Why OSPF Demand Circuit Keeps Bringing Up the Link

Document ID: 4808

Contents

Introduction
Prerequisites
Requirements
Components Used
Conventions
Sample Network
Reason 1: Change in the Network Topology
Solution
Reason 2: Network Type Defined as Broadcast
Solution
Reason 3: One or More Routers Do not Understand Demand Circuit
Reason 4: Host Route Is Redistributed into the OSPF Database
Solution 1: Use the no peer neighbor−route Command
Solution 2: Use the route−map Command
Solution 3: Use a Different Major Net
Reason 5: OSPF Demand Circuit Is Configured Over an Asynchronous Interface
Solution
Reason 6: OSPF Demand Circuit Is Configured Over a Multilink PPP
Solution
Related Information

Introduction

When an Open Shortest Path First (OSPF) link is configured as demand circuit, OSPF Hellos are suppressed and periodic LSA refreshes are not flooded over the link. These packets bring up the link only when they are exchanged for the first time, or when a change occurs in the information they contain. This allows the underlying Data Link Layer to be closed when the network topology is stable. A demand circuit that goes up and down indicates a problem that needs to be investigated. This document demonstrates some possible causes and offers solutions.

For more information on demand circuit, refer to OSPF Demand Circuit Feature.

Prerequisites

Requirements

There are no specific requirements for this document.

Components Used

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

Conventions

For more information on document conventions, see the Cisco Technical Tips Conventions.

Sample Network

The problem mentioned above is described with the following network diagram and configuration.

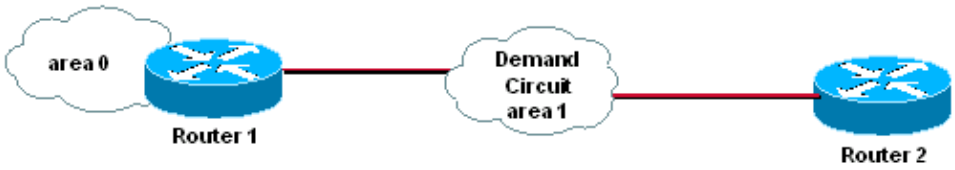

Router 1	Router 2
interface BRI1/1 ip address 192.158.254.13 255.255.255.252 **ip ospf demand-circuit** router ospf 20 network 192.158.254.0 0.0.0.255 area 0	interface BRI1/0 ip address 192.158.254.14 255.255.255.252 router ospf 20 network 192.158.254.0 0.0.0.255 area 0

Note: You need to configure the demand circuit at one end of the link only. However, if you configure this command on both ends it does not cause any harm.

In the diagram above, Routers 1 and 2 are running OSPF demand circuit across the ISDN link. The link between Routers 1 and 2 keeps coming up, which defeats the purpose of OSPF demand circuit. The output of the **show dialer** command shows that the link came up because of the OSPF multicast Hello packet.

```
Router1# show dialer
BRI1/1:1 - dialer type = ISDN
Idle timer (120 secs), Fast idle timer (20 secs)
Wait for carrier (30 secs), Re-enable (2 secs)
Dialer state is data link layer up
Dial reason: ip (s=192.168.254.13, d=224.0.0.5)
```

The link can be brought up for several reasons. Below we explore several common cases and offer solutions.

Reason 1: Change in the Network Topology

Whenever there is a change in an OSPF network topology, the OSPF routers must be notified. In this situation the OSPF demand circuit should be brought up so that the neighbors can exchange the new information. Once the new database is exchanged the link can go down again and the adjacency remains in the FULL state.

Solution

To determine if the link is brought up due to a change in network topology, use the **debug ip ospf monitor** command. It shows which LSA is changing, as seen below:

```
Router1# debug ip ospf monitor
```

```
OSPF: Schedule SPF in area 0.0.0.0
      Change in LS ID 192.168.246.41, LSA type R,
OSPF: schedule SPF: spf_time 1620348064ms wait_interval 10s
```

The output above shows there was a change in the router LSA with the router ID of 192.168.246.41, which causes the database to be resynchronized. If the network is stable, then this debug output shows nothing.

