

OSPF as PE–CE Protocol and Loop–Prevention Techniques in MPLS L3 VPN Configuration Example



Document ID: 118800

Contributed by Naveen Bansal and Rahul Kukreja, Cisco TAC Engineers
Mar 06, 2015

Contents

Introduction

Prerequisites

- Requirements

- Components Used

Background Information

Configure

- Network Diagram

- Configurations

 - DN Bit

 - Domain Tag

Verify

Troubleshoot

Introduction

This document describes the loop prevention features and the minimum configuration steps when you run the Open Shortest Path First (OSPF) Routing Protocol between Provider Edge (PE) and Customer Edge (CE) Routers. It presents a network scenario that depicts the use of Downward Bit (DN), which is an option in the Link State Advertisement (LSA) and Domain Tag.

Prerequisites

Requirements

Cisco recommends that you have knowledge of OSPF and Multiprotocol Label Switching (MPLS) Layer 3 VPN.

Components Used

This document is not restricted to specific software and hardware versions.

The information in this document was created from the devices in a specific lab environment. All of the devices used in this document started with a cleared (default) configuration. If your network is live, make sure that you understand the potential impact of any command.

Background Information

The Service Provider (SP) and the CE Router exchange routes with a routing protocol to which the SP and Customer jointly agree. The scope of this document is to describe the loop-prevention mechanism when OSPFv2 is used.


When OSPFv2 is used on a PE-CE link that belongs to a particular Virtual Routing and Forwarding (VRF) or VPN, the PE router:

- Redistributes the routes received via OSPF for that VPN into Multiprotocol-Border Gateway Protocol (MP-BGP) and advertises it to the other PE routers.
- Redistributes the BGP routes installed in the VPN via MP-BGP into the OSPF Instance for that VPN and advertises it to the CE Routers.

Configure

Network Diagram

Consider this network topology in order to understand the loop-prevention techniques.



In this setup, there is a possibility of a loop. For example, if CE1 advertises the OSPF LSA Type 1 to PE1, which redistributes the route into VPNv4 and advertises it to PE2, then PE2 in turn advertises the Summary LSA to CE2. This route received by CE2 could be advertised back to PE3. The third PE router learns the OSPF route, which is better than the BGP route, and re-advertises the route into BGP as local to the Customer Site 2. PE3 never learns that the route that was advertised was not originated from Customer Site 2.

In order to overcome this situation, when the routes are redistributed from MP-BGP into OSPF, then they are marked with a DN Bit in LSA Type 3, 5, or 7 and have the domain tag for Type 5 and 7 LSA.

Configurations

Here is the sample configuration on PE routers. This configuration includes the VRF configuration, the OSPF Process 2 that runs between the PE-CE Routers, the OSPF Process 1 that runs as Interior Gateway Protocol (IGP) in the MPLS Core, and the MP-BGP configuration.

```

ip vrf A
rd 1:1
route-target both 65000:1
route-target import 65000:2
route-target import 65000:3
! VRF A configuration with Route Distinguisher and Route Targets
! 2:2 and 3:3 import route-target is configured as export route-target on PE2 and PE3

interface Ethernet0/0
ip vrf forwarding A
ip address 10.10.23.3 255.255.255.0
! Eth0/0 Interface - CE1 Facing

router ospf 1
router-id 10.1.1.1
network 0.0.0.0 255.255.255.255 area 0
! OSPF Process 1 running in MPLS Core and Loopback1


router ospf 2 vrf A
redistribute bgp 65000 subnets
network 10.10.23.3 0.0.0.0 area 0
! OSPF Process 2 in VRF A and redistribution of BGP Routes installed via MP-BGP in the VRF A into OSPF 2

router bgp 65000
no bgp default ipv4-unicast
neighbor 10.2.2.2 remote-as 65000
neighbor 10.2.2.2 update-source Loopback1
neighbor 10.3.3.3 remote-as 65000
neighbor 10.3.3.3 update-source Loopback1
!
address-family vpnv4
neighbor 10.2.2.2 activate
neighbor 10.2.2.2 send-community extended
neighbor 10.3.3.3 activate
neighbor 10.3.3.3 send-community extended
exit-address-family
!
address-family ipv4 vrf A
redistribute ospf 2 metric 10 match internal external 1 external 2
exit-address-family
! BGP VPNv4 and OSPF Process 2 configuration
! Redistribution of OSPF Process 2 into BGP, so that the routes could be advertised via MP BGP to PE2 and PE3

