IS-IS Hello Padding Behavior

TAC **Document ID: 119399**

Contributed by Luc De Ghein, Cisco TAC Engineer.
Dec 15, 2015

Contents

Introduction

This document describes the behavior of Integrated Intermediate System-to-Intermediate System (IS-IS) Hello packet padding in the Cisco IOS®.

Background Information

The IS-IS by default pads the Hello packets to the full interface Maximum Transmission Unit (MTU). This is in order to detect MTU mismatches. The MTU on either side of the link should match. The padding can also be used in order to detect the real MTU value of the technology that lies beneath. For example, for Layer 2 (L2) transport over Multi Protocol Label Switching (MPLS) scenarios, the MTU of the transport technology might be much lower than the MTU on the edge. For example, the MTU can be 9,000 bytes on the edge, while the MPLS transport technology has an MTU of 1,500 bytes.

If the MTU values match on either side, then the padding can be disabled. As such, unnecessary usage of bandwidth and buffers by IS-IS Hello packets can be avoided. The router command that is used in order to disable the Hello padding is **no hello padding [multi-point|point-to-point]**. The interface command that is used in order to disable the Hello padding is **no isis hello padding**.

If the padding is disabled at the start, the router still sends Hello packets at full MTU. In order to avoid this, disable the padding with the interface command and use the *always* keyword. In this case, all of the IS-IS Hello packets are not padded.

Note: Cisco recommends that you do not disable the IS-IS Hello padding in order to ensure that two routers form an IS-IS adjacency on a link that has mismatched MTU values on either side.

Padding TLVs

The IS-IS Hello packets have a padding Type Length Value (TLV). For a Point-to-Point (P2P) IH, the TLV for the padding is 8. For the LAN IIH, the TLV for the padding is 8.

Padding TLV Example

The example that is provided in the next image is used in this section in order to explain MTU and the disablement of Hello padding in IS-IS:

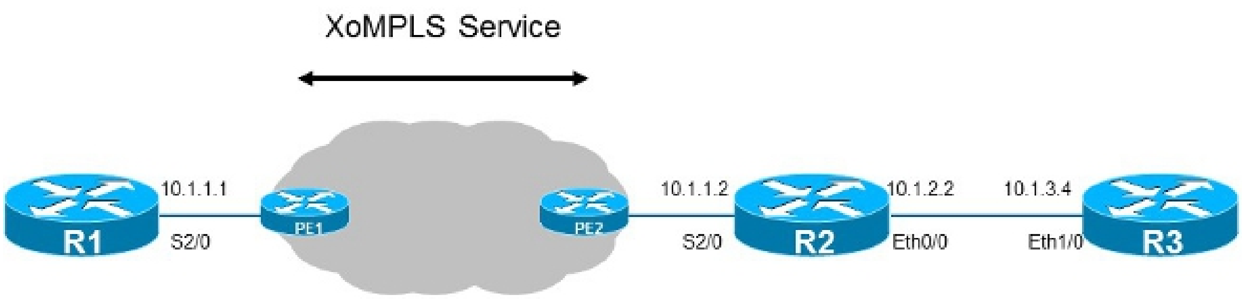

In this example, PE1 and PE2 have set up a Virtual Circuit (VC) 100 between them in order to connect the routers R1 and R2 at L2. This VC is an Ethernet over MPLS (EoMPLS) VC.

```
PE1#show xconnect all
Legend:   XC ST=Xconnect State S1=Segment1 State S2=Segment2 State
UP=Up       DN=Down           AD=Admin Down     IA=Inactive
SB=Standby HS=Hot Standby     RV=Recovering     NH=No Hardware

XC ST Segment 1                       S1 Segment 2                         S2
------+--------------------------------+--+--------------------------------+--
UP pri   ac Se2/0(HDLC)               UP mpls 10.100.1.5:100              UP
```

```
PE1#show mpls l2transport vc 100

Local intf     Local circuit              Dest address    VC ID      Status
------------- -------------------------- --------------- ---------- ----------
Se2/0          HDLC                       10.100.1.5      100        UP
```

Here is the output for the router R1:

```
interface Serial2/0
ip address 10.1.1.1 255.255.255.0
ip router isis 1
serial restart-delay 0
```

Here is the output for the router R2:

```
interface Serial2/0
ip address 10.1.1.2 255.255.255.0
ip router isis 1
serial restart-delay 0
```

