

Nexus 7000: Understand "hardware ip glean throttle" Feature

Contents

[Introduction](#)

[Prerequisites](#)

[Requirements](#)

[Components Used](#)

[Background Information](#)

[Lab Test](#)

[Related Information](#)

Introduction

This document describes how the **hardware ip glean throttle** feature works with examples and the intention of this feature.

Prerequisites

Requirements

Cisco recommends that you have basic knowledge of Nexus 7000 Series Switches configuration.

Components Used

The information in this document is based on these software and hardware versions:

- Nexus 7000 with Release 6.2.x and later
- F2e series line card

The information in this document was created from the devices in a specific lab environment. All of the devices used in this document started with a cleared (default) configuration. If your network is live, make sure that you understand the potential impact of any command.

Background Information

When you forward an incoming IP packet in a line card, if the Address Resolution Protocol (ARP) request for the next hop is not resolved, the line card forwards the packets to the supervisor in order to generate an ARP request. Once the ARP request responds to the supervisor, it resolves the MAC address for the next hop and programs the hardware.

If the supervisor cannot resolve the ARP entry, then the linecard sends all packets destined to that address to the supervisor. The supervisor generates ARP requests indefinitely until the ARP entry is resolved. There is a hardware rate limiter called glean placed in order to protect the supervisor's

processor (CPU) from excessive traffic.

An issue that can arise is a single destination IP drops off the network due to maintenance or a hardware problem, and suddenly all traffic destined to it is being sent to the CPU. Since the rate limiter is in place, the CPU does not go high but this single destination IP can consume the entire rate limiter and not give other legitimate IP's access to the CPU. It is for this scenario that **hardware ip glean throttle** was created.

With the **hardware ip glean throttle** configuration, routed traffic for each unknown destination IP reaches the CPU post Hardware Rate Limiter (HWRL) action for ARP resolution. Unreachable destination will result in a /32 drop adjacency to be created in hardware. This prevents additional packets to the same next-hop IP address to be forwarded to the supervisor. While this drop adjacency is added, subsequent packets are dropped yet the supervisor continues to generate ARP requests until the next-hop is resolved. The drop adjacency is installed for a short period of time, which is configurable. Once the timer expires, one packet is again sent to the CPU and the process repeats. The number of entries that is installed in this fashion is limited to 1000 by default, but is configurable to a larger number of desired. This is to limit the impact on the Routing Information Base (RIB) table size.

Lab Test

In this case, you have a server, 172.28.191.200, which is down due to a hardware failure, and is currently unavailable to service traffic.

Note: There is no ARP entry for the host and no adjacency is created.

```
N7K# show ip route vrf VRF_ABC 172.28.191.200
IP Route Table for VRF "VRF_ABC"
 '*' denotes best ucast next-hop
 *** denotes best mcast next-hop
 '[x/y]' denotes [preference/metric]
 '%<string>' in via output denotes VRF <string>
```

```
172.28.191.192/28, ubest/mbest: 1/0, attached >>> There is no /32 entry
 *via 172.28.191.195, Vlan1601, [0/0], 02:01:17, direct
```

