Control Plane Protection in EOS

NLNOG Live - Joris Claassen
joris@arista.com

Control Plane Protection in Arista EOS

Consists out of the following features:

- Control Plane ACL
  - Protects using the kernel
- Control Plane Policers
  - Rate-limits any traffic destined to the the CPU
- CPU Traffic Policy
  - Protects using the data plane (TCAM) - before traffic hits the kernel

Access Control List

- Software-based
  - ACL entries → iptables
  - Dynamic BGP neighbor ACL
    - ≫ GTSM
- Applies to all traffic hitting the kernel
- Allows the following protocols by default:

Telnet	IGMP
SSH	OSPF
HTTP(S)	BGP
BOOTP	VRRP
SNMP	PIM
ICMP	AHP
MLAG	VNC

```
IP Access List default-control-plane-acl [readonly]
        counters per-entry
        10 permit icmp any any
        20 permit ip any any tracked
        30 permit udp any any eq bfd ttl eq 255
        40 permit udp any any eq bfd-echo ttl eq 254
        50 permit udp any any eq multihop-bfd
        60 permit udp any any eq micro-bfd
        70 permit udp any any eq sbfd
        80 permit udp any eq sbfd any eq sbfd-initiator
        90 permit ospf any any
        100 permit tcp any any eq ssh telnet www snmp
      bgp https msdp ldp netconf-ssh gnmi
        110 permit udp any any eq bootps bootpc snmp rip
      ntp ldp
        120 permit tcp any any eq mlag ttl eq 255
        130 permit udp any any eq mlag ttl eq 255
        140 permit vrrp any any
        150 permit ahp any any
        160 permit pim any any
        170 permit igmp any any
        180 permit tcp any any range 5900 5910
        190 permit tcp any any range 50000 50100
        200 permit udp any any range 51000 51100
        210 permit tcp any any eq 3333
        220 permit tcp any any eq nat ttl eq 255
        230 permit tcp any eq bgp any
        240 permit rsvp any any
        250 permit tcp any any eq 6040
```

Nmap results

```
Starting Nmap 7.70 ( https://nmap.org ) at 2020-09-02 21:11 BST
Nmap scan report for switch
Host is up (0.00026s latency).
Not shown: 9977 filtered ports
PORT     STATE  SERVICE
22/tcp   open   ssh
23/tcp   closed telnet
80/tcp   open   http
161/tcp  open   snmp
179/tcp  closed bgp
443/tcp  open   https
639/tcp  closed msdp
646/tcp  closed ldp
830/tcp  closed netconf-ssh
3333/tcp closed dec-notes
5900/tcp closed vnc
5901/tcp closed vnc-1
5902/tcp closed vnc-2
5903/tcp closed vnc-3
5904/tcp closed unknown
5905/tcp closed unknown
5906/tcp closed unknown
5907/tcp closed unknown
5908/tcp closed unknown
5909/tcp closed unknown
5910/tcp closed cm
6030/tcp closed x11
6040/tcp closed x11
MAC Address: 00:1C:73:EE:E8:BA (Arista Networks)
Device type: general purpose
Running: Linux 3.X|4.X
OS CPE: cpe:/o:linux:linux_kernel:3 cpe:/o:linux:linux_kernel:4
OS details: Linux 3.2 - 4.9
Network Distance: 1 hop

OS detection performed. Please report any incorrect results at https://nmap.org/submit/ .
Nmap done: 1 IP address (1 host up) scanned in 188.41 seconds
```

Policing

- Hardware-based
  - Exact implementation differs slightly per chipset family
- Traffic classification
  - Global policy-map
    - Per protocol class-maps
- Rate-limiting
  - Minimum reserved bandwidth
  - Maximum allowed bandwidth

something like this

```
Service-policy input: copp-system-policy
  Class-map: copp-system-bpdu (match-any)
  Class-map: copp-system-lacp (match-any)
  Class-map: copp-system-bfd (match-any)
  Class-map: copp-system-mlag (match-any)
<...>
```

