CoPP on Nexus 7000 Series Switches

TAC **Document ID: 116043**

Contributed by Viral Bhutta, Cisco TAC Engineer.
Sep 04, 2014

Contents

Introduction
Prerequisites
  Requirements
  Components Used
CoPP on the Nexus 7000 Series Switch Overview
Why CoPP on the Nexus 7000 Series Switch
Control Plane Processing on the Nexus 7000 Series Switch
CoPP Best Practices Policy
  How to Customize a CoPP Policy
CoPP Data Structure
CoPP Scale Factor
CoPP Monitoring and Management
  CoPP Counters
ACL Counters
CoPP Configuration Best Practices
CoPP Monitoring Best Practices
Conclusions
Unsupported Features

Introduction

This document describes what, how, and why Control Plane Policing (CoPP) is used on the Nexus 7000 Series Switches, which include the F1, F2, M1, and M2 Series Modules and line cards (LCs). It also includes best practice policies, as well as how to customize a CoPP policy.

Prerequisites

Requirements

Cisco recommends that you have knowledge of Nexus operating system CLI.

Components Used

The information in this document is based on the Nexus 7000 Series Switches with Supervisor 1 Module.

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.

CoPP on the Nexus 7000 Series Switch Overview

The CoPP is critical to network operation. A Denial of Service (DoS) attack to the Control/Management Plane, which can be perpetrated either inadvertently or maliciously, typically involves high rates of traffic that result in excessive CPU utilization. The Supervisor module spends an inordinate amount of time handling the packets.

Examples of such attacks include:

- Internet Control Message Protocol (ICMP) echo requests.
- Packets sent with ***ip−options*** set.

This can lead to:

- Loss of keep−alive messages and routing protocol updates.
- Filling of packet queues, which results in indiscriminate drops.
- Slow or unresponsive interactive sessions.

Attacks can overwhelm the network stability and availability and lead to business−impacting network outages.

CoPP is a hardware−based feature that protects the Supervisor from DoS attacks. It controls the rate at which packets are allowed to reach the Supervisor. The CoPP feature is modeled like an input QoS policy attached to the special interface called the ***control−plane***. However, CoPP is a security feature and not part of QoS. In order to protect the Supervisor, the CoPP separates data plane packets from the control plane packets (Exception Logic). It identifies DoS attack packets from valid packets (Classification). CoPP allows for classification of these packets:

- Receive packets
- Multicast packets
- Exception packets
- Redirect packets
- Broadcast MAC + non−IP packets
- Broadcast MAC + IP packets (See Cisco Bug ID CSCub47533 − Packets in L2 Vlan (No SVI) hitting CoPP)
- Mcast MAC + IP packets
- Router MAC + non−IP packets
- ARP packets

After a packet is classified, the packet can also be marked and used to assign different priorities based on the type of packets. Conform, exceed, and violate actions (transmit, drop, mark−down) can be set. If no policer is attached to a class, then a default policer is added whose conform action is drop. Glean packets are policed with default−class. One rate, two color, and two rate, three color policing are supported.

Traffic that hits the CPU on the Supervisor module can come in through four paths:

1. Inband interfaces (front panel port) for traffic sent by line cards.
2. Management Interface (mgmt0) used for management traffic.
3. Control and Monitoring Processor (CMP) interface used for the console.
4. Switched Ethernet Out Band Channel (EOBC) to control the line cards from the Supervisor module and exchange status messages.

reaches the Supervisor module through the forwarding engines (FEs) on the line cards. The Nexus 7000 Series Switch implementation of CoPP is hardware−based only, which means that CoPP is not performed in software by the Supervisor module. CoPP functionality (policing) is implemented on each FE independently. When the various rates are configured for CoPP policy−map, consideration must be taken in regard to the number of line cards in the system.

The total traffic received by the Supervisor is N times X, where N is the number of FEs on the Nexus 7000 system, and X is the rate allowed for the particular class. The configured policer values apply on a per FE basis, and the aggregate traffic prone to hit the CPU is the sum of the conformed and transmitted traffic on all of the FEs. In other words, traffic that hits the CPU equals the configured conform rate multiplied by the number of FEs.