To reduce the affect of link flaps on the demand circuit, configure the area that contains the demand circuit as totally stub. If this isn't doable, and there is a constant link flap within the network, demand circuit might not be a good choice for you.

Reason 2: Network Type Defined as Broadcast

When you configure demand circuit on a link, the link type must be defined as point−to−point or point−to−multipoint. Any other link type can cause the link to come up unnecessarily because the OSPF Hellos are not suppressed if the network type is anything other than point−to−point or point−to−multipoint. Following is a sample configuration to illustrate this problem on Routers 1 and 2.

Router 1	Router 2
interface BRI1/1 ip address 192.158.254.13 255.255.255.252 **ip ospf network broadcast** router ospf 20 network 192.158.254.0 0.0.0.255 area 0	interface BRI1/0 ip address 192.158.254.14 255.255.255.252 **ip ospf network broadcast** router ospf 20 network 192.158.254.0 0.0.0.255 area 0

With the network type defined as broadcast, the OSPF Hellos bring the link up at every Hello interval. The **show dialer** output shows that the last time the link was brought up was because of an OSPF Hello.

```
Router1# show dialer
BRI1/1:1 − dialer type = ISDN
Idle timer (120 secs), Fast idle timer (20 secs)
Wait for carrier (30 secs), Re−enable (2 secs)
Dialer state is data link layer up
Dial reason: ip (s=192.168.254.13, d=224.0.0.5)
Interface bound to profile Di1
Current call connected 00:00:08
Connected to 57654 (R2)
```

Solution

To solve this problem, either change the network type to point−to−point or point−to−multipoint. Here we remove the network type broadcast so by default it is configured as point−to−point.

Router 1	Router 2
interface BRI1/1 ip address 192.158.254.13 255.255.255.252 router ospf 20 network 192.158.254.0 0.0.0.255 area 0	interface BRI1/0 ip address 192.158.254.14 255.255.255.252 router ospf 20 network 192.158.254.0 0.0.0.255 area 0

By changing the network type to point−to−point or point−to−multipoint, the OSPF Hellos are suppressed on the link, and the demand circuit link stops flapping.

Reason 3: One or More Routers Do not Understand Demand Circuit

When one or more routers in the OSPF domain do not understand demand circuit, a periodic LSA refresh occurs. See the When Is a Periodic LSA Refresh Sent Over an OSPF Demand Circuit? section of this document to learn how to resolve this issue.

Reason 4: Host Route Is Redistributed into the OSPF Database

Let's consider the following network diagram as an example:

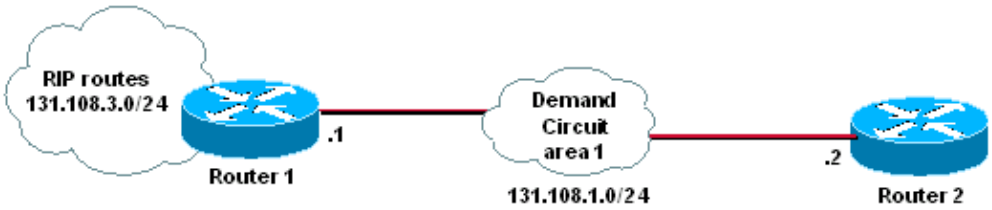

The link between Routers 1 and 2 is 131.108.1.0/24, and demand circuit is configured between Routers 1 and 2. Router 1 is redistributing Routing Information Protocol (RIP) routes into OSPF.