Traffic is sent to the supervisor in order to generate an ARP request:

```
N7K# show system internal forwarding vrf VRF_ABC ipv4 route 172.28.191.200 detail
slot 1
=====
RPF Flags legend:
  S - Directly attached route (S_Star)
  V - RPF valid
  M - SMAC IP check enabled
  G - SGT valid
  E - RPF External table valid
172.28.191.192/28 , sup-eth2
Dev: 0 , Idx: 0x65fb , Prio: 0x8487 , RPF Flags: VS , DGT: 0 , VPN: 9
RPF_Intf_5: Vlan1601 (0x19 )
AdjIdx: 0x5a , LIFB: 0 , LIF: sup-eth2 (0x1fe1 ), DI: 0xc01
DMAC: 0000.0000.0000 SMAC: 0000.0000.0000
172.28.191.192/28 , sup-eth2
```

```

Dev: 1 , Idx: 0x65fb , Prio: 0x8487 , RPF Flags: VS , DGT: 0 , VPN: 9
RPF_Intf_5: Vlan1601 (0x19 )
AdjIdx: 0x5a , LIFB: 0 , LIF: sup-eth2 (0x1fe1 ), DI: 0xc01
DMAC: 0000.0000.0000 SMAC: 0000.0000.0000
172.28.191.192/28 , sup-eth2
Dev: 2 , Idx: 0x65fb , Prio: 0x8487 , RPF Flags: VS , DGT: 0 , VPN: 9
RPF_Intf_5: Vlan1601 (0x19 )
AdjIdx: 0x5a , LIFB: 0 , LIF: sup-eth2 (0x1fe1 ), DI: 0xc01
DMAC: 0000.0000.0000 SMAC: 0000.0000.0000
172.28.191.192/28 , sup-eth2
Dev: 5 , Idx: 0x65f1 , Prio: 0x84f2 , RPF Flags: VS , DGT: 0 , VPN: 9
RPF_Intf_5: Vlan1601 (0x19 )
AdjIdx: 0x5a , LIFB: 0 , LIF: sup-eth2 (0x1fe1 ), DI: 0xc01
DMAC: 0000.0000.0000 SMAC: 0000.0000.0000

```

The glean rate limiter for the specific module throttles the traffic to 100 packets per second, per module. You can see that some of the packets gets dropped.

```

N7K# show hardware rate-limiter
Units for Config: packets per second
Allowed, Dropped & Total: aggregated since last clear counters
rl-1: STP and Fabricpath-ISIS
rl-2: L3-ISIS and OTV-ISIS
rl-3: UDLD, LACP, CDP and LLDP
rl-4: Q-in-Q and ARP request
rl-5: IGMP, NTP, DHCP-Snoop, Port-Security, Mgmt and Copy traffic

```

```
Module: 1
```

R-L Class	Config	Allowed	Dropped	Total
L3 mtu	500	0	0	0
L3 ttl	500	0	0	0
L3 control	10000	0	0	0
L3 glean	100	3326	3190	6516
L3 mcast dirconn	3000	0	0	0
L3 mcast loc-grp	3000	0	0	0
L3 mcast rpf-leak	500	0	0	0
L2 storm-ctrl	Disable			
access-list-log	100	0	0	0
copy	30000	1877	0	1877
receive	30000	318	0	318

When the **hardware ip glean throttle** command is configured:

```

N7K# show hardware rate-limiter
Units for Config: packets per second
Allowed, Dropped & Total: aggregated since last clear counters
rl-1: STP and Fabricpath-ISIS
rl-2: L3-ISIS and OTV-ISIS
rl-3: UDLD, LACP, CDP and LLDP
rl-4: Q-in-Q and ARP request
rl-5: IGMP, NTP, DHCP-Snoop, Port-Security, Mgmt and Copy traffic

```

```
Module: 1
```

R-L Class	Config	Allowed	Dropped	Total
L3 mtu	500	0	0	0
L3 ttl	500	0	0	0
L3 control	10000	0	0	0
L3 glean	100	3326	3190	6516
L3 mcast dirconn	3000	0	0	0
L3 mcast loc-grp	3000	0	0	0

```

L3 mcast rpf-leak          500          0          0          0
L2 storm-ctrl              Disable
access-list-log            100          0          0          0
copy                       30000         1877         0         1877
receive                    30000         318          0         318

```

An adjacency is installed in the RIB:

```

N7K# show hardware rate-limiter
Units for Config: packets per second
Allowed, Dropped & Total: aggregated since last clear counters
r1-1: STP and Fabricpath-ISIS
r1-2: L3-ISIS and OTV-ISIS
r1-3: UDLD, LACP, CDP and LLDP
r1-4: Q-in-Q and ARP request
r1-5: IGMP, NTP, DHCP-Snoop, Port-Security, Mgmt and Copy traffic