- N7K−M148GT−11/L LC has 1 FE
- N7K−M148GS−11/L LC has 1 FE
- N7K−M132XP−12/L LC has 1 FE
- N7K−M108X2−12L LC has 2 FE
- N7K−F248XP−15 LC has 12 FE (SOC)
- N7K−M235XP−23L LC has 2 FE
- N7K−M206FQ−23L LC has 2 FE
- N7K−M202CF−23L LC has 2 FE

CoPP configuration is only implemented in the default virtual device context (VDC); however, the CoPP policies are applicable for all VDCs. The same global policy is applied for all line cards. CoPP applies resource sharing between VDCs if ports of the same FEs belong to different VDCs (M1 Series or M2 Series LC). For example, ports of one FE, even in different VDCs, count against the same threshold for CoPP.

If the same FE is shared between different VDCs and a given class of control plane traffic exceeds the threshold, this affects all VDCs on the same FE. It is recommended to dedicate one FE per VDC in order to isolate CoPP enforcement, if possible.

When the switch comes up first time, the default policy must be programmed to protect the ***control−plane***. CoPP provides the default policies, which are applied to ***control−plane*** as part of the initial startup sequence.

Why CoPP on the Nexus 7000 Series Switch

The Nexus 7000 Series Switch is deployed as an aggregation or core switch. Hence, it is the ear and brain of the network. It handles the maximum load in the network. It must handle frequent and burst requests. Some of requests include:

- ***Spanning Tree Bridge Protocol Data Unit (BPDU) Processing*** − Default is every two seconds.
- ***First Hop Redundancy*** − This includes Hot Standby Router Protocol (HSRP), Virtual Router Redundancy Protocol (VRRP), and Gateway Load Balancing Protocol (GLBP) − Default is every three seconds.
- ***Address Resolution*** − This includes Address Resolution Protocol/Neighbor−Discovery (ARP/ND), Forwarding Information Base (FIB) Glean − Up to one request per second, per host , such as network interface controller (NIC) teaming.
- ***Dynamic Host Control Protocol (DHCP)*** − DHCP Request, Relay − Up to one request per second, per host.
- ***Routing Protocols*** for Layer 3 (L3).
- ***Data Center Interconnect*** − Overlay Transport Virtualization (OTV), Multiprotocol Label Switching (MPLS), and Virtual Private LAN Service (VPLS).

allows the CPU to have enough cycle to process critical control plane messages.

Control Plane Processing on the Nexus 7000 Series Switch

The Nexus 7000 Series Switch takes a distributed control plane approach. It has a multi−core on each I/O module, as well as a multi−core for switch control plane on the Supervisor module. It offloads intensive tasks to the I/O module CPU for access control lists (ACL) and FIB programming. It scales the control plane capacity with the number of line cards. This avoids Supervisor CPU bottleneck, which is seen in a centralized approach. Hardware rate limiters and hardware−based CoPP protects the control plane from bad or malicious activity.

CoPP Best Practices Policy

CoPP Best Practices Policy (BPP) was introduced in Cisco NX−OS Release 5.2. The ***show running−config*** command output does not display the content of the CoPP BPP. The ***show run all*** command displays the content of CoPP BPP.

```
------------------------------------SNIP-----------------------------------------
SITE1-AGG1# show run copp

!! Command: show running-config copp
!! Time: Mon Nov  5 22:21:04 2012

version 5.2(7)
copp profile strict


SITE1-AGG1# show run copp all

!! Command: show running-config copp all
!! Time: Mon Nov  5 22:21:15 2012

version 5.2(7)
------------------------------SNIP---------------------
control-plane
  service-policy input copp-system-p-policy-strict
copp profile strict
```

CoPP provides four options to the user for default policies:

- Strict
- Moderate
- Lenient
- Dense (introduced in Release 6.0(1))

If no option is selected or if set up is skipped, then strict policing is applied. All of these options use the same class−maps and classes, but different Committed Information Rate (CIR) and Burst Count (BC) values for policing. In Cisco NX−OS releases earlier than 5.2.1, the ***setup*** command was used to change the option. Cisco NX−OS Release 5.2.1 introduced an enhancement to the CoPP BPP so that the option can be changed without the ***setup*** command; use the ***copp profile*** command.

```
SITE1-AGG1# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
SITE1-AGG1(config)# copp profile ?
  dense     The Dense   Profile
  lenient   The Lenient Profile
  moderate  The Moderate Profile
```

```
SITE1-AGG1(config)# copp profile strict
SITE1-AGG1(config)# exit
```

