


## 某局点S6900端口配置突然丢失问题案例

ADDC方案 张文宁 2020-02-28 发表

### 组网及说明

拓扑如下:



### 问题描述

如图所示, 两台S6800作为ADDC弱控方案中的leaf设备, 现场反馈突然右侧的服务器无法被访问, 登录设备查看leaf和二层设备相连聚合口下的vxlan服务实例配置丢失了, 而且很多成员口无法被选中。客户通过自行创建聚合组并且在vcfc上点接口绑定vlan-vxlan映射关系后下发服务实例解决的。客户反馈故障前后没有人对拓扑中的设备做过任何操作, 然后找到这边问设备为什么配置突然会丢失。

### 过程分析

首先在VCFC上查找故障前的配置文件, 确认了leaf和二层设备之间的聚合口是agg 11:

S6800侧接口配置:

```
#
interface Bridge-Aggregation11
description TO-MOD2-18-HWS6855
port link-type trunk
undo port trunk permit vlan 1
port trunk permit vlan 2001 3003
vtep access port
#
service-instance 2001
encapsulation s-vid 2001
xconnect vsi SDN_VSI_12001
#
service-instance 3003
encapsulation s-vid 3003
xconnect vsi SDN_VSI_13003
```

华为侧配置:

```
#
interface Eth-Trunk0
port link-type trunk
undo port trunk allow-pass vlan 1
port trunk allow-pass vlan 2001
#
```

查看到接口有vtep access port命令, 且dir看到设备有自动化相关的模板:

```
=====dir /all /all-filestems=====
```

Directory of flash:

```
0 -rw- 4845 Jun 18 2019 15:24:34 545_leaf.template
1 -rw- 135489536 Apr 17 2019 15:38:31 S6900F-CMW710-R2612P02.ipe
```

**进一步查看故障时候S6800的配置发现确实有大量undo 配置的操作, 根据Line=-IPAddr=-\*-User=-\*\*信息可以判断为非人为操作, 而且删除配置前有VCF\_DOWN\_LINK提示, 因此基本确认了是自动化进程重走了一遍自动聚合过程:**

```
%@441758%Sep 2 14:48:40:378 2019 XM-JM601-MOD2-45U-TOR-H3CS6900 CFGMAN/5/CFGMAN_CFGCHANGED: -EventIndex=10024-CommandSource=snmp-COnfigSource=startup-COnfigDestination=running; Configuration changed.
```

```
%@441759%Sep 2 14:48:43:838 2019 XM-JM601-MOD2-45U-TOR-H3CS6900 VCF/6/VCF_DOWN_LINK: In phase 2.0.8, device with MAC address 7c1e-0627-38a2 discovered downlink interface F0/52.
```

```
%@441760%Sep 2 14:48:44:410 2019 XM-JM601-MOD2-45U-TOR-H3CS6900 VCF/6/VCF_DOWN
```

N\_LINK: In phase 2.0.8, device with MAC address 7c1e-0627-38a2 discovered downlink interface FortyGigE1/0/51.

```
%@441761%Sep 2 14:48:48:841 2019 XM-JM601-MOD2-45U-TOR-H3CS6900 SHELL/6/SHELL_C
MD: -Line=-IPAddr=**-User=**; Command is system-view
%@441762%Sep 2 14:48:50:701 2019 XM-JM601-MOD2-45U-TOR-H3CS6900 SHELL/6/SHELL_C
MD: -Line=-IPAddr=**-User=**; Command is interface FortyGigE1/0/52
%@441763%Sep 2 14:48:50:892 2019 XM-JM601-MOD2-45U-TOR-H3CS6900 SHELL/6/SHELL_C
MD: -Line=-IPAddr=**-User=**; Command is system-view
%@441764%Sep 2 14:48:50:900 2019 XM-JM601-MOD2-45U-TOR-H3CS6900 SHELL/6/SHELL_C
MD: -Line=-IPAddr=**-User=**; Command is interface FortyGigE1/0/52
%@441765%Sep 2 14:48:50:905 2019 XM-JM601-MOD2-45U-TOR-H3CS6900 SHELL/6/SHELL_C
MD: -Line=-IPAddr=**-User=**; Command is undo lldp compliance admin-status cdp
%@441766%Sep 2 14:48:50:913 2019 XM-JM601-MOD2-45U-TOR-H3CS6900 SHELL/6/SHELL_C
MD: -Line=-IPAddr=**-User=**; Command is undo port trunk permit vlan 2001 3003
%@441767%Sep 2 14:48:50:952 2019 XM-JM601-MOD2-45U-TOR-H3CS6900 SHELL/6/SHELL_C
MD: -Line=-IPAddr=**-User=**; Command is port trunk permit vlan 1
```