The output of the **debug isis adj-packets** debug command provides information about the IS-IS adjacency:

```
R1#debug isis adj-packets
IS-IS Adjacency related packets debugging is on for router process 1

R1#
13:00:59.978: ISIS-Adj: Sending serial IIH on Serial2/0, 3way state:DOWN, length 1499
13:01:07.758: ISIS-Adj: Sending serial IIH on Serial2/0, 3way state:DOWN, length 1499
13:01:16.280: ISIS-Adj: Sending serial IIH on Serial2/0, 3way state:DOWN, length 1499
```

```
R2#
13:01:50.100: ISIS-Adj: Sending serial IIH on Serial2/0, 3way state:DOWN, length 1499
13:02:00.062: ISIS-Adj: Sending serial IIH on Serial2/0, 3way state:DOWN, length 1499
13:02:07.899: ISIS-Adj: Sending serial IIH on Serial2/0, 3way state:DOWN, length 1499
```

In this scenario, the IS-IS adjacency fails.

```
R1#show isis neighbors

Tag 1:
System Id      Type Interface    IP Address      State Holdtime Circuit Id
R1#
```

```
R1#show clns interface Serial 2/0
Serial2/0 is up, line protocol is up
Checksums enabled, MTU 1500, Encapsulation HDLC
ERPDUs enabled, min. interval 10 msec.
CLNS fast switching enabled
CLNS SSE switching disabled
DEC compatibility mode OFF for this interface
Next ESH/ISH in 18 seconds
Routing Protocol: IS-IS
   Circuit Type: level-1-2
   Interface number 0x1, local circuit ID 0x101
   Level-1 Metric: 10, Priority: 64, Circuit ID: R1.01
   Level-1 IPv6 Metric: 10
   Number of active level-1 adjacencies: 0
   Next IS-IS Hello in 5 seconds
   if state DOWN
```

The MTU on the serial interfaces for the routers R1 and R2 are the default 1,500 bytes.

The IS-IS adjacency fails because the IS-IS Hello packets are 1,499 bytes in size. The MPLS network only allows 1,500-byte packets, minus eight bytes (two MPLS labels for the MPLS service), which equals 1,492 bytes (the packet size that is allowed to pass through). For transport of L2 over MPLS, the size of the L2 header must be subtracted from the 1,492 bytes that result as well.

No Hello Padding

In this scenario, the **no isis hello padding** command is used on the interface of Serial2/0 on the router R1:

```
interface Serial2/0
ip address 10.1.1.1 255.255.255.0
ip router isis 1
serial restart-delay 0
no isis hello padding
```

```
R1#
13:03:46.712: ISIS-Adj: Sending serial IIH on Serial2/0, 3way state:DOWN, length 1499
13:03:54.717: ISIS-Adj: Sending serial IIH on Serial2/0, 3way state:DOWN, length 1499
13:04:03.057: ISIS-Adj: Sending serial IIH on Serial2/0, 3way state:DOWN, length 1499
13:04:11.538: ISIS-Adj: Sending serial IIH on Serial2/0, 3way state:DOWN, length 1499
13:04:21.301: ISIS-Adj: Sending serial IIH on Serial2/0, 3way state:DOWN, length 1499
13:04:30.636: ISIS-Adj: Sending serial IIH on Serial2/0, 3way state:DOWN, length 1499
13:04:39.958: ISIS-Adj: Sending serial IIH on Serial2/0, 3way state:DOWN, length 1499
```

As shown, more than five IS-IS Hello packets are sent with full MTU size (1,497 bytes). The router continues to send the Hello packets with padding until the IS-IS adjacency comes up. However, unless the MTU issue is fixed, the adjacency does not come up.

The MTU is lowered to 1,400 bytes on the interface Serial2/0 on the router R1. Thus, the packets that are up to 1,400 bytes in size can surely pass through the MPLS network over the pseudo-wire.

