

## 知 CR16-F路由器与思科N9K对接ISIS邻居因为MTU不一致导致邻居无法建立问题

IS-IS 何理 2016-03-23 发表

我司CR16-F路由器与Cisco N9K路由对接ISIS时，发现ISIS邻居一直无法建立成功，两边配置完全一致；现场收集了debug信息以及思科侧信息；



通过现场debug信息查看，IIH报文收发正常，但是不断提示ISIS hello超时，如此反复不断；

```
*Mar 2 07:21:40:899 2016 H3C ISIS/7/ISISDBG:
```

```
ISIS-1-TMR: Level-1 hello timer expired on the circuit GigabitEthernet0/0.
```

```
*Mar 2 07:21:40:899 2016 H3C ISIS/7/ISISDBG:
```

```
ISIS-1-ADJ: Send a Lan L1 Hello packet on circuit(GigabitEthernet0/0)
```

```
*Mar 2 07:21:40:921 2016 H3C ISIS/7/ISISDBG:
```

```
ISIS-1-ADJ: Receive a Lan L1 Hello packet from(0000.0000.0001) on circuit(GigabitEthernet0/0)
```

```
*Mar 2 07:21:40:922 2016 H3C ISIS/7/ISISDBG:
```

```
ISIS-1-ADJ:
```

```
0000: 83 1b 01 06 0f 01 00 00 01 00 00 00 00 00 01 00  
0010: 1e 00 35 40 00 00 00 00 00 01 01 01 02 01 10 84  
0020: 04 0a 00 00 01 81 01 cc d3 03 00 00 00 06 06 0e  
0030: cd 0a 5c 01 05
```

```
*Mar 2 07:21:40:922 2016 H3C ISIS/7/ISISDBG:
```

```
ISIS-1-ADJ: Level-1 NBR(0000.0000.0001) two way pass.
```

那么我们可以通过打印出来的hello报文查看一下，根据ISIS IIH报文结构，我们可以看出我们发出的ISIS报文为1541大小：

```
*Mar 2 12:20:19:597 2016 ZhangJiang-H3C-P2 ISIS/7/ISISDBG: -MDC=1;
```

```
ISIS-9812-ADJ:
```

```
0000: 83 1b 01 06 10 01 00 00 02 00 00 00 00 00 59 00  
0010: 1e 06 05 40 00 00 00 00 00 59 02 01 02 01 00 84  
// 0x605 = 1541，也就是mtu=1544
```

```
0020: 04 db e9 2b 8e 81 01 cc d3 01 00 06 06 6c 9c ed
```

```
0030: 0d ed ea 08 ff 00 00 00 00 00 00 00 00 00 00 00
```

但是我们收到对端的ISIS IIH报文长度只有1527：

接收的报文长度：

```
*Mar 2 12:20:22:397 2016 ZhangJiang-H3C-P2 ISIS/7/ISISDBG: -MDC=1;
```

ISIS-9812-ADJ: Receive a Lan L2 Hello packet from(0000.0002.0530) on circuit(TenGigabitEthernet2/0/4)

\*Mar 2 12:20:22:399 2016 ZhangJiang-H3C-P2 ISIS/7/ISISDBG: -MDC=1;

ISIS-9812-ADJ:

0000: 83 1b 01 00 10 01 00 00 02 00 00 00 02 05 30 00

0010: 1e 05 f7 40 00 00 00 02 05 30 3d 81 01 cc d3 03

//但接收的报文长度为, 0x5f7 = 1527

0020: 00 00 00 01 02 01 00 84 04 db e9 2b 8d 06 00 08

此时查看Cisco侧的ISIS信息:

#show isis interface tenGigE 0/0/0/0

Wed Mar 2 17:20:39.287 CST

TenGigE0/0/0/0 Enabled

Adjacency Formation: Enabled

Prefix Advertisement: Enabled

IPv4 BFD: Disabled

IPv6 BFD: Disabled

BFD Min Interval: 150

BFD Multiplier: 3

Circuit Type: level-2-only (Interface circuit type is level-1-2)

Media Type: LAN

Circuit Number: 37

Level-2

Adjacency Count: 1

LAN ID:

Priority (Local/DIS): 64/64

Next LAN IIH in: 226 ms

LSP Pacing Interval: 33 ms

PSNP Entry Queue Size: 0

CLNS I/O

Protocol State: Up

MTU: 1527 //发现对端只有1527

SNPA: 6c9c.ed0d.eda8

Layer-2 MCast Groups Membership:

All Level-2 ISs: Yes

IPv4 Unicast Topology: Enabled

Adjacency Formation: Running

Prefix Advertisement: Running

Metric (L1/L2): 10/10

MPLS LDP Sync (L1/L2): Enabled/Enabled

Sync Status: Achieved

IPv4 Address Family: Enabled

Protocol State: Up

Forwarding Address(es): 219.233.13.249

Global Prefix(es): 219.233.13.248/30

LSP transmit timer expires in 0 ms

LSP transmission is idle

Can send up to 9 back-to-back LSPs in the next 288 ms

但是实际接口配置确实为1544

```
#sh int tenGigE 0/0/0/0
```

Wed Mar 2 17:20:50.379 CST

TenGigE0/0/0/0 is up, line protocol is up

Interface state transitions: 27

Hardware is TenGigE, address is

Layer 1 Transport Mode is LAN

Description:

Internet address is

MTU 1544 bytes, BW 10000000 Kbit (Max: 10000000 Kbit)

reliability 254/255, txload 0/255, rxload 0/255

Encapsulation ARPA,

Full-duplex, 10000Mb/s, ZR, link type is force-up

output flow control is off, input flow control is off

loopback not set,

ARP type ARPA, ARP timeout 04:00:00

Last input 00:00:00, output 00:00:00

Last clearing of "show interface" counters never

由此我们可以推断出思科的MTU算法与我司不一致，思科的MTU计算了以太头部而我司不计算以太头；

将我司接口MTU配置为1544-14=1530后，MTU就一样了，这样ISIS邻居就可以正常建立；