因为设备走自动化聚合的条件是lldp模块发现多根线和对端设备相连的sysname和设备角色都一样，会判断到这几个口实际连的是一台设备应该做聚合因此联动到自动化模块走创建聚合的脚本。一般是首次接线，或者对端重启接口全down了，或者虽然接口没有down但由于一些原因lldp邻居收不到对端的信息老化后突然又收到才会触发这个过程。

**但是从日志中可以看到故障前后，成员口并没有updown或者lldp邻居超时的信息。**

所以现在疑点有两个：

- 1、为什么触发了重走聚合的过程。
- 2、为什么触发了重走聚合的过程，仍然无法恢复业务。

**设备上/var/log/fabric.log记录了设备的自动化日志（probe视图下执行 view /var/log/fabric.log > fbr ic.log 然后将flash下的fabric.log导出），通过日志可以看到故障时间点突然收到一个拓扑改变的事件，原因是发现对端设备系统名称改变了：**

```
2019-09-02 14:48:43,545 topo.py[line:87] DEBUG [topology]--Notified--topology change {'FortyGigE1
/0/52': {'add': [], 'del': [], 'mod': [{'PeerIP': None, 'Oldrole': 'computer', 'sysName': 'JM601-MOD2-18-U4
5-HWCE6855', 'sysdesc': '', 'Loopback1IP': '0.0.0.0', 'OldsysName': 'M2-18-U45-HWCE6855', 'mac': 'e
0cc-7aec-bc21', 'PeerIntf': '40GE1/0/1', 'role': 'computer', 'LoopbackIP': '10.104.51.74', 'OldPeerIntf': '4
0GE1/0/1', 'OldPeerIP': None, 'OldLoopbackIP': '10.104.51.74', 'OldLoopback1IP': '0.0.0.0', 'Oldmac': '
e0cc-7aec-bc21', 'oldsysdesc': ''}]}}
```

```
2019-09-02 14:48:43,553 base.py[line:257] DEBUG [underlay]--Notified--topology change {'data': {'u'F
ortyGigE1/0/52': {'u'add': [], 'u'del': [], 'u'mod': [{'u'PeerIP': None, 'u'Oldrole': 'u'computer', 'u'sysName': 'u'J
M601-MOD2-18-U45-HWCE6855', 'u'sysdesc': 'u'', 'u'Loopback1IP': 'u'0.0.0.0', 'u'OldsysName': 'u'M2-
18-U45-HWCE6855', 'u'mac': 'u'e0cc-7aec-bc21', 'u'PeerIntf': 'u'40GE1/0/1', 'u'role': 'u'computer', 'u'Loop
backIP': 'u'10.104.51.74', 'u'OldPeerIntf': 'u'40GE1/0/1', 'u'OldPeerIP': None, 'u'OldLoopbackIP': 'u'10.10
4.51.74', 'u'OldLoopback1IP': 'u'0.0.0.0', 'u'Oldmac': 'u'e0cc-7aec-bc21', 'u'oldsysdesc': 'u'}}], 'event': 9}
module [<underlay.allocation.updown_link.UpDownLinkIF object at 0xffe8a87990>, <underlay.allocati
on.aggregation.AggrAllocation object at 0xffe8a87d10>, <underlay.allocation.irf.IRFAllocation object
at 0xffe8a87cd0>, <underlay.underlay_route.bgp object at 0xffe8a87c10>]
```

```
2019-09-02 14:48:43,557 topo.py[line:87] DEBUG [topology]--Notified--topology change {'FortyGigE1
/0/51': {'add': [], 'del': [], 'mod': [{'PeerIP': None, 'Oldrole': 'computer', 'sysName': 'JM601-MOD2-18-U4
5-HWCE6855', 'sysdesc': '', 'Loopback1IP': '0.0.0.0', 'OldsysName': 'M2-18-U45-HWCE6855', 'mac': 'e
0cc-7aec-bc21', 'PeerIntf': '40GE2/0/1', 'role': 'computer', 'LoopbackIP': '10.104.51.74', 'OldPeerIntf': '4
0GE2/0/1', 'OldPeerIP': None, 'OldLoopbackIP': '10.104.51.74', 'OldLoopback1IP': '0.0.0.0', 'Oldmac': '
e0cc-7aec-bc21', 'oldsysdesc': ''}]}}
```

```
2019-09-02 14:48:43,559 command_executor.py[line:37] DEBUG Execute command: system ; interfa
ce FortyGigE1/0/52 ; return
```