Router 1

```
router ospf 1
 redistribute rip subnets
 network 131.108.1.0 0.0.0.255 area 1
!
router rip
 network 131.108.0.0
```

Since the link encapsulation type is PPP, both routers install a host route for the other side of the link as shown below.

```
Router1# show ip route 131.108.1.2
Routing entry for 131.108.1.2/32
  Known via "connected", distance 0, metric 0 (connected, via interface)
  Routing Descriptor Blocks:
  * directly connected, via BRI1/1
      Route metric is 0, traffic share count is 1
```

Interior Gateway Routing Protocol (IGRP) and RIP are classful routing protocols, and therefore the network statement in the configuration is for a classful network of 131.108.0.0. Because of this the host route of 131.108.1.2/32 is considered to be originated by RIP and gets redistributed into OSPF as an external route as shown below.

```
Router1# show ip ospf database external 131.108.1.2

       OSPF Router with ID (131.108.3.1) (Process ID 1)


                Type-5 AS External Link States
```

```
LS age: 298
Options: (No TOS-capability, DC)
LS Type: AS External Link
Link State ID: 131.108.1.2 (External Network Number )
Advertising Router: 131.108.3.1
LS Seq Number: 80000001
Checksum: 0xDC2B
Length: 36
Network Mask: /32
      Metric Type: 2 (Larger than any link state path)
      TOS: 0
      Metric: 20
      Forward Address: 0.0.0.0
      External Route Tag: 0
```

When the link goes down the /32 disappears and OSPF understands this as a change in topology. The demand circuit brings the link up again to propagate the MAXAGE version of the /32 mask to its neighbor. When the link comes up, the /32 mask becomes valid again so the LSA age gets reset. Then after the dead timer of the link kicks in, the link goes down again. This process repeats itself and the demand circuit link keeps flapping. There are three ways to solve this problem shown below.

Solution 1: Use the no peer neighbor−route Command

Under the BRI interface that is running demand circuit, configure **no peer neighbor−route**. This prevents the /32 mask from being installed. You can use the configuration shown below on Router 1 only, but we recommend configuring this command on both sides for consistency.

```
R1# configure terminal
R1(config)# interface BRI1/1
R1(config-if)# no peer neighbor-route
```

Solution 2: Use the route−map Command

When redistributing from RIP into OSPF, use the **route−map** command and deny /32 so it does not get injected into the OSPF database. This configuration command is required only on the router that is doing the redistribution, which in our example is Router 1.

First we have to create an access list to match the /32 mask. Then we apply this access list to the route map and use the route map when applying the **redistribution** command as shown below.

```
R1# configure terminal
R1(config)# access-list 1 deny host 131.108.1.2
R1(config)# access-list 1 permit any

R1# configure terminal
R1(config)# route-map rip-ospf
R1(config-route-map)# match ip address 1

R1(config)# router ospf 1
R1(config-router)# redistribute rip subnets route-map rip-ospf
```

Solution 3: Use a Different Major Net

Use a different major net for either the RIP or OSPF domain. The idea is to have a different major net on the demand circuit link so when the link comes up under PPP encapsulation it installs the host route for the other side of the link. If the host route is in a different major net than the one being used in RIP, RIP does not own

this PPP−installed host route because it does not have a network statement for the major net. The network diagram below shows an example.

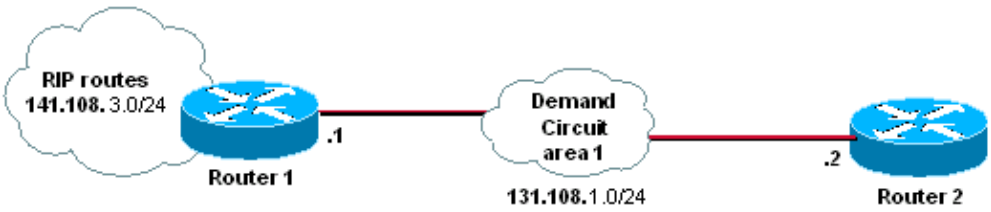

The RIP domain is now under the 141.108.0.0 network while the OSPF domain (and the demand circuit link) is under the 131.108.0.0 network.

Reason 5: OSPF Demand Circuit Is Configured Over an Asynchronous Interface

When you configure a demand circuit over an asynchronous (async) interface, then when the Layer 2 goes down, the actual physical interface goes down. This triggers a change in the OSPF database and the async interface comes back up again to exchange the database. The Layer 2 goes down again, and this will trigger the change in database again, so this process keeps on repeating itself.

The following scenario is used to reproduce the above problem.