Here is the output for the router R1:

```
!
interface Serial2/0
mtu 1400
ip address 10.1.1.1 255.255.255.0
ip router isis 1
serial restart-delay 0
no isis hello padding

R1#
13:07:19.428: ISIS-Adj: Sending serial IIH on Serial2/0, 3way state:DOWN, length 1399
13:07:29.024: ISIS-Adj: Sending serial IIH on Serial2/0, 3way state:DOWN, length 1399
13:07:38.185: ISIS-Adj: Sending serial IIH on Serial2/0, 3way state:DOWN, length 1399
13:07:45.715: ISIS-Adj: Sending serial IIH on Serial2/0, 3way state:DOWN, length 1399
13:07:55.351: ISIS-Adj: Sending serial IIH on Serial2/0, 3way state:DOWN, length 1399
13:08:04.814: ISIS-Adj: Sending serial IIH on Serial2/0, 3way state:DOWN, length 1399
13:08:14.216: ISIS-Adj: Sending serial IIH on Serial2/0, 3way state:DOWN, length 1399
13:08:23.447: ISIS-Adj: Sending serial IIH on Serial2/0, 3way state:DOWN, length 1399
13:08:31.676: ISIS-Adj: Sending serial IIH on Serial2/0, 3way state:DOWN, length 1399
13:08:39.966: ISIS-Adj: Sending serial IIH on Serial2/0, 3way state:DOWN, length 1399
```

The router R1 continues to transmit the Hello packets with padding. The size is now 1,400 bytes minus one.

Once the MTU is lowered on the interface Serial 2/0 on the router R2, the padding is disabled.

Here is the output for the router R2:

```
interface Serial2/0
mtu 1400
ip address 10.1.1.2 255.255.255.0
ip router isis 1
serial restart-delay 0
```

Once the router R1 sees the IS-IS Hello packet arrive from the router R2, it brings up the IS-IS adjacency. Because the router R2 also sees the IS-IS Hello packets from the router R1, eventually the IS-IS adjacency moves to the *UP* state, which means that a three-way adjacency is created. At this point, the router R1 (with Hello padding disabled on the interface Serial 2/0) lowers the size of the Hello packet to the minimum.

```
R1#
13:08:47.010: ISIS-Adj: Rec serial IIH from *HDLC* (Serial2/0), cir type L1, cir id 01,
 length 1399
13:08:47.010: ISIS-Adj: newstate:1, state_changed:1, going_up:0, going_down:0
13:08:47.010: ISIS-Adj: Action = GOING UP, new type = L1
13:08:47.010: ISIS-Adj: New serial adjacency
13:08:47.010: ISIS-Adj: rcvd state INIT, old state DOWN, new state INIT, nbr usable TRUE
13:08:47.011: ISIS-Adj: Sending serial IIH on Serial2/0, 3way state:INIT, length 1399
13:08:47.055: ISIS-Adj: Rec serial IIH from *HDLC* (Serial2/0), cir type L1, cir id 01,
 length 1399
13:08:47.055: ISIS-Adj: rcvd state UP, old state INIT, new state UP, nbr usable TRUE
13:08:47.056: ISIS-Adj: newstate:0, state_changed:1, going_up:1, going_down:0
13:08:47.056: ISIS-Adj: Action = GOING UP, new type = L1
13:08:47.056: ISIS-Adj: L1 adj count 1
13:08:47.056: ISIS-Adj: Sending serial IIH on Serial2/0, 3way state:UP, length 43
```

As shown, the router R1 sends an IS-IS Hello packet with **length 43** and receives the Hello packets from the router R2 with **length 1399**. This is because the Hello padding is still active on the router R2.

In this example, the IS-IS adjacency does not come up if either side of the link still has the MTU set to 1,500 bytes on the interface Serial 2/0. This is the case even when the **no isis hello padding** command is enabled. The interface only comes up after the MTU is set to the correct value on either side of the link.

Thus, if you only disable the IS-IS Hello padding, it is not enough to bring up the IS-IS adjacency. The MTU must be low enough so that the MTU-sized IS-IS Hello packets are sent and received properly by the routers on either side of the link.