```

DN Bit

The previously unused bit in the OSPF LSA Options Field is referred to as the DN Bit. This bit is set on Type 3, 5, and 7 LSA when the MP-BGP routes are redistributed into OSPF. When the other PE Router receives the LSA from a CE router Type 3, 5, or 7 LSA with the DN Bit set, the information from that LSA is not used in the OSPF route calculation.



Based on the network topology, PE2 sets the DN Bit for the redistributed LSA and this LSA is never considered for route calculation in OSPF Process 2 on PE3. So PE3 never redistributes this route back into MP-BGP.

Here is an example of the OSPF Header that shows the DN Bit Set, when the route was advertised by PE Router for Type 3 LSA:

Open Shortest Path First

OSPF Header

```
Version: 2
Message Type: LS Update (4)
Packet Length: 56
Source OSPF Router: 10.10.23.3 (10.10.23.3)
Area ID: 0.0.0.0 (0.0.0.0) (Backbone)
Checksum: 0x4034 [correct]
Auth Type: Null (0)
Auth Data (none): 0000000000000000
```

LS Update Packet

```
Number of LSAs: 1
Summary-LSA (IP network)
.000 1110 0001 0000 = LS Age (seconds): 3600
0... .. = Do Not Age Flag: 0
Options: 0xa2 (DN, DC, E)
1... .. = DN: Set
.0... .. = O: Not set
..1. .... = DC: Demand Circuits are supported
...0 .... = L: The packet does NOT contain LLS data block
.... 0... = NP: NSSA is NOT supported
.... .0.. = MC: NOT Multicast Capable
.... ..1. = E: External Routing Capability
.... ...0 = MT: NO Multi-Topology Routing
```

Domain Tag

The Domain Tag is applicable only for the OSPF Type 5 and Type 7 LSA. When the VPNv4 routes are redistributed from MP-BGP into OSPF on PE Router, the Domain Tag is set for OSPF External Routes. The tag could either be manually set with the *domain-tag* command under OSPF Process or a 32-bit value can be automatically generated:

Manually configured tags:

```

0 0 0 0 0 0 0 0 0 0 1 1 1 1 1 1 1 1 1 1 2 2 2 2 2 2 2 2 2 2 3 3
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-----+-----+-----+-----+-----+-----+-----+-----+-----+
|0|                                     LocalInfo                               |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+

```

```

Command:      router ospf
              domain-tag <1-4294967295>
              OSPF domain tag - 32-bit value