和现场沟通确认，对端华为设备确实是在故障时间点做了sysname的修改操作。

经确认当前设备的自动化实现确实如此，当对端sysname变化的时候，认为聚合对端的设备发生了改变，所以会重新走自动聚合。（但实际上此处有待优化，因为修改系统名称是常常有的事，已提需求建议后续如果仅仅是sysname发生变化，不触发重新聚合过程）。

**以下是设备成员口离开聚合和重新创建了聚合口2047的过程日志：**

**离开原有聚合组：**

```
2019-09-02 14:48:43,687 command_executor.py[line:37] DEBUG Execute command: system-view ;i
```

```
interface Bridge-Aggregation11 ; port link-type trunk ; vtep access port
2019-09-02 14:48:43,845 base.py[line:36] DEBUG AggrAllocation <process_neighbor_mod:642>: Port
FortyGigE1/0/52 process neighbor {u'PeerIP': None, u'Oldrole': u'computer', u'sysName': u'JM601-
MOD2-18-U45-HWCE6855', u'sysdesc': u'', u'Loopback1IP': u'0.0.0.0', u'OldsysName': u'M2-18-U45-
HWCE6855', u'mac': u'e0cc-7aec-bc21', u'PeerIntf': u'40GE1/0/1', u'role': u'computer', u'LoopbackIP':
u'10.104.51.74', u'OldPeerIntf': u'40GE1/0/1', u'OldPeerIP': None, u'OldLoopbackIP': u'10.104.51.74',
u'OldLoopback1IP': u'0.0.0.0', u'Oldmac': u'e0cc-7aec-bc21', u'oldsysdesc': u''} modify
2019-09-02 14:48:43,854 base.py[line:36] DEBUG AggrAllocation <process_neighbor_del:608>: Port
FortyGigE1/0/52 process neighbor {u'PeerIP': None, u'Oldrole': u'computer', u'sysName': u'M2-18-U4
5-HWCE6855', u'sysdesc': u'', u'Loopback1IP': u'0.0.0.0', u'OldsysName': u'M2-18-U45-HWCE6855',
u'mac': u'e0cc-7aec-bc21', u'PeerIntf': u'40GE1/0/1', u'role': u'computer', u'LoopbackIP': u'10.104.51.7
4', u'OldPeerIntf': u'40GE1/0/1', u'OldPeerIP': None, u'OldLoopbackIP': u'10.104.51.74',
u'OldLoopback1IP': u'0.0.0.0', u'Oldmac': u'e0cc-7aec-bc21', u'oldsysdesc': u''} delete
2019-09-02 14:48:43,859 base.py[line:36] DEBUG AggrAllocation
<process_delete_aggregation:584>: Remove member port FortyGigE1/0/52 from [u'FortyGigE5/0/52',
u'FortyGigE5/0/51', u'FortyGigE1/0/51', u'FortyGigE1/0/52']
2019-09-02 14:48:43,863 command_executor.py[line:37] DEBUG Execute command: system ; interfa
ce FortyGigE1/0/52 ; undo port link-aggregation group
2019-09-02 14:48:44,068 base.py[line:36] DEBUG AggrAllocation
<process_delete_aggregation:589>: Remove member FortyGigE1/0/52's config
```

#### 重新加入新增聚合组:

```
%@441795%Sep 2 14:48:58:493 2019 XM-JM601-MOD2-45U-TOR-H3CS6900 VCF/6VCF_AGGR
_CREATE: In phase 2.0.5, device with MAC address 7c1e-0627-38a2 created aggregation group204
7. The member port list is FortyGigE1/0/52,FortyGigE1/0/51.
```

```
2019-09-02 14:48:45,156 command_executor.py[line:37] DEBUG Execute command: system ; interfa
ce bridge-aggregation 2047
2019-09-02 14:48:45,205 command_executor.py[line:37] DEBUG Execute command: system ; interfa
ce FortyGigE1/0/52 ; default
2019-09-02 14:48:51,039 command_executor.py[line:37] DEBUG Execute command: system ; interfa
ce FortyGigE1/0/52 ; lldp compliance admin-status cdp txrx ;
2019-09-02 14:48:51,093 command_executor.py[line:37] DEBUG Execute command: system ; interfa
ce FortyGigE1/0/52 ; port link-aggregation group 2047 force
2019-09-02 14:48:51,359 aggregation.py[line:248] DEBUG add_members result: ['system', 'System V
iew: return to User View with Ctrl+Z.', 'XM-JM601-MOD2-45U-TOR-H3CS6900', 'interface FortyGig
E1/0/52', 'XM-JM601-MOD2-45U-TOR-H3CS6900-FortyGigE1/0/52', 'port link-aggregation group 2
047 force', 'XM-JM601-MOD2-45U-TOR-H3CS6900-FortyGigE1/0/52']]
2019-09-02 14:48:51,363 command_executor.py[line:37] DEBUG Execute command: system ; interfa
ce FortyGigE1/0/51 ; default
2019-09-02 14:48:57,185 command_executor.py[line:37] DEBUG Execute command: system ; interfa
ce FortyGigE1/0/51 ; lldp compliance admin-status cdp txrx ;
2019-09-02 14:48:57,238 command_executor.py[line:37] DEBUG Execute command: system ; interfa
ce FortyGigE1/0/51 ; port link-aggregation group 2047 force
```