No Hello Padding Always

With the MTU set to 1,500 bytes on the interface Serial2/0 on the router R1, the adjacency does not come up because the transmitted IS-IS Hello packets are still the full MTU size. In order to work around this issue, you can configure the **no isis hello padding always** interface command on the interface Serial2/0 in order to disable the padding always.

```
!
interface Serial2/0
ip address 10.1.1.1 255.255.255.0
ip router isis 1
serial restart-delay 0
no isis hello padding always
```

As soon as this command is configured, the IS-IS Hello packets have the minimum size. The IS-IS adjacency between the routers R1 and R2 immediately comes up.

```
R1#
13:25:47.284: ISIS-Adj: Sending serial IIH on Serial2/0, 3way state:INIT,
 length 43, never pad
13:25:47.328: ISIS-Adj: Rec serial IIH from *HDLC* (Serial2/0), cir type L1,
 cir id 01, length 1399
13:25:47.328: ISIS-Adj: rcvd state INIT, old state INIT, new state UP,
 nbr usable TRUE
13:25:47.328: ISIS-Adj: newstate:0, state_changed:1, going_up:1, going_down:0
13:25:47.328: ISIS-Adj: Action = GOING UP, new type = L1
13:25:47.329: ISIS-Adj: L1 adj count 1
13:25:47.330: ISIS-Adj: Sending serial IIH on Serial2/0, 3way state:UP,
 length 43, never pad
13:25:47.374: ISIS-Adj: Rec serial IIH from *HDLC* (Serial2/0), cir type L1,
 cir id 01, length 1399
13:25:47.374: ISIS-Adj: rcvd state UP, old state UP, new state UP,
 nbr usable TRUE
13:25:47.375: ISIS-Adj: newstate:0, state_changed:0, going_up:0, going_down:0
13:25:47.375: ISIS-Adj: Action = ACCEPT
13:25:47.375: ISIS-Adj: ACTION_ACCEPT:
```

The Problem with IS-IS and Interface MTU

If the interface MTU is mismatched, then the IS-IS adjacency does not come up. For a quick fix, you can disable the IS-IS Hello padding with the *always* keyword. However, this might not be a real fix.

Here is the output for the router R1:

```
interface Serial2/0
ip address 10.1.1.1 255.255.255.0
ip router isis 1
serial restart-delay 0
no isis hello padding always
```

The IS-IS adjacency is up.

```
R1#show isis neighbors

Tag 1:
System Id      Type Interface   IP Address      State Holdtime Circuit Id
R2              L1   Se2/0       10.1.1.2        UP    22       01
```

Here is a ping that is sent from the router R1 to the router R3 in order to check the traffic that crosses the link:

```
R1#ping 10.100.1.3 source 10.100.1.1 size 1400 repeat 1
Type escape sequence to abort.
Sending 1, 1400-byte ICMP Echos to 10.100.1.3, timeout is 2 seconds:
Packet sent with a source address of 10.100.1.1
!
Success rate is 100 percent (1/1), round-trip min/avg/max = 44/44/44 ms

R1#ping 10.100.1.3 source 10.100.1.1 size 1500 repeat 1
Type escape sequence to abort.
Sending 1, 1500-byte ICMP Echos to 10.100.1.3, timeout is 2 seconds:
Packet sent with a source address of 10.100.1.1
.
Success rate is 0 percent (0/1)
```

As shown, the packets with a size of 1,500 bytes do not make it through. This is because the router R1 believes that the MTU is 1,500 bytes on the interface Serial2/0:

```
R1#show interfaces Serial2/0
Serial2/0 is up, line protocol is up
Hardware is M4T
Internet address is 10.1.1.1/24
MTU 1500 bytes, BW 1544 Kbit/sec, DLY 20000 usec,
     reliability 255/255, txload 1/255, rxload 1/255
Encapsulation HDLC, crc 16, loopback not set
Keepalive set (10 sec)
Restart-Delay is 0 secs
Last input 00:00:01, output 00:00:01, output hang never
Last clearing of "show interface" counters never
Input queue: 0/75/0/0 (size/max/drops/flushes); Total output drops: 0
Queueing strategy: weighted fair
Output queue: 0/1000/64/0 (size/max total/threshold/drops)
     Conversations 0/1/256 (active/max active/max total)
     Reserved Conversations 0/0 (allocated/max allocated)
     Available Bandwidth 1158 kilobits/sec
5 minute input rate 0 bits/sec, 0 packets/sec
5 minute output rate 0 bits/sec, 0 packets/sec
     590 packets input, 283131 bytes, 0 no buffer
     Received 567 broadcasts (0 IP multicasts)
     0 runts, 0 giants, 0 throttles
     0 input errors, 0 CRC, 0 frame, 0 overrun, 0 ignored, 0 abort
     693 packets output, 313789 bytes, 0 underruns
     0 output errors, 0 collisions, 2 interface resets
     0 unknown protocol drops
     0 output buffer failures, 0 output buffers swapped out
     3 carrier transitions     DCD=up DSR=up DTR=up RTS=up CTS=up
```