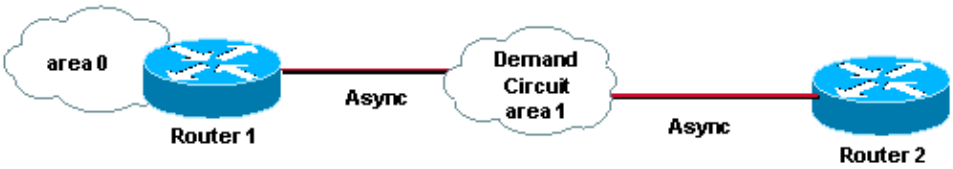

The following configuration is used for the above scenario.

Router 1	Router 2
interface Async 1 ip address 192.158.254.13 255.255.255.252 encapsulation ppp ip ospf demand−circuit dialer in−band async default routing async mode dedicated ppp authentication chap ppp chap hostname Router1 ppp chap password 7 13061E010803 ! router ospf 20 network 192.158.254.0 0.0.0.255 area 1	interface Async 1 ip address 192.158.254.14 255.255.255.252 encapsulation ppp ip ospf demand−circuit dialer in−band dialer map ip 192.158.254.13 broadcast 12345 dialer−group 2 async default routing async mode dedicated ppp authentication chap callin ! dialer−list 2 protocol ip permit ! router ospf 20 network 192.158.254.0 0.0.0.255 area 1

The OSPF default network type is point−to−point on an async interface, but still the demand circuit keeps bringing up the link.

```
Rouer1# show ip ospf interface Async1
 Async1 is up, line protocol is up (spoofing)
   Internet Address 192.158.254.13/32, Area 1
   Process ID 1, Router ID 1.1.1.1, Network Type POINT_TO_POINT, Cost:869
   Transmit Delay is 1 sec, State POINT_TO_POINT,
   Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
    Hello due in 00:00:02
   Index 1/2, flood queue length 0
   Next 0x0(0)/0x0(0)
   Last flood scan length is 0, maximum is 1
   Last flood scan time is 0 msec, maximum is 0 msec
   Neighbor Count is 0, Adjacent neighbor count is 0
   Suppress hello for 0 neighbor(s)
```

Solution

The reason the demand circuit keeps bringing up the link is because when the Layer 2 goes down after the idle timeout expires, the whole interface goes down. But in the case of BRI or PRI, when one of the channels goes down, the interface still remains up (in spoofing mode). To solve the problem you must configure a dialer interface because it never goes down. A dialer interface remains up (in spoofing mode).

Router 1	Router 2
interface Async 1 no ip address encapsulation ppp async default routing async mode dedicated dialer in−band dialer rotary−group 0 ! **interface Dialer0** **ip address 192.158.254.13 255.255.255.252** **encapsulation ppp** **ip ospf demand−circuit** **ppp authentication chap** **ppp chap hostname Router1** **ppp chap password 7 13061E010803** ! router ospf 20 network 192.158.254.0 0.0.0.255 area 0	interface Async 1 no ip address encapsulation ppp async default routing async mode dedicated dialer in−band dialer rotary−group 0 ! **interface Dialer0** **ip address 192.158.254.14 255.255.255.252** **encapsulation ppp** **ip ospf demand−circuit** **dialer map ip 192.158.254.13 broadcast 12345** **dialer−group 2** **ppp authentication callin** **!** **dialer−list 2 protocol ip permit** ! router ospf 20 network 192.158.254.0 0.0.0.255 area 0

Since the dialer interface never goes down, it will not create the problem that is created when an async interface goes down.

Reason 6: OSPF Demand Circuit Is Configured Over a Multilink PPP

The multilink PPP feature can be used for load balancing purposes in cases there are multiple WAN links. One important thing to remember in terms of OSPF is the bandwidth of the multilink PPP. When multiple links are combined, the bandwidth of the multilink interface will change.

The following scenario is used to reproduce the above problem.