**但是第二个疑点，为什么重走了自动聚合过程，业务仍然无法恢复？看日志聚合2047的配置理论上是能够让服务器重新上线的：**

```
2019-09-02 15:43:07,203 command_executor.py[line:37] DEBUG Execute command: system ; interfa
ce Bridge-Aggregation2047 ; port link-type trunk
2019-09-02 15:43:07,349 command_executor.py[line:37] DEBUG Execute command: system ; interfa
ce Bridge-Aggregation2047 ; undo port trunk permit vlan 1
2019-09-02 15:43:07,463 topo.py[line:87] DEBUG [topology]--Notified--topology change {'FortyGigE5
/0/51': {'add': [{'PeerIP': None, 'sysName': 'JM601-MOD2-18-U45-HWCE6855', 'sysdesc': '',
'Loopback1IP': '0.0.0.0', 'mac': 'e0cc-7aec-bc21', 'PeerIntf': '40GE1/0/2', 'role': 'computer',
'LoopbackIP': '10.104.51.74'}], 'del': [], 'mod': []}}
2019-09-02 15:43:07,468 topo_collect.py[line:822] DEBUG ifIndex 703 peerChassisID e0cc-7aec-bc2
1 peerPortID 40GE1/0/2 lldpAgent 1 event 1
2019-09-02 15:43:07,498 command_executor.py[line:37] DEBUG Execute command: system ; interfa
ce Bridge-Aggregation2047 ; port trunk permit vlan 2001 3003
2019-09-02 15:43:07,644 command_executor.py[line:37] DEBUG Execute command: system ; interfa
ce Bridge-Aggregation2047 ; vtep access port
2019-09-02 15:43:07,706 aggregation.py[line:226] DEBUG do config [port link-type trunk', 'undo port
trunk permit vlan 1', 'port trunk permit vlan 2001 3003'] on aggr 2047
```

```
2019-09-02 15:43:07,711 command_executor.py[line:37] DEBUG Execute command: system ; interface Bridge-Aggregation2047 ; port link-type trunk
2019-09-02 15:43:07,875 command_executor.py[line:37] DEBUG Execute command: system ; interface Bridge-Aggregation2047 ; undo port trunk permit vlan 1
2019-09-02 15:43:08,023 command_executor.py[line:37] DEBUG Execute command: system ; interface Bridge-Aggregation2047 ; port trunk permit vlan 2001 3003
2019-09-02 15:43:08,172 command_executor.py[line:37] DEBUG Execute command: system ; interface bridge-aggregation 2047 ; link-aggregation mode dynamic
```

因为理论接口有vtep access port命令，设备会将arp上送VCFC，VCFC会根据vlan-vxlan映射表下发服务实例的。而且原先的聚合口agg11编号很小，不像是自动化模块创建的，像是手工配置的。（设备自动化聚合编号从2048编号往前分）。

重新和现场进一步沟通确认，聚合11是手工配置的，并且是静态聚合，上线的时候是先和设备手工创建聚合组，然后在VCFC侧手工绑定这个聚合11接口到vlan-vxlan关系表中，然后插上S6800和華為设备之间的线。

因此第二点也清晰明了了，现场在vlan-vxlan映射表中没有绑定这个leaf设备，而且设备自动创建的具体聚合组在vcfc上也不会有绑定关系，因此VCFC是不会给2047下发服务实例的。而且也是因为对端華為是静态聚合，我们这边变成了动态聚合，所以很多接口都是非选中状态。

综上，现场S6800聚合配置被删除是因为对端设备修改了sysname触发了自动化进程重新聚合，重新聚合仍无法恢复业务是因为原本对端是静态聚合而且VCFC的vlan-vxlan关系表没有绑定这台leaf设备。

## 解决方法

### 解决方法:

在VCFC上手工将聚合2047绑定到vlan-vxlan关系表中解决。

### 总结:

现场S6800接口不足，因此中间挂了二层交换机，但是此时是静态聚合，用不到设备的自动化进程，因此建议leaf设备停掉自动化聚合。如果想用自动聚合，建议不要手工配置静态聚合。