If the MTU is lowered to 1,400 bytes on the interface Serial2/0, then the router R1 can fragment the packets if the packets do not have the Do Not fragment (DF) Bit set. If the packets have the DF bit set, then the router can send back an ICMP 3/4 message, which is used by the Path MTU Discovery. This allows the sender of the packets to lower the size of the packets that it sends out. The correct setting of the MTU is important for the traffic that traverses the router, but also for the traffic that originates from the router and crosses that link. An example of the latter is Border Gateway Protocol (BGP), which uses TCP and can use the Path MTU

Discovery.

IS-IS Flooding

In order to fix the IS-IS adjacency issue, the operator of the network can disable the Hello padding with the *always* keyword. The MTU of the serial link is left at 1, 500 bytes.

There is still the issue of the IS-IS flooding. When the IS-IS database is small, there is no issue.

```
R1#debug isis update-packets
IS-IS Update related packet debugging is on for router process 1
```

When the router R3 adds a prefix and floods this, the router R1 receives the router R3 Link State PDU (LSP) from the router R2.

```
R1#
*Nov 19 13:53:58.227: ISIS-Upd: Rec L1 LSP 0000.0000.0003.00-00, seq B, ht 1197,
*Nov 19 13:53:58.227: ISIS-Upd: from SNPA *HDLC* (Serial2/0)
*Nov 19 13:53:58.227: ISIS-Upd: LSP newer than database copy
*Nov 19 13:53:58.227: ISIS-Upd: TLV contents different, code 130
*Nov 19 13:53:58.228: ISIS-Upd: TID 0 leaf routes changed
```

When the number of prefixes that are advertised by the router R3 increases, the LSP of the router R3 is so large that it is split into several fragments:

```
R3#show isis database

Tag 1:
IS-IS Level-1 Link State Database:
LSPID                 LSP Seq Num LSP Checksum LSP Holdtime      ATT/P/OL
R1.00-00                0x0000000C   0x5931       1137              0/0/0
R2.00-00                0x0000000B   0xCB7D       1162              0/0/0
R3.00-00              * 0x0000000D   0xF637       1104              0/0/0
R3.00-01              * 0x00000001   0x6AD8       1104              0/0/0
R3.00-02              * 0x00000001   0xB58A       1104              0/0/0
R3.01-00              * 0x00000002   0x9BB1       387               0/0/0

Tag null:
```

The **R3.00-00** is the first fragment, the **R3.00-01** is the second fragment, and so on.

```
R2#
14:22:15.584: ISIS-Upd: Retransmitting L1 LSP 0000.0000.0003.00-00 on Serial2/0
14:22:15.624: ISIS-Upd: Sending L1 LSP 0000.0000.0003.00-00, seq E, ht 467 on
 Serial2/0
14:22:18.352: ISIS-Snp: Rec L1 CSNP from 0000.0000.0003 (Ethernet1/0)
14:22:20.625: ISIS-Upd: Retransmitting L1 LSP 0000.0000.0003.00-00 on Serial2/0
14:22:20.657: ISIS-Upd: Sending L1 LSP 0000.0000.0003.00-00, seq E, ht 462 on
 Serial2/0
```

This is the LSP that is retransmitted by the router R2 over the interface Serial2/0. The PDU length is 1,490 bytes, so the size of this packet does not allow it to reach the router R1.

```
▷ Frame 9 (1495 bytes on wire, 1495 bytes captured)
▽ Cisco HDLC
    Address: Multicast (0x8f)
    Protocol: OSI (0xfefe)
    CLNS Padding: 0x03
▽ ISO 10589 ISIS InTRA Domain Routeing Information Exchange Protocol
    Intra Domain Routing Protocol Discriminator: ISIS (0x83)
    PDU Header Length  : 27
    Version (==1)      : 1
    System ID Length   : 0
    PDU Type           : L1 LSP (R:000)
    Version2 (==1)     : 1
    Reserved (==0)     : 0
    Max.AREAs: (0==3)  : 0
  ▽ ISO 10589 ISIS Link State Protocol Data Unit
      PDU length: 1490
      Remaining lifetime: 754
      LSP-ID: 0000.0000.0003.00-00
      Sequence number: 0x0000000e
    ▷ Checksum: 0xf438 [correct]
    ▷ Type block(0x03): Partition Repair:0, Attached bits:0, Overload bit:0, IS type:3
    ▷ Area address(es) (4)
    ▷ Protocols supported (1)
    ▷ Hostname (2)
    ▷ IP Interface address(es) (4)
    ▷ IP Internal reachability (24)
    ▷ IS Reachability (12)
    ▷ IP External reachability (252)
    ▷ IP External reachability (252)
    ▷ IP External reachability (252)
    ▷ IP External reachability (252)
    ▷ IP External reachability (252)
    ▷ IP External reachability (132)
```

