Contents

Overview

This document will help understand how the "hardware ip glean throttle" feature works with examples and the intention of this feature.

When forwarding 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 to generate an ARP request. Once the ARP request is responded to the supervisor resolves the MAC address for the next hop and programs the hardware.

If the supervisor cannot resolve the ARP entry then the linecard keeps sending all packets destined to that address to the supervisor. The supervisor will keep generating ARP requests indefinitely until the ARP entry is resolved. There is a hardware rate limiter called glean in place 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, a single packet is sent to the CPU for each destination IP to generate the ARP request. Then the software adds a /32 drop adjacency in the hardware to prevent additional packets to the same next-hop IP address to be forwarded to the supervisor. While this drop adjacency is added all 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 will again be sent to the CPU and the process repeats. The number of entries being 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.

Prerequisites

Requirements

Ensure that you meet these requirements before you attempt this configuration:

- Have a 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.

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. **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 will be throttling the traffic to 100 packets per second, per module. You can see that some of the packets are getting 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 in installed in the RIB

```
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 looking 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 is not seeing 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
```

Further Reading

[Configuring IP Glean Throttling]