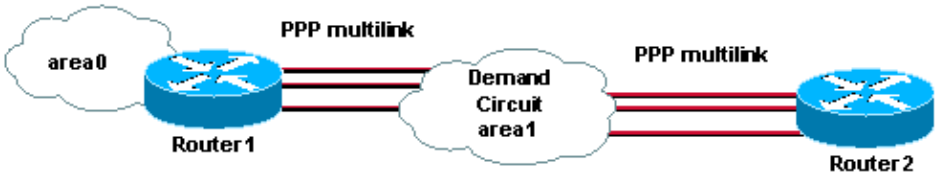

The following configuration is used for the above scenario.

Router 1	Router 2
interface Multilink1 ip address 192.158.254.1 255.255.255.0 no cdp enable ppp multilink no ppp multilink fragmentation multilink-group 1 ! interface Serial0/1/0:0 no ip address ip route-cache distributed encapsulation ppp tx-queue-limit 26 no fair-queue ppp multilink multilink-group 1 ! interface Serial0/1/1:0 no ip address ip route-cache distributed encapsulation ppp tx-queue-limit 26 no fair-queue ppp multilink multilink-group 1 ! interface Serial0/1/2:0 no ip address ip route-cache distributed encapsulation ppp tx-queue-limit 26 no fair-queue ppp multilink multilink-group 1 ! router ospf 20 network 192.158.254.0 0.0.0.255 area 1	interface Multilink1 ip address 192.158.254.2 255.255.255.0 no cdp enable ppp multilink no ppp multilink fragmentation multilink-group 1 ! interface Serial0/1/0:0 no ip address ip route-cache distributed encapsulation ppp tx-queue-limit 26 no fair-queue ppp multilink multilink-group 1 ! interface Serial0/1/1:0 no ip address ip route-cache distributed encapsulation ppp tx-queue-limit 26 no fair-queue ppp multilink multilink-group 1 ! interface Serial0/1/2:0 no ip address ip route-cache distributed encapsulation ppp tx-queue-limit 26 no fair-queue ppp multilink multilink-group 1 ! router ospf 20 network 192.158.254.0 0.0.0.255 area 1

The following output shows that there are three serial interfaces bundled together in the multilink PPP.

```
Router1# show ppp multilink
 Multilink1, bundle name is Router2
   Bundle up for 00:05:35
   Bundle is Distributed
   0 lost fragments, 0 reordered, 0 unassigned
   0 discarded, 0 lost received, 3/255 load
   0x1226 received sequence, 0x1226 sent sequence
   Member links: 3 active, 0 inactive (max not set, min not set)
     Serial1/0/0:0, since 00:05:35, no frags rcvd
     Serial1/0/1:0, since 00:05:35, no frags rcvd
     Serial1/0/2:0, since 00:05:35, no frags rcvd
```

The interface bandwidth will represent the aggregated bandwidth of the link, and this bandwidth will be used in the OSPF cost calculation.

```
Router1# show interface multilink 1
Multilink1 is up, line protocol is up
  Hardware is multilink group interface
  Internet address is 192.168.254.1/24
  MTU 1500 bytes, BW 5952 Kbit, DLY 100000 usec,
     reliability 255/255, txload 3/255, rxload 3/255
  Encapsulation PPP, loopback not set
  Keepalive set (10 sec)
  DTR is pulsed for 2 seconds on reset
  LCP Open, multilink Open
  Open: IPCP
  Last input 00:00:00, output never, output hang never
  Last clearing of "show interface" counters 00:06:39
  Input queue: 1/75/0/0 (size/max/drops/flushes); Total output drops: 0
  Queueing strategy: fifo
  Output queue :0/40 (size/max)
  5 minute input rate 241000 bits/sec, 28 packets/sec
  5 minute output rate 241000 bits/sec, 28 packets/sec
     6525 packets input, 9810620 bytes, 0 no buffer
     Received 0 broadcasts, 0 runts, 0 giants, 0 throttles
     0 input errors, 0 CRC, 0 frame, 0 overrun, 0 ignored, 0 abort
     6526 packets output, 9796112 bytes, 0 underruns
     0 output errors, 0 collisions, 0 interface resets
     0 output buffer failures, 0 output buffers swapped out
     0 carrier transitions
```

