

Lag Resilient hash

From 2.9.2,pica8 switch support lag resilient hash. If one port removed from the lag or link_down,the traffic of this port will hash to the other ports of the lag without affect original traffic of the other ports.

Command

ovs-vsctl -- set Interface ae1 options:hash-mapping=resilient: Once enable resilient hash,all the hash fields in advance mode are enabled by default.

ovs-vsctl set-max-resilient-hash-lag-count COUNT : the valid value of COUNT is 1, 2, 4, 8, 16, 32, 64; The default value is 1.The COUNT means the max count of lags which can be set to resilient hash.



support platform

At present, trident2, trident2 plus support resilient hash.Eg:4048,4806,P5401,P5101,as5712,AS6701_32x,as6712_32x,as5812

Example

set hash-mapping-fields to mac,ip,vlan

Test Name	Test basic flow
Test Configuration	
Test Procedure	<p>Step1:configure bridge and port</p> <pre>ovs-vsctl add-br br0 -- set bridge br0 datapath_type=pica8</pre> <pre>ovs-vsctl add-port br0 te-1/1/1 vlan_mode=trunk tag=1 -- set interface te-1/1/1 type=pica8</pre> <pre>ovs-vsctl add-port br0 te-1/1/2 vlan_mode=trunk tag=1 -- set interface te-1/1/2 type=pica8</pre> <pre>ovs-vsctl add-port br0 te-1/1/3 vlan_mode=trunk tag=1 -- set interface te-1/1/3 type=pica8</pre> <pre>ovs-vsctl add-port br0 ae1 vlan_mode=trunk tag=1 -- set interface ae1 type=pica8_lag options:members=te-1/1/2,te-1/1/3,te-1/1/4 options:lag_type=static</pre> <pre>ovs-vsctl -- set Interface ae1 options:hash-mapping=resilient</pre> <p>step2:add flow</p> <pre>ovs-ofctl add-flow br0 in_port=1,actions=1025</pre> <p>step3:set hash-mapping-field to dl_dst:</p> <pre>ovs-vsctl set-lag-advance-hash-mapping-fields dl_dst</pre> <p>Send 12packets/sec ipv4 packets to te-1/1/1,with different dst mac(22:22:22:22:22:10 to 22:22:22:22:22:15):</p> <pre>MAC: Destination Address : 22 22 22 22 22 10</pre> <pre>MAC: Source Address : 22 11 11 11 11 40</pre> <pre>Vlan-----20</pre> <p>Te-1/1/2,te-1/1/3,te-1/1/4 will transmit the packets,eg:te-1/1/2 transmit packets with dst mac 22:11:11:11:11:10,13,15,te-1/1/3 transmit packets with dst mac 11,14,te-1/1/4 transmit packets with dst mac 12.</p> <p>send 12packets/sec ip packets with src ip increasement to te-1/1/1,then one of the lag port transmit all the packets.</p> <p>Down te-1/1/2 then the traffic of te-1/1/2 will be hash to te-1/1/3 and te-1/1/4, the original traffic of te-1/1/3 and te-1/1/4 will not be redistribute.</p> <p>Then up the te-1/1/2 port, the traffic of te-1/1/2 is some traffic of the other two ports.</p> <p>Step4:set hash-mapping-field to dl_src</p>

ovs-vsctl set-lag-advance-hash-mapping-fields dl_src

Send 12packets/sec ipv4 packets to te-1/1/1,with different src mac(22:11:11:11:11:40 to 22:11:11:11:11:45):

MAC: Destination Address : 22 22 22 22 22 10

MAC: Source Address : 22 11 11 11 11 40

Vlan-----20

Te-1/1/2,te-1/1/3,te-1/1/4 will transmit the packets,eg:te-1/1/2 transmit packets with dst mac 22:11:11:11:11:40,43,45,te-1/1/3 transmit packets with dst mac 41,44,te-1/1/4 transmit packets with dst mac 42.

Down te-1/1/2 then the traffic of te-1/1/2 will be hash to te-1/1/3 and te-1/1/4, the original traffic of te-1/1/3 and te-1/1/4 will not be redistribute.

Then up the te-1/1/2 port, the traffic of te-1/1/2 is some traffic of the other two ports

Step5:set hash-mapping-field to nw_src

ovs-vsctl set-lag-advance-hash-mapping-fields nw_src

send 12packets/s contents with src ip 1.1.1.1-1.1.6:

MAC: Destination Address : 22 22 22 22 22 10

MAC: Source Address : 22 11 11 11 11 40

Vlan-----20

IP: Source Address = 1.1.1.1

Te-1/1/2,te-1/1/3,te-1/1/4 will transmit the packets,eg:te-1/1/2 transmit packets with source ip 1.1.1.1 and 1.1.1.3. te-1/1/3 transmit packets with source ip 1.1.1.2 and 1.1.1.4 ,te-1/1/4 transmit packets with srouce ip 1.1.1.5 and 1.1.1.6.

Down te-1/1/2 then the traffic of te-1/1/2 will be hash to te-1/1/3 and te-1/1/4, the original traffic of te-1/1/3 and te-1/1/4 will not be redistribute.

Then up the te-1/1/2 port, the traffic of te-1/1/2 is some traffic of the other two ports

Step6:set hash-mapping-field to nw_dst

ovs-vsctl set-lag-advance-hash-mapping-fields nw_dst

send 12packets/s contents with dst ip 4.4.4.1-4.4.4.6:

MAC: Destination Address : 22 22 22 22 22 10

MAC: Source Address : 22 11 11 11 11 40

Vlan-----20

IP: Source Address = 1.1.1.1

IP: dst Address = 4.4.4.1

Te-1/1/2,te-1/1/3,te-1/1/4 will transmit the packets,eg:te-1/1/2 transmit packets with dst ip 4.4.4.1 and 4.4.4.3. te-1/1/3 transmit packets with source ip 4.4.4.2 and 4.4.4.4 ,te-1/1/4 transmit packets with srouce ip 4.4.4.5 and 4.4.4.6.

Down te-1/1/2 then the traffic of te-1/1/2 will be hash to te-1/1/3 and te-1/1/4, the original traffic of te-1/1/3 and te-1/1/4 will not be redistribute.

Then up the te-1/1/2 port, the traffic of te-1/1/2 is some traffic of the other two ports

Step6:set hash-mapping-field to vlan

ovs-vsctl set-lag-advance-hash-mapping-fields vlan

send 12packets/s contents with vlan from 20 to 25

send 12packets/s contents with vlan20-25:

MAC: Destination Address : 22 22 22 22 22 10

MAC: Source Address : 22 11 11 11 11 40

Vlan-----20

IP: Source Address = 1.1.1.1

IP: dst Address = 4.4.4.1

Te-1/1/2,te-1/1/3,te-1/1/4 will transmit the packets,eg:te-1/1/2 transmit packets with vlan20,22. te-1/1/3 transmit packets with vlan 21 ,te-1/1/4 transmit packets with vlan 23,24,25.

Down te-1/1/3 then the traffic of te-1/1/3 will be hash to te-1/1/2 or te-1/1/4, the original traffic of te-1/1/2 and te-1/1/4 will not be redistribute.

Then up the te-1/1/3 port, the traffic of te-1/1/3 is some traffic of the other two ports.