```

```

Module: 1
R-L Class          Config          Allowed          Dropped          Total
+-----+-----+-----+-----+-----+
L3 mtu             500             0                0                0
L3 ttl             500             0                0                0
L3 control         10000           0                0                0
L3 glean          100            3326            3190            6516
L3 mcast dirconn  3000            0                0                0
L3 mcast loc-grp   3000            0                0                0
L3 mcast rpf-leak  500             0                0                0
L2 storm-ctrl     Disable
access-list-log    100             0                0                0
copy               30000           1877             0                1877
receive            30000           318              0                318

```

When you look at the hardware programming, a drop index is installed:

```

N7K# show system internal forwarding vrf VRF_ABC ipv4 route 172.28.191.200 detail

slot 1
=====

RPF Flags legend:
  S - Directly attached route (S_Star)
  V - RPF valid
  M - SMAC IP check enabled
  G - SGT valid
  E - RPF External table valid
172.28.191.200/32 , Drop
Dev: 0 , Idx: 0x1a5 , Prio: 0x8b61 , RPF Flags: VS , DGT: 0 , VPN: 9
RPF_Intf_5: Vlan1601 (0x19 )
AdjIdx: 0x8913 , LIFB: 0 , LIF: Drop (0x0 ), DI: 0x0
DMAC: 0000.0000.0000 SMAC: 0000.0000.0000
172.28.191.200/32 , Drop
Dev: 1 , Idx: 0x1a5 , Prio: 0x8b61 , RPF Flags: VS , DGT: 0 , VPN: 9
RPF_Intf_5: Vlan1601 (0x19 )
AdjIdx: 0x8913 , LIFB: 0 , LIF: Drop (0x0 ), DI: 0x0
DMAC: 0000.0000.0000 SMAC: 0000.0000.0000
172.28.191.200/32 , Drop
Dev: 2 , Idx: 0x1a5 , Prio: 0x8b61 , RPF Flags: VS , DGT: 0 , VPN: 9
RPF_Intf_5: Vlan1601 (0x19 )
AdjIdx: 0x8913 , LIFB: 0 , LIF: Drop (0x0 ), DI: 0x0
DMAC: 0000.0000.0000 SMAC: 0000.0000.0000
172.28.191.200/32 , Drop
Dev: 5 , Idx: 0x1e1 , Prio: 0x88ee , RPF Flags: VS , DGT: 0 , VPN: 9

```

```
RPF_Intf_5:  Vlan1601      (0x19      )
AdjIdx: 0x8914 , LIFB: 0      , LIF: Drop      (0x0      ), DI: 0x0
DMAC: 0000.0000.0000 SMAC: 0000.0000.0000
```

You can now see that the hardware rate-limiter does not see any drops.

```
N7K# show hardware rate-limiter
```

```
Units for Config: packets per second
Allowed, Dropped & Total: aggregated since last clear counters
rl-1: STP and Fabricpath-ISIS
rl-2: L3-ISIS and OTV-ISIS
rl-3: UDLD, LACP, CDP and LLDP
rl-4: Q-in-Q and ARP request
rl-5: IGMP, NTP, DHCP-Snoop, Port-Security, Mgmt and Copy traffic
```

```
Module: 1
```

R-L Class	Config	Allowed	Dropped	Total
L3 mtu	500	0	0	0
L3 ttl	500	0	0	0
L3 control	10000	0	0	0
L3 glean	100	0	0	0
L3 mcast dirconn	3000	0	0	0
L3 mcast loc-grp	3000	0	0	0
L3 mcast rpf-leak	500	0	0	0
L2 storm-ctrl	Disable			
access-list-log	100	0	0	0
copy	30000	1877	0	1877
receive	30000	318	0	318

Related Information

- [Configuring IP Glean Throttling](#)
- [Technical Support & Documentation - Cisco Systems](#)