Contents

Introduction
Background
Root Cause
Solution

Introduction

This document describes how to troubleshoot the ARP storm, without any inband ARP traffic.

Background

ARP storm is a common denial-of-service (DoS) attack you would see in the data center enviroment.

The common switch logic to handle ARP packet is that:

- ARP packet with broadcast destination Media Access Control (MAC)
- ARP packet with unicast destination MAC, which belongs to the switch

will be processed by ARP process in software if the Switch Virtual Interface (SVI) is up in the receiving Vlan.

By this logic, if there is one or more malicous hosts keep sending ARP request in a Vlan, where a switch is the gateway of that Vlan. The ARP request will be processed in software hence causes the switch being overwhelmed. In some older Cisco switch model and version, you will see that ARP process takes the CPU usage up to high level and system is too busy to handle other control plane traffic.The common way to trace such attack is to run inband capture to identify the source MAC of the ARP storm.

In the data center where Nexus 7000 is acting as the aggregation gateway, such impact is reduced by CoPP on Nexus 7000 Series Switches. You could still run inband capture Ethanalyzer on Nexus 7000 Troubleshooting Guide to identify the source MAC of the ARP storm since Control Plane Policing (CoPP) is just a bandit slowing down but not elminating the ARP storm rushing to the CPU.

How about this scenario where:

- SVI is down
- No excessive ARP packet being punt to CPU
- No high CPU due to ARP process

The switch however still see ARP related problem, e.g. direct connected host has incomplete ARP. Is it possiblely caused by ARP storm?

The answer is yes on Nexus 7000.

Root Cause

In the nexus 7000 linecard design, to support ARP packet process in CoPP, ARP request will drive a special logical interface (LIF) then be rate limited by CoPP in forwarding engine (FE). This happens no matter you have a SVI up for the Vlan or not.

Hence, while the final forwarding decision made by FE is to not send the ARP request to inband CPU (in the case no SVI up for the vlan), the CoPP counter is still updated. It leads to CoPP saturated with excessive ARP request and dropping legitimate ARP request/reply. In this scenario, you will not see any excessive inband ARP packets but still being affected by ARP storm.

We have an enhanced bug CSCub47533 filed for this CoPP day one behavior.

Solution

There could be a few options to identify the source of ARP storm in this scenario. One effective option is:

- First identify which module the ARP storm comes from

```
N7K# sh policy-map interface control-plane class copp-system-p-class-normal
Control Plane
service-policy input copp-system-p-policy-strict

class-map copp-system-p-class-normal (match-any)
match access-group name copp-system-p-acl-mac-dot1x
match exception ip multicast directly-connected-sources
match exception ipv6 multicast directly-connected-sources
match protocol arp
set cos 1
police cir 680 kbps bc 250 ms
conform action: transmit
violate action: drop
 module 3:
conformed 4820928 bytes,
5-min offered rate 0 bytes/sec
peak rate 104 bytes/sec at Thu Aug 25 08:12:12 2016
 violated 9730978848 bytes,
       5-min violate rate 6983650 bytes/sec
       peak rate 7632238 bytes/sec at Thu Aug 25 00:43:33 2016
module 4:
conformed 4379136 bytes,
5-min offered rate 0 bytes/sec
peak rate 38 bytes/sec at Wed Aug 24 07:12:09 2016
violated 0 bytes,
5-min violate rate 0 bytes/sec
peak rate 0 bytes/sec
...
```

- Second use ELAM Procedure to capture all the ARP packet hitting the module. You might need to do it several times. But if there is a storm going on, the chance you capture the violate ARP packet is much better than legetimate ARP packet. Identify the source MAC and Vlan from the ELAM capture.