

一 组网:

[REDACTED]

二 问题描述:

某保险公司LJZ分公司两台友商cisco75xx系列路由器更换为SR6604X路由器后, OSPF邻居关系不断UP、DOWN震荡, 每几分钟到20多分钟不等就会重复一次。

三 过程分析:

现场收集了SR6604X-2的诊断信息和日志文件 (SR6604X-1也有类似的问题, 分析同SR6604-2), 诊断信息经过分析未发现设备异常。从日志文件中发现如下信息:

```
%@976%Sep 15 00:02:14:089 2014 HQ-HQ-NB-7F-R-CO-SR6604-2 OSPF/5/OSPF_NBR_CHG: OSPF 100 Neighbor 10.200.123.190(Pos3/2/0) from Full to Init.
```

```
%@977%Sep 15 00:03:09:722 2014 HQ-HQ-NB-7F-R-CO-SR6604-2 OSPF/5/OSPF_NBR_CHG: OSPF 100 Neighbor 10.200.123.190(Pos3/2/0) from Loading to Full.
```

```
%@978%Sep 15 00:10:06:662 2014 HQ-HQ-NB-7F-R-CO-SR6604-2 OSPF/5/OSPF_NBR_CHG: OSPF 100 Neighbor 10.200.123.190(Pos3/2/0) from Full to Init.
```

```
%@979%Sep 15 00:11:09:710 2014 HQ-HQ-NB-7F-R-CO-SR6604-2 OSPF/5/OSPF_NBR_CHG: OSPF 100 Neighbor 10.200.123.190(Pos3/2/0) from Loading to Full.
```

```
%@980%Sep 15 00:26:55:377 2014 HQ-HQ-NB-7F-R-CO-SR6604-2 OSPF/5/OSPF_NBR_CHG: OSPF 100 Neighbor 10.200.123.190(Pos3/2/0) from Full to Init.
```

```
%@981%Sep 15 00:27:59:940 2014 HQ-HQ-NB-7F-R-CO-SR6604-2 OSPF/5/OSPF_NBR_CHG: OSPF 100 Neighbor 10.200.123.190(Pos3/2/0) from Loading to Full.
```

```
%@982%Sep 15 00:52:50:927 2014 HQ-HQ-NB-7F-R-CO-SR6604-2 OSPF/5/OSPF_NBR_CHG: OSPF 100 Neighbor 10.200.123.190(Pos3/2/0) from Full to Init.
```

```
%@983%Sep 15 00:53:49:133 2014 HQ-HQ-NB-7F-R-CO-SR6604-2 OSPF/5/OSPF_NBR_CHG: OSPF 100 Neighbor 10.200.123.190(Pos3/2/0) from Loading to Full.
```

从日志的信息可以看到, SR6604X-2路由器的OSPF 邻居关系一直在Full/Init和Loading/Full间切换, 由于客户业务需要回复, 故障环境不存在了, 只能搜集了SR6604X-2路由器的邻居csico 75XX路由器日志, 信息如下:

```
Sep 15 00:27:46.086: %OSPF-5-ADJCHG: Process 100, Nbr 10.206.1.2 on POS7/3/1 from FULL to DOWN, Neighbor Down: Too many retransmissions
```

```
Sep 15 00:28:46.087: %OSPF-5-ADJCHG: Process 100, Nbr 10.206.1.2 on POS7/3/1 from DOWN to DOWN, Neighbor Down: Ignore timer expired
```

```
Sep 15 00:28:52.639: %OSPF-5-ADJCHG: Process 100, Nbr 10.206.1.2 on POS7/3/1 from LOADING to FULL, Loading Done
```

```
Sep 15 00:36:17.924: %OSPF-5-ADJCHG: Process 100, Nbr 10.206.1.2 on POS7/3/1 from FULL to DOWN, Neighbor Down: Too many retransmissions
```

```
Sep 15 00:37:17.924: %OSPF-5-ADJCHG: Process 100, Nbr 10.206.1.2 on POS7/3/1 from DOWN to DOWN, Neighbor Down: Ignore timer expired
```

```
Sep 15 00:37:22.628: %OSPF-5-ADJCHG: Process 100, Nbr 10.206.1.2 on POS7/3/1 from LOADING to FULL, Loading Done
```

```
Sep 15 00:53:10.962: %OSPF-5-ADJCHG: Process 100, Nbr 10.206.1.2 on POS7/3/1 from FULL to DOWN, Neighbor Down: Too many retransmissions
```

```
Sep 15 00:54:10.962: %OSPF-5-ADJCHG: Process 100, Nbr 10.206.1.2 on POS7/3/1 from DOWN to DOWN, Neighbor Down: Ignore timer expired
```

```
Sep 15 00:54:12.718: %OSPF-5-ADJCHG: Process 100, Nbr 10.206.1.2 on POS7/3/1 from LOADING to FULL, Loading Done
```

从上面的日志分析, Cisco 75xx的邻居关系震荡的原因是Too many retransmissions, 根据以往的经验, 应该和两端的MTU设置不一致有关系。Cisco 75xx的pos口的默认MTU是4770字节, SR6604的默认MTU是1500字节。这样导致从Cisco 7500发送的大量Isa的报文被SR6604侧丢弃, 从而导致OSPF邻居关系震荡。

在实验室也成功复现了这种情况。实验室中使用SR6604-X FIP210 和Cisco ASR1002-X通过GE口 (实验室使用GE口替代POS验证) 建立OSPF邻居 (Cisco 侧MTU设置为4770)。使用测试仪TestCenter向ASR注入1000条本地路由、50000万条外部路由。过了一会, 故障复现了, SR6604-X路由器的邻居From Full To init, Cisco ASR因多次重传导致OSPF 邻居Down了, 同时, SR6604-x接口下也记录了丢弃超大报文的统计:

SR6604X-2的POS接口上记录了Giants超大包的统计如下:

Last clearing of counters: 09:00:07 Sun 09/14/2014
Last 300 seconds input rate 0 bytes/sec, 0 bits/sec, 0 packets/sec
Last 300 seconds output rate 0 bytes/sec, 0 bits/sec, 0 packets/sec
Input: 68828183 packets, 25650167182 bytes, 0 no buffers
1858 errors, 0 runts, 1754 giants, 103 CRC
0 overruns, 1 aborts
Output: 13852976 packets, 3304450475 bytes
0 errors, 0 underruns, 0 aborts

四 解决方法:

修改两侧的MTU一致即可, 由于SR6604X-2的fip200板卡为旧卡, 其上的POS接口MTU最多只能配置到2000字节, 所以需要修改对端的MTU为1500字节。