The output of **show ip ospf interface** shows the current OSPF cost, which is 16.

```
Router1# show ip ospf interface multilink 1
      Multilink1 is up, line protocol is up
        Internet Address 192.158.254.13/24, Area 1
       Process ID 1, Router ID 1.1.1.1, Network Type POINT_TO_POINT, Cost:16
       Transmit Delay is 1 sec, State POINT_TO_POINT,
       Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
        Hello due in 00:00:02
       Index 1/2, flood queue length 0
       Next 0x0(0)/0x0(0)
       Last flood scan length is 0, maximum is 1
       Last flood scan time is 0 msec, maximum is 0 msec
       Neighbor Count is 0, Adjacent neighbor count is 0
       Suppress hello for 0 neighbor(s)
```

Now a link goes down and we can see that in the log:

```
Router1# show log | include down

%LINK-3-UPDOWN: Interface Serial1/0/0:0, changed state to down
%LINEPROTO-5-UPDOWN: Line protocol on Interface Serial1/0/0:0, changed state to down
```

If we check the bandwidth again it will be different than what we saw previously. Now it is showing 3968 and the bundle has only two interfaces instead of three because one interface went down. Note below the interface is still up:

```
Router1# show ppp multilink
 Multilink1, bundle name is Router2
   Bundle up for 00:05:35
   Bundle is Distributed
   0 lost fragments, 0 reordered, 0 unassigned
   0 discarded, 0 lost received, 3/255 load
   0x1226 received sequence, 0x1226 sent sequence
```

```
    Member links: 2 active, 1 inactive (max not set, min not set)
      Serial1/0/1:0, since 00:05:35, no frags rcvd
      Serial1/0/2:0, since 00:05:35, no frags rcvd
      Serial1/0/0:0 (inactive)
```

Also, the PPP multilink is still showing up, but the OSPF cost is now changed to 25 since one link is down

```
Router1# show ip ospf interface multilink 1
      Multilink1 is up, line protocol is up
        Internet Address 192.158.254.13/24, Area 1
       Process ID 1, Router ID 1.1.1.1, Network Type POINT_TO_POINT, Cost:25
       Transmit Delay is 1 sec, State POINT_TO_POINT,
       Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
        Hello due in 00:00:02
       Index 1/2, flood queue length 0
       Next 0x0(0)/0x0(0)
       Last flood scan length is 0, maximum is 1
       Last flood scan time is 0 msec, maximum is 0 msec
       Neighbor Count is 0, Adjacent neighbor count is 0
       Suppress hello for 0 neighbor(s)
```

This what will trigger SPF calculation, and OSPF will bring up the demand circuit. If the link keeps flapping then we may see the demand circuit keep flapping because the cost will be changed every time a link adds up or gets deleted from the the multilink PPP bundle.

Solution

The PPP multilink is supported in OSPF, but as long as all the link within the bundle stays up, the demand circuit will be stable. As soon as a link goes down, even though there is no IP address associated with it, it will affect the OSPF cost calculation, and because of that, OSPF will run the SPF to recalculate the best paths. To solve this problem, the only solution is to configure the OSPF cost manually with the following command.

Router 1	Router 2
interface Multilink1 ip address 192.158.254.1 255.255.255.0 no cdp enable **ip ospf cost 10** ppp multilink no ppp multilink fragmentation multilink-group 1 ! router ospf 20 network 192.158.254.0 0.0.0.255 area 1	interface Multilink1 ip address 192.158.254.2 255.255.255.0 no cdp enable **ip ospf cost 10** ppp multilink no ppp multilink fragmentation multilink-group 1 ! router ospf 20 network 192.158.254.0 0.0.0.255 area 1

This command will make sure that every time there is a link that is added or deleted in the multilink PPP bundle, OSPF cost will not be affected. This will stabilizes the OSPF demand circuit over the PPP multilink.

Related Information

- **OSPF Support Page**
- **Technical Support − Cisco Systems**

Contacts & Feedback | Help | Site Map
© 2014 − 2015 Cisco Systems, Inc. All rights reserved. Terms & Conditions | Privacy Statement | Cookie Policy | Trademarks of Cisco Systems, Inc.