```

Automatic Tag Generation: 32 bits

When the tag is automatically generated, the high order bit is set to 1
 c bit is set when Origin is EGP or IGP
 p1 2 bits are for Path Length information
 ArbitraryTag 12 bits defaults to 0
 AutonomousSystem 16 bits indicating the AS number
 The other bits are defined below:

```

0 0 0 0 0 0 0 0 0 0 1 1 1 1 1 1 1 1 1 1 2 2 2 2 2 2 2 2 2 2 3 3
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-----+-----+-----+-----+-----+-----+-----+-----+-----+
|1|c|p 1|           ArbitraryTag           |           AutonomousSystem           |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+

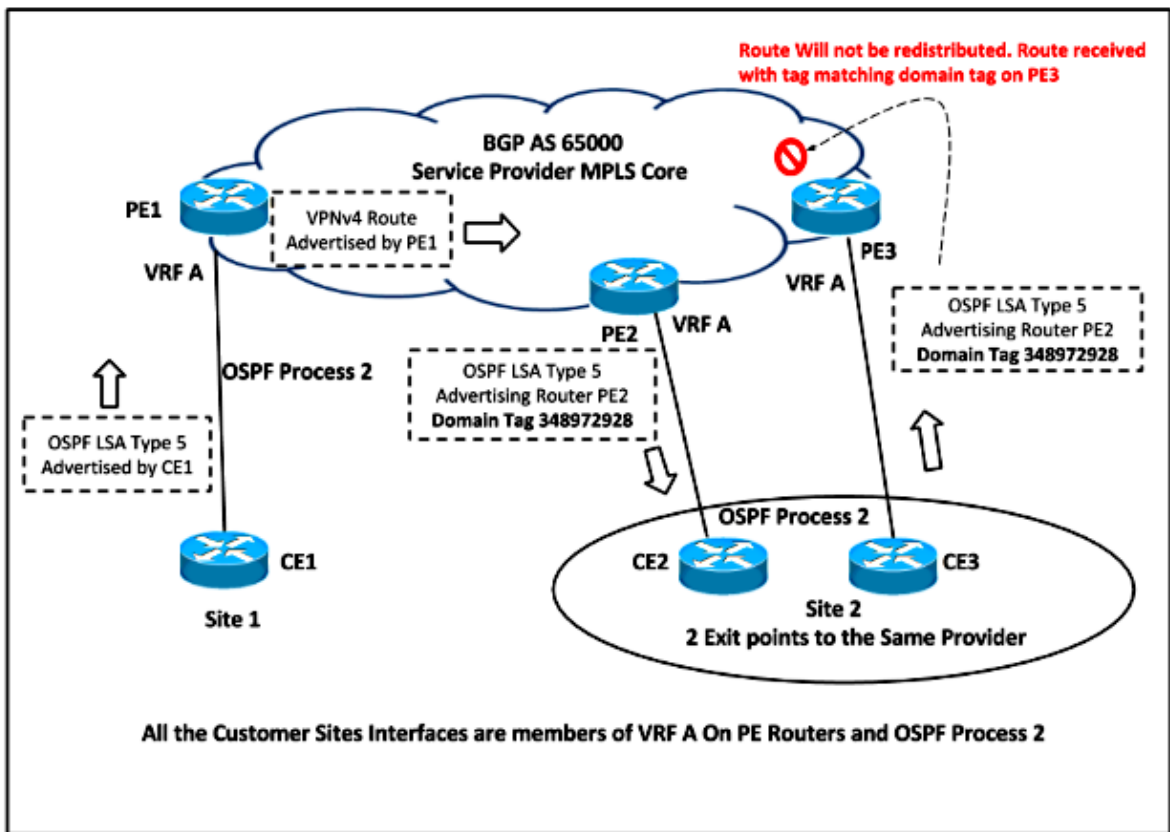
```

In our example the routes received on CE2 from PE1, the tag is set to **3489725928**
 Binary Representation:

```

11010000 00000000 11111101 11101000
          <-----65000----->
          Autonomous System Number

```



Based on the network topology, PE2 sets the Domain Tag for Type 5 and Type 7 LSA when it redistributes the VPNv4 route into OSPF. This LSA is never considered for route calculation because the DN Bit is already set, but it also has the Domain Tag set, so the LSA is ignored because the Domain Tag matches the VPN / VRF Tag. Hence the route is never redistributed into OSPF.

This example shows LSA Type 5 ignored when it is received with the Domain Tag Set the same as the local VRF Domain Tag on PE3 from CE3:

```
*Jan 31 00:29:23.947: OSPF-2 EXTER: adv_rtr 10.10.57.5, age 3, seq 0x80000001,
metric 10, metric-type 2, fw-addr 0.0.0.0
*Jan 31 00:29:23.947: OSPF-2 EXTER: Tag equals to VPN Tag, ignoring the LSA
*Jan 31 00:29:23.947: OSPF-2 EXTER: Process partial nssa spf queue
```

```
PE3#show ip ospf database external 192.168.5.5
```

```
OSPF Router with ID (10.3.3.3) (Process ID 1)
```

```
OSPF Router with ID (10.10.68.6) (Process ID 2)
```

```
Type-5 AS External Link States
```

```
LS age: 38
Options: (No TOS-capability, DC)
LS Type: AS External Link
Link State ID: 192.168.5.5 (External Network Number )
Advertising Router: 10.10.57.5
LS Seq Number: 80000001
Checksum: 0x89A3
Length: 36
Network Mask: /32
Metric Type: 2 (Larger than any link state path)
```

```
MTID: 0
Metric: 10
Forward Address: 0.0.0.0
External Route Tag: 3489725928
```

