ve-id 1
  exit
  remote-ve-name "PE-2"
  ve-id 2
  exit
  no shutdown
exit
site "siteB" create
  site-id 1
  sap 1/1/4:44
  site-preference 200
  no shutdown
exit
sap 1/1/4:44 create
exit
no shutdown
exit
exit
```

## PE-3 configuration:

Note that the local VE-ID is 3 (different from previous example).

```
A:PE-3# configure service
pw-template 3 create
exit
epipe 4 customer 1 create
  bgp
    route-distinguisher 65551:43
    route-target export target:65551:41 import target:65551:41
    pw-template-binding 3
  exit
exit
bgp-vpws
  ve-name "PE-3"
  ve-id 3
  exit
  remote-ve-name "PE-2"
  ve-id 2
  exit
  no shutdown
exit
site "siteB" create
  site-id 1
  sap 1/1/4:44
  site-preference 10
  no shutdown
exit
sap 1/1/4:44 create
exit
no shutdown
exit
exit
```

## PE-2 configuration:

Note that there are two remote VE names configured, PE-1 and PE-3 (this is the maximum number allowed).

```
A:PE-2# configure service
pw-template 3 create
exit
epipe 4 customer 1 create
  bgp
    route-distinguisher 65551:42
    route-target export target:65551:41 import target:65551:41
    pw-template-binding 3
  exit
exit
bgp-vpws
  ve-name "PE-2"
  ve-id 2
  exit
  remote-ve-name "PE-1"
```

## Dual Homed BGP VPWS with Active/Standby Pseudowire

```
                ve-id 1
            exit
            remote-ve-name "PE-3"
                ve-id 3
            exit
            no shutdown
        exit
        sap 1/1/4:44 create
        exit
        no shutdown
    exit
exit
```

Compared to the single pseudowire solution, both pseudowires are signaled and up on all PEs. The pseudowire with the higher preference is forwarding traffic (to PE-1), while the Tx Status on the other one is set to inactive.

Verified on PE-2:

```
A:PE-2# show service l2-route-table bgp-vpws detail
=====
Services: L2 Bgp-Vpws Route Information - Summary
=====
<snipped>

Svc Id      : 4
VeId       : 1
PW Temp Id : 3
RD         : *65551:41
Next Hop   : 192.0.2.1
State (D-Bit) : up(0)
Path MTU   : 1514
Control Word : 0
Seq Delivery : 0
Status     : active
Tx Status  : active
CSV       : 0
Preference : 200
Sdp Bind Id : 17407:4294967294

Svc Id      : 4
VeId       : 3
PW Temp Id : 3
RD         : *65551:43
Next Hop   : 192.0.2.3
State (D-Bit) : up(0)
Path MTU   : 1514
Control Word : 0
Seq Delivery : 0
Status     : active
Tx Status  : inactive
CSV       : 0
Preference : 10
Sdp Bind Id : 17406:4294967292
=====
A:PE-2#
```

The choice of pseudowire to be used to transmit traffic from PE-2 to PE-1 can also be seen in the endpoint created in the BGP VPWS service. Endpoints are automatically created for the pseudowires within a BGP VPWS service regardless of whether active/standby pseudowires are used; these endpoints are created with a system generated name that ends with the BGP VPWS service id.

```
A:PE-2# show service id 4 endpoint
=====
Service 4 endpoints
=====
Endpoint name           : _tmnx_BgpVpws-4
Description             : Automatically created BGP-VPWS endpoint
Creation Origin        : bgpVpws
Revert time            : 0
Act Hold Delay         : 0
Standby Signaling Master : false
Standby Signaling Slave  : false
Tx Active (SDP)        : 17407:4294967294
Tx Active Up Time      : 0d 00:24:26
Revert Time Count Down : N/A
Tx Active Change Count : 1
Last Tx Active Change  : 03/10/2014 12:06:37
-----
Members
-----
Spoke-sdp: 17406:4294967292 Prec:4           Oper Status: Up
Spoke-sdp: 17407:4294967294 Prec:4           Oper Status: Up
=====
A:PE-2#
```

Note that the following command has no effect on an automatically created VPWS endpoint.

```
tools perform service id <service-id> endpoint <endpoint-name> force-switchover
```

## Conclusion

BGP VPWS allows the delivery of Layer 2 virtual private wire services to customers where BGP is commonly used. This example shows the configuration of single and dual-homed BGP VPWS services together with the associated show output, which can be used to verify and troubleshoot them.