While the IS-IS adjacency between the routers R1 and R2 is active, the router R1 has less IP prefixes in its routing table:

```
R1#show isis neighbors

Tag 1:
System Id      Type Interface   IP Address      State Holdtime Circuit Id
R2              L1   Se2/0       10.1.1.2        UP    25       01

R2#show isis neighbors

Tag 1:
System Id      Type Interface   IP Address      State Holdtime Circuit Id
R1              L1   Se2/0       10.1.1.1        UP    26       01
R3              L1   Et1/0       10.1.2.3        UP    8        R3.01

R2#show ip route summary
IP routing table name is default (0x0)
IP routing table maximum-paths is 32
Route Source    Networks    Subnets     Replicates Overhead    Memory (bytes)
connected        0           5           0          360         900
static           0           0           0          0           0
application      0           0           0          0           0
isis 1           0           252         0          18144       45360
Level 1: 252 Level 2: 0 Inter-area: 0
internal         1                                              10620
```

```
Total           1           257         0           18504       56880
```

```
R1#show ip route summary
IP routing table name is default (0x0)
IP routing table maximum-paths is 32
Route Source    Networks    Subnets     Replicates Overhead    Memory (bytes)
connected         0           3           0           216         540
static            0           0           0           0           0
application       0           0           0           0           0
isis 1            0           2           0           144         360
Level 1: 2 Level 2: 0 Inter-area: 0
internal           1                                              560
Total           1           5           0           360         1460
```

This is because the LSP R3.00-00 from the router R3 does not reach the router R1.

```
R3#show isis database

Tag 1:
IS-IS Level-1 Link State Database:
LSPID                 LSP Seq Num LSP Checksum LSP Holdtime     ATT/P/OL
R1.00-00              0x0000000E   0x5533       1009             0/0/0
R2.00-00              0x0000000C   0xC97E       453              0/0/0
R3.00-00            * 0x0000000F   0xF239       1045             0/0/0
R3.00-01            * 0x00000003   0x66DA       1098             0/0/0
R3.00-02            * 0x00000003   0xB18C       1060             0/0/0
R3.01-00            * 0x00000004   0x97B3       554              0/0/0

Tag null:
```

```
R1#show isis database

Tag 1:
IS-IS Level-1 Link State Database
LSPID                 LSP Seq Num LSP Checksum LSP Holdtime     ATT/P/OL
R1.00-00            * 0x0000000E   0x5533       1008             0/0/0
R2.00-00              0x0000000C   0xC97E       449              0/0/0
R3.00-01              0x00000002   0x68D9       223              0/0/0
R3.00-02              0x00000002   0xB38B       246              0/0/0
R3.01-00              0x00000004   0x97B3       545              0/0/0
```

The router R1 does not have the first fragment of the L1 LSP (R3.00-00) of the router R3. This first fragment is the largest and holds the most prefixes in this case. For this reason, the router R1 does not have some of the prefixes, which causes black-holing of the traffic.

In order to resolve this issue, you can lower the LSP MTU via the **lsp-mtu <128-4352>** router IS-IS command. If you configure this command only at the router R2, then the router R2 does not change the LSPs that are received from the router R3 in any way. This means that if the router R2 receives an LSP with a size of 1,490 bytes, then the router R2 does not fragment it. If you configure the **lsp-mtu 1400** command on the router R3, then the router R3 creates smaller LSPs, which are small enough to cross the link between the routers R2 and R1.

