

知 某局点S7606X堆叠口流量不均衡问题

IRF2 董智敏 2021-12-14 发表

组网及说明

两台S7606-X设备做堆叠，堆叠口为37到48口。

问题描述

两台S7606-X做堆叠，堆叠口包括37到48口，现场发现37口的流量偏大，大约占用带宽20%，其他堆叠口只有2%-4%。看收发包计数，发现只有37口有广播流量，虽然广播流量都走主堆叠口，但从数据看，广播流量和单播、组播流量也不是一个量级的，大部分都是单播流量，单播流量应该是在堆叠口哈希的

过程分析

堆叠跨板hash, 组播/广播/未知单播都走主堆叠口, 单播会进行hash, 但是这里客户配置了mac不老化, 导致mac超规格, 单播变成未知单播, 都走了主堆叠口。

#

```
mac-address timer no-aging
```

#

```
=====display mac-address=====
```

MAC Address	VLAN ID	State	Port/Nickname	Aging
10e8-78ce-6a92	10	Learned	BAGG3	N
10e8-78cd-5692	11	Learned	BAGG4	N
10e8-78ce-6a92	50	Learned	BAGG3	N

```
%Dec 24 00:34:31:729 2020 SHNQ-PA-CMNET-SW03-lyljj7606X RESMON/3/RESMON_SEVERE: -  
Chassis=1-Slot=4; -Resource=mac-Total=163840-Used=163570-Free=270; Free resource  
decreased to or below severe threshold 10%.
```

```
%Dec 24 00:34:32:154 2020 SHNQ-PA-CMNET-SW03-lyljj7606X RESMON/3/RESMON_SEVERE: -  
Chassis=2-Slot=0; -Resource=mac-Total=163840-Used=163572-Free=268; Free resource  
decreased to or below severe threshold 10%.
```

```
%Dec 24 00:58:11:770 2020 SHNQ-PA-CMNET-SW03-lyljj7606X RESMON/3/RESMON_SEVERE: -  
Chassis=2-Slot=4; -Resource=mac-Total=163840-Used=163559-Free=281; Free resource  
decreased to or below severe threshold 10%.
```

```
%Dec 24 01:40:59:546 2020 SHNQ-PA-CMNET-SW03-lyljj7606X RESMON/3/RESMON_SEVERE: -  
Chassis=1-Slot=0; -Resource=mac-Total=163840-Used=163578-Free=262; Free resource  
decreased to or below severe threshold 10%.
```

可以通过如下命令查看哪个接口是主堆叠口

```
====debug stack show portinfo chassis 1 slot 1====
```

```
=====  
Port Information of STACK Module  
=====
```

```
Stack 1
```

```
Portnum: 0, ActiveNum:0, IsShutDown:0, Mtu:0, FwdStat:???
```

```
PortOnBoard:
```

```
Stack Ports: Master~255/255/255:255 IsChecked:1
```

```
Stack 2
```

```
Portnum: 12, ActiveNum:12, IsShutDown:0, Mtu:0, FwdStat:FOWARD
```

```
PortOnBoard: Slot 1:(12,12)
```

```
Stack Ports: Master~1/0/37:0 IsChecked:1
```

```
1: 1/0/37:0, PortType:2, IsActive:1, HgMode:1 (1), IsShutDown:0, LogSlot: 1  
2: 1/0/38:0, PortType:2, IsActive:1, HgMode:1 (1), IsShutDown:0, LogSlot: 1  
3: 1/0/39:0, PortType:2, IsActive:1, HgMode:1 (1), IsShutDown:0, LogSlot: 1  
4: 1/0/40:0, PortType:2, IsActive:1, HgMode:1 (1), IsShutDown:0, LogSlot: 1  
5: 1/0/41:0, PortType:2, IsActive:1, HgMode:1 (1), IsShutDown:0, LogSlot: 1  
6: 1/0/42:0, PortType:2, IsActive:1, HgMode:1 (1), IsShutDown:0, LogSlot: 1  
7: 1/0/43:0, PortType:2, IsActive:1, HgMode:1 (1), IsShutDown:0, LogSlot: 1  
8: 1/0/44:0, PortType:2, IsActive:1, HgMode:1 (1), IsShutDown:0, LogSlot: 1  
9: 1/0/45:0, PortType:2, IsActive:1, HgMode:1 (1), IsShutDown:0, LogSlot: 1  
10: 1/0/46:0, PortType:2, IsActive:1, HgMode:1 (1), IsShutDown:0, LogSlot: 1  
11: 1/0/47:0, PortType:2, IsActive:1, HgMode:1 (1), IsShutDown:0, LogSlot: 1  
12: 1/0/48:0, PortType:2, IsActive:1, HgMode:1 (1), IsShutDown:0, LogSlot: 1
```

解决方法

建议现场将mac地址配置成可老化，防止mac超规格产生大量未知单播流量。补充：同一堆叠单板，同一芯片的堆叠口可以进行聚合hash；但对于跨单板聚合的情况，如果当本堆叠单板的堆叠口流量打满后，不会hash再到其它单板的堆叠口，应该避免有过大的跨框流量，尽量本框优先转发。