Use the ***show copp profile <profile−type>*** command to view the default CoPP BPP configuration. Use the ***show copp status*** command to verify that the CoPP policy has been applied correctly.

```
SITE1-AGG1# show copp status
Last Config Operation: copp profile strict
Last Config Operation Timestamp: 20:40:27 PST Nov  5 2012
Last Config Operation Status: Success
Policy-map attached to the control-plane: copp-system-p-policy-strict
```

In order to view the difference between two CoPP BPPs, use the ***show copp diff profile <profile−type 1> profile <profile−type 2>*** command:

```
SITE1-AGG1# show copp diff profile strict profile moderate
A '+' represents a line that has been added and
a '-' represents a line that has been removed.
    -policy-map type control-plane copp-system-p-policy-strict
    -  class copp-system-p-class-critical
    -    set cos 7
    -    police cir 39600 kbps bc 250 ms conform transmit violate drop
    -  class copp-system-p-class-important
    -    set cos 6
    -    police cir 1060 kbps bc 1000 ms conform transmit violate drop
----------------------SNIP---------------------------------------
    +policy-map type control-plane copp-system-p-policy-moderate
    +  class copp-system-p-class-critical
    +    set cos 7
    +    police cir 39600 kbps bc 310 ms conform transmit violate drop
    +  class copp-system-p-class-important
    +    set cos 6
    +    police cir 1060 kbps bc 1250 ms conform transmit violate drop
----------------------SNIP---------------------------------------
```

How to Customize a CoPP Policy

Users can create a customized CoPP policy. Clone the default CoPP BPP, and attach it to the ***control−plane*** interface because the CoPP BPP is read−only.

```
SITE2-AGG1(config)# policy-map type control-plane copp-system-p-policy-strict
                                                                             ^
% String is invalid, 'copp-system-p-policy-strict' is not an allowed string at
   '^' marker.
```

The ***copp copy profile <profile−type> <prefix> [suffix]*** command creates a clone of the CoPP BPP. This is used in order to modify the default configurations. The ***copp copy profile*** command is an ***exec mode*** command. User can choose a prefix or suffix for the access−list, class−maps, and policy−map name. For instance, ***copp−system−p−policy−strict*** is changed to ***[prefix]copp−policy−strict[suffix]***. Cloned configurations are treated as user configurations and are included in the ***show run*** output.

```
SITE1-AGG1# copp copy profile ?
  dense      The Dense   Profile
  lenient    The Lenient Profile
  moderate   The Moderate Profile
  strict     The Strict Profile
SITE1-AGG1# copp copy profile strict ?
  prefix  Prefix for the copied policy
  suffix  Suffix for the copied policy
SITE1-AGG1# copp copy profile strict suffix ?
```

Explore Litigation Insights

Docket Alarm provides insights to develop a more informed litigation strategy and the peace of mind of knowing you're on top of things.

Real-Time Litigation Alerts

Keep your litigation team up-to-date with **real-time alerts** and advanced team management tools built for the enterprise, all while greatly reducing PACER spend.

Our comprehensive service means we can handle Federal, State, and Administrative courts across the country.

Advanced Docket Research

With over 230 million records, Docket Alarm's cloud-native docket research platform finds what other services can't. Coverage includes Federal, State, plus PTAB, TTAB, ITC and NLRB decisions, all in one place.

Identify arguments that have been successful in the past with full text, pinpoint searching. Link to case law cited within any court document via Fastcase.

Analytics At Your Fingertips

Learn what happened the last time a particular judge, opposing counsel or company faced cases similar to yours.

Advanced out-of-the-box PTAB and TTAB analytics are always at your fingertips.

API

Docket Alarm offers a powerful API (application programming interface) to developers that want to integrate case filings into their apps.

LAW FIRMS

Build custom dashboards for your attorneys and clients with live data direct from the court.

Automate many repetitive legal tasks like conflict checks, document management, and marketing.

FINANCIAL INSTITUTIONS

Litigation and bankruptcy checks for companies and debtors.

E-DISCOVERY AND LEGAL VENDORS

Sync your system to PACER to automate legal marketing.

WHAT WILL YOU BUILD? | **sales@docketalarm.com** | **1-866-77-FASTCASE**