The PDU length is now 1,394 bytes if you configure the **lsp-mtu 1400** command on the router R3:

```
Frame 9 (1399 bytes on wire, 1399 bytes captured)
Cisco HDLC
    Address: Multicast (0x8f)
    Protocol: OSI (0xfefe)
    CLNS Padding: 0x03
ISO 10589 ISIS InTRA Domain Routeing Information Exchange Protocol
    Intra Domain Routing Protocol Discriminator: ISIS (0x83)
    PDU Header Length  : 27
    Version (==1)      : 1
    System ID Length   : 0
    PDU Type           : L1 LSP (R:000)
    Version2 (==1)     : 1
    Reserved (==0)     : 0
    Max.AREAs: (0==3)  : 0
  ISO 10589 ISIS Link State Protocol Data Unit
      PDU length: 1394
      Remaining lifetime: 1197
      LSP-ID: 0000.0000.0003.00-00
      Sequence number: 0x00000012
    Checksum: 0xb7e0 [correct]
    Type block(0x03): Partition Repair:0, Attached bits:0, Overload bit:0, IS type:3
    Area address(es) (4)
    Protocols supported (1)
    Hostname (2)
    IP Interface address(es) (4)
    IP Internal reachability (24)
    IS Reachability (12)
    IP External reachability (252)
    IP External reachability (252)
    IP External reachability (252)
    IP External reachability (252)
    IP External reachability (252)
    IP External reachability (36)
```

In conclusion, if you have one link with a smaller MTU and use the **no isis hello padding always** command, it can lead to traffic flooding and black-holing. In order to resolve the flooding issue, you can lower the maximum size of the LSPs, but you must also configure the **lsp-mtu** router IS-IS command on every IS-IS router.

Changes to the MTU

This section describes the effects of changes that are made to the underlying MTU.

Hello Padding Enabled

In this scenario, the network functions properly from the start. The MTU is set to 1,400 bytes on the interface Serial2/0 on the routers R1 and R2. The IS-IS Hello padding is enabled, which is the default behavior.

Here is the output for the router R1:

```
interface Serial2/0
mtu 1400
ip address 10.1.1.1 255.255.255.0
ip router isis 1
serial restart-delay 0
```

Here is the output for the router R2:

```
interface Serial2/0
mtu 1400
ip address 10.1.1.2 255.255.255.0
ip router isis 1
serial restart-delay 0
```

```
R1#show isis neighbors

Tag 1:
System Id      Type Interface   IP Address      State Holdtime Circuit Id
R2              L1   Se2/0       10.1.1.2        UP    23       01
```

```
R2#show isis neighbors

Tag 1:
System Id      Type Interface   IP Address      State Holdtime Circuit Id
R1              L1   Se2/0       10.1.1.1        UP    27       01
0000.0000.0003 L1   Et1/0       10.1.2.3        UP    7       0000.0000.0003.01
```

The IS-IS adjacency across the serial is up, and the IS-IS flooding is fine.

At a certain point in time, an issue occurs in the MPLS service provider network that causes the end-to-end MTU between the PE1 and PE2 to drop below 1,400 bytes.

Because the Hello padding is enabled (the default behavior), the IS-IS adjacency quickly goes down on the interface Serial2/0. This indicates that there is an issue across the MPLS cloud. Because the IS-IS adjacency goes down, the routing no longer points to this MPLS cloud, and no traffic is black-holed across it.

Hello Padding Disabled

In this scenario, the network functions properly from the start. The MTU is set to 1,400 bytes on the interface Serial2/0 on the routers R1 and R2. The IS-IS Hello padding is disabled.

Here is the output for the router R1:

```
!
interface Serial2/0
mtu 1400
ip address 10.1.1.1 255.255.255.0
ip router isis 1
serial restart-delay 0
no isis hello padding
```

Here is the output for the router R2:

```
!
interface Serial2/0
mtu 1400
ip address 10.1.1.2 255.255.255.0
ip router isis 1
serial restart-delay 0
no isis hello padding
```

The IS-IS adjacency across the serial is up, and the IS-IS flooding is fine.

This is the database of the router R1:

```
R1#show isis database
```

```
Tag 1:
IS-IS Level-1 Link State Database:
LSPID                 LSP Seq Num LSP Checksum LSP Holdtime     ATT/P/OL
R1.00-00            * 0x0000001D   0x3742       1148              0/0/0
R2.00-00              0x0000001D   0xA78F       1161              0/0/0
R3.00-00              0x00000016   0xAFE4       454               0/0/0
R3.00-01              0x0000000B   0x0A0B       393               0/0/0
R3.00-02              0x0000000B   0xC2A5       451               0/0/0
R3.01-00              0x00000009   0x8DB8       435               0/0/0
```

At a certain point in time, an issue occurs in the MPLS service provider network that causes the end-to-end MTU between the PE1 and PE2 to drop below 1,400 bytes.