Verify

The commands to discover if the DN Bit is set for the LSA and the Domain Tag applied are the same that are used in order to check the LSA Database.

This output shows the example for OSPF Type 3 and Type 5 LSA and highlights the DN Bit and Tag Set when the VPNv4 routes are redistributed into OSPF on PE2:

<pre>LSA Type 3 PE2#sh ip ospf 2 database summary 192.168.1.1 OSPF Router with ID (10.10.57.5) (Process ID 2) Summary Net Link States (Area 0) LS age: 1735 Options: (No TOS-capability, DC, Downward) LS Type: Summary Links(Network) Link State ID: 192.168.1.1 (summary Network Number) Advertising Router: 10.10.57.5 LS Seq Number: 80000001 Checksum: 0x46AE Length: 28 Network Mask: /32 MTID: 0 Metric: 10 LS age: 1738 Options: (No TOS-capability, DC, Downward) LS Type: Summary Links(Network) Link State ID: 192.168.1.1 (summary Network Number) Advertising Router: 10.10.68.6 LS Seq Number: 80000001 Checksum: 0xF2F5 Length: 28 Network Mask: /32 MTID: 0 Metric: 10</pre>	<pre>LSA Type 5 PE2#sh ip ospf 2 database external 192.168.5.5 OSPF Router with ID (10.10.57.5) (Process ID 2) Type-5 AS External Link States LS age: 1756 Options: (No TOS-capability, DC, Downward) LS Type: AS External Link Link State ID: 192.168.5.5 (External Network Number) Advertising Router: 10.10.57.5 LS Seq Number: 80000001 Checksum: 0x2AA Length: 36 Network Mask: /32 Metric Type: 2 (Larger than any link state path) MTID: 0 Metric: 10 Forward Address: 0.0.0.0 External Route Tag: 3489725928 LS age: 1759 Options: (No TOS-capability, DC, Downward) LS Type: AS External Link Link State ID: 192.168.5.5 (External Network Number) Advertising Router: 10.10.68.6 LS Seq Number: 80000001 Checksum: 0xAEF1 Length: 36 Network Mask: /32 Metric Type: 2 (Larger than any link state path) MTID: 0 Metric: 10 Forward Address: 0.0.0.0 External Route Tag: 3489725928</pre>
---	---

Note: MPLS VPN OSPF PE-CE always includes the loop-prevention mechanism in order to handle issues. In the older Cisco IOS®, the Per original IETF draft Type 3 LSAs use the DN Bit in LSA and Type 5 LSAs use a tag. The newer RFC 4576 mandates use of DN Bit for both Type 3 and Type 5 LSAs.

This was committed via Cisco bug ID CSCtw79182.

The PE routers with Cisco IOS images with the fix of this defect will originate Type 5 external LSAs with both DN Bit and a tag as loop-prevention mechanisms. Previous Cisco IOS versions advertised the only tag for this purpose for External Routes.

The change in behavior was made because a tag is possible to rewrite (by changing VPN domain ID or via route-map) but the DN Bit is not user-controllable. In some corner-case designs, some customers might have deliberately disabled the loop-prevention mechanism with an overwrite of tags of external LSAs in order for the PE router to prefer the OSPF route over the BGP route.

In newer Cisco IOS versions, this is not possible. The vast majority of customers that use PE-CE OSPF in a textbook configuration will not be affected. Customers that override tags MIGHT see a change in behavior.

Troubleshoot

There is currently no specific troubleshooting information available for this configuration.

Updated: Mar 06, 2015

Document ID: 118800
