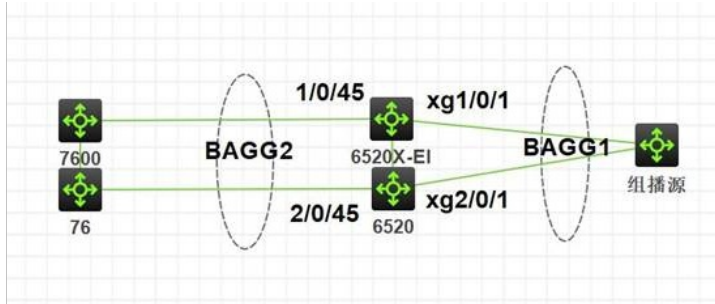


# 某局点S6520X-EI跨框聚合组播流量不hash

PIM 二层链路聚合 许家豪 2022-03-17 发表

## 组网及说明

组网如下图：



## 问题描述

问题描述: 两台6520做堆叠, 设备开启PIM-SM跑组播流量, 流量从聚合口1进入, 从聚合口2发出, 但是聚合口2中只有一个链路Ten1/0/45口有流量

Ten-GigabitEthernet1/0/1

Current state: UP

Line protocol state: UP

...

Last 300 second input: 32091 packets/sec 44222049 bytes/sec 36%

Last 300 second output: 0 packets/sec 77 bytes/sec 0%

Ten-GigabitEthernet2/0/1

Current state: UP

Line protocol state: UP

...

Last 300 second input: 46122 packets/sec 63556982 bytes/sec 52%

Last 300 second output: 0 packets/sec 106 bytes/sec 0%

Ten-GigabitEthernet1/0/45

Current state: UP

Line protocol state: UP

...

Last 300 second input: 0 packets/sec 15 bytes/sec 0%

Last 300 second output: 77245 packets/sec 106443727 bytes/sec 9%

Ten-GigabitEthernet2/0/45

Current state: UP

Line protocol state: UP

...

Last 300 second input: 1 packets/sec 146 bytes/sec 0%

Last 300 second output: 0 packets/sec 27 bytes/sec 0%

## 过程分析

### 过程分析:

聚合口选中正常, 端口收发光正常, 端口正常放通业务vlan, 组播邻居及组播(S,G)表项正常  
成员口状态正常:

Aggregate Interface: Bridge-Aggregation2

Creation Mode: Manual

Aggregation Mode: Dynamic

Loadsharing Type: Shar

Management VLANs: None

System ID: 0x8000, f474-88d2-9c10

Local:

Port Status Priority Index Oper-Key Flag

XGE1/0/45 S 32768 3 2 {ACDEF}

XGE2/0/45 S 32768 4 2 {ACDEF}

Remote:

Actor Priority Index Oper-Key SystemID Flag

XGE1/0/45(R) 32768 3 3 0x8000, 84d9-31d8-0a00 {ACDEF}

XGE2/0/45 32768 31 3 0x8000, 84d9-31d8-0a00 {ACDEF}

### 收发光正常:

Ten-GigabitEthernet1/0/45 transceiver diagnostic information:

Current diagnostic parameters:

Temp.(j) Voltage(V) Bias(mA) RX power(dBm) TX power(dBm)

33 3.31 29.52 -4.71 -3.45

Alarm thresholds:

Temp.(j) Voltage(V) Bias(mA) RX power(dBm) TX power(dBm)

High 73 3.80 88.00 2.50 1.00

Low -3 2.81 1.00 -16.40 -11.25

Ten-GigabitEthernet2/0/45 transceiver diagnostic information:

Current diagnostic parameters:

Temp.(j) Voltage(V) Bias(mA) RX power(dBm) TX power(dBm)

33 3.35 35.66 -2.90 -2.64

Alarm thresholds: Temp.(j) Voltage(V) Bias(mA) RX power(dBm) TX power(dBm)

High 73 3.80 99.00 2.50 2.50

Low -8 2.81 1.00 -16.40 -11.20

### 端口放通vlan正常

VLAN ID: 3333

VLAN type: Static

Route interface: Configured

IPv4 address: 172.26.252.154

IPv4 subnet mask: 255.255.255.252

Description: LingShiHuLian-6520

Name: VLAN 3333

Tagged ports:

Bridge-Aggregation2

Ten-GigabitEthernet1/0/45

Ten-GigabitEthernet2/0/45

Untagged ports: None

尝试修改hash方式, 如下中所有的方法都尝试过, 全都无法hash

## 2. 配置限制和指导

同一聚合接口下, 同时配置了静态聚合负载分担和动态聚合负载分担时, 动态聚合负载分担生效。

目前, 在系统视图下进行全局聚合负载分担类型配置, 交换机只支持:

- 根据报文类型自动匹配负载分担类型;
- 根据源IP地址进行聚合负载分担;
- 根据目的IP地址进行聚合负载分担;
- 根据源MAC地址进行聚合负载分担;
- 根据目的MAC地址进行聚合负载分担;
- 根据报文入端口进行聚合负载分担;

- 根据源IP地址与目的IP地址进行聚合负载分担;
- 根据源IP地址与源端口进行聚合负载分担;
- 根据目的IP地址与目的端口进行聚合负载分担;

问题原因: 已知组播报文出接口为聚合接口目的报文的转发是单播转发, 本地优先转发对已知单播报文转发, 因此本地已知单播报文生效目的地址之向不同物理接口进行聚合转发, 没有实现负责分担。因此, 关闭本框优先转发可实现流量hash转发

关闭本框优先转发 (undo link-aggregation load-sharing mode local-first)后, 流量可正常hash

[S6520]dis cri

```
Usage: Bandwidth utilization in percentage
```

Interface	Usage (%)	Total (pps)	Broadcast (pps)	Multicast (pps)
BAGG1	0	1	--	--
BAGG2	5	83372	--	--
XGE1/0/1	0	0	--	--
XGE1/0/45	5	48674	--	--
XGE1/0/47	3	23271	--	--
XGE1/0/48	2	21779	--	--
XGE2/0/1	0	1	--	--
XGE2/0/45	4	34698	--	--
XGE2/0/46	4	38317	--	--
XGE2/0/47	3	29230	--	--
XGE2/0/48	3	29360	--	--