The IS-IS is not affected immediately, but the IP traffic might be. If there is traffic with packets that are 1,400 bytes in size, they are dropped in the MPLS network.

If the network is stable, there is no flooding for a large amount of time. This remains as long as the LSP refresh time. Once it is time to refresh the LSP(s), the flooding is broken across the MPLS network.

```
R2#
15:27:07.848: ISIS-Upd: Retransmitting L1 LSP 0000.0000.0003.00-01 on Serial2/0
15:27:07.880: ISIS-Upd: Sending L1 LSP 0000.0000.0003.00-01, seq C, ht 1147 on
 Serial2/0
15:27:12.883: ISIS-Upd: Retransmitting L1 LSP 0000.0000.0003.00-01 on Serial2/0
15:27:12.924: ISIS-Upd: Sending L1 LSP 0000.0000.0003.00-01, seq C, ht 1142 on
 Serial2/0
```

This is the IS-IS database of the router R1 after the issue occurs in the MPLS network:

```
R1#show isis database

Tag 1:
IS-IS Level-1 Link State Database:
LSPID                 LSP Seq Num LSP Checksum LSP Holdtime     ATT/P/OL
R1.00-00            * 0x0000001D   0x3742       725               0/0/0
R2.00-00              0x0000001D   0xA78F       737               0/0/0
R3.00-00              0x00000016   0xAFE4       30                0/0/0
R3.00-01              0x0000000B   0xCE1F       0 (30)            0/0/0
R3.00-02              0x0000000C   0xC0A6       895               0/0/0
R3.01-00              0x0000000A   0x8BB9       906               0/0/0
```

This is the database after the holdtime has expired for some of the LSP fragments from the router R3:

```
R1#show isis database

Tag 1:
IS-IS Level-1 Link State Database:
LSPID                LSP Seq Num LSP Checksum LSP Holdtime     ATT/P/OL
R1.00-00            * 0x0000001D   0x3742       605               0/0/0
R2.00-00              0x0000001D   0xA78F       618               0/0/0
R3.00-02              0x0000000C   0xC0A6       775               0/0/0
R3.01-00              0x0000000A   0x8BB9       787               0/0/0
```

The fragments R3.00-00 and R3.00-01 no longer appear on the router R1, and the routes from the router R3 are no longer on the router R1:

```
R1#show ip route summary
IP routing table name is default (0x0)
IP routing table maximum-paths is 32
Route Source    Networks    Subnets     Replicates Overhead    Memory (bytes)
connected        0           3           0            216         540
```

```
static          0           0           0           0           0
application     0           0           0           0           0
isis 1          0           2           0           144         360
Level 1: 2 Level 2: 0 Inter-area: 0
internal        1                                               560
Total           1           5           0           360         1460
```

As shown, some of the router R3 LSP fragments are timed-out and do not appear. This causes some of the routes to not appear in the routing table.

If you disable the Hello padding, it can hide a future issue in the network. When the underlying MTU changes, it can cause a routing issue that is much harder to troubleshoot because you must examine the routing table and the IS-IS database on multiple routers in order to pinpoint the issue. With the Hello padding enabled, the fact that the IS-IS adjacency goes down makes it much easier to determine the location of the issue.

Important Notes

The best solution is to set the MTU to the correct value on the links and ensure that it is equal on both sides of the links. This ensures that the IS-IS flooding works properly and that the router is able to perform fragmentation correctly or behave correctly when it assists with the Path MTU Discovery.

The issue with the IS-IS flooding might only become obvious when the LSPs become larger (when the network grows). When the IS-IS Hello padding is disabled, it fixes the issue where the IS-IS adjacencies do not come up. However, the issue of flooding, black-holing traffic, and perhaps broken Path MTU Discovery, can potentially arise much later than the time at which the IS-IS Hello padding is disabled. This makes the issue much harder to troubleshoot, which takes much more time.

Updated: Dec 15, 2015 Document ID: 119399