CPU Traffic Policy

- Hardware-based
  - Platform dependent; EOS 4.22.0+
  - Supported by all “full-table” capable devices
- Rejecting, policing or counting traffic
  - Before it hits the kernel
- Allowing traffic from trusted sources
  - Automatic policies deferred from BGP neighbor configuration
  - Deduplication for efficient TCAM usage
  - VRF aware

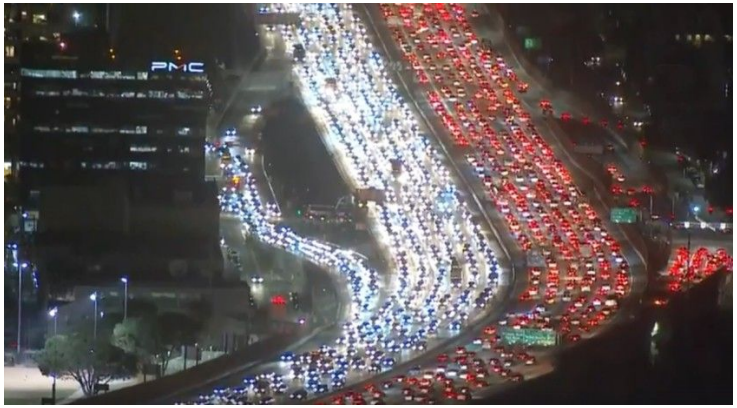

CPU Traffic Policy

To optimize TCAM resources, a default permit rule is never installed in the TCAM, as the default behavior for a packet processed by the TCAM is to be permitted.

```
traffic-policies
   cpu traffic-policy CPU-DEFAULT vrf all
   traffic-policy CPU-DEFAULT
      match ICMPV6 ipv6
         protocol icmpv6
      match BFD-DPORT ipv4
         protocol udp destination port 3784-3785
      match BFD-SPORT ipv4
         protocol udp source port 49152
      match BFD-DPORT6 ipv6
         protocol udp destination port 3784-3785
      match BFD-SPORT6 ipv6
         protocol udp source port 49152
      match BGP ipv4
         protocol neighbors bgp
      match BGP-OTHER ipv4
         protocol bgp
         actions
            drop
```

protocol tcp udp destination port 179

for each configured neighbor, protocol tcp source + destination port 179

```
match BGP6 ipv6
   protocol neighbors bgp
match BGP-OTHER6 ipv6
   protocol bgp
   actions
      drop
match OSPF ipv4
   protocol ospf
match PIM4 ipv4
   protocol pim
match PIM6 ipv6
   protocol pim
match ipv4-all-default ipv4
   actions
      drop
match ipv6-all-default ipv6
   actions
      drop
```

Per BGP-peer policing and “show” commands

```
traffic-policies
   cpu traffic-policy CPU vrf all
   traffic-policy CPU
      match BGP ipv4
         protocol neighbors bgp
         actions
            police rate 15 kbps
```

```
router bgp 65000
   <...>
   vrf vrf1
      local-as 65101
      neighbor 1.0.0.1 remote-as 65102
      neighbor 2.0.0.1 remote-as 65103
      <...>
   !
   vrf vrf2
      local-as 65201
      neighbor 1.0.0.1 remote-as 65202
      redistribute connected include leaked
      <...>
```

```
switch(config-traffic-policies)#show traffic-policy protocol neighbors bgp
Traffic Policy CPU-policy
      VRF             Source IP           Source Port    Destination Port
--------------- ---------------------- ----------------- ----------------
   vrf1               1.0.0.1/32              any              179
   vrf1               1.0.0.1/32              179              any
   vrf1               2.0.0.1/32              any              179
   vrf1               2.0.0.1/32              179              any
   vrf2               1.0.0.1/32              any              179
   vrf2               1.0.0.1/32              179              any
<...>
```

Polices ***each BGP neighbor*** to a maximum of 15 kbps (per neighbor)

Resources

* - requires login with Arista account (not a device!)

CoPP resources:

https://www.arista.com/en/um-eos/eos-traffic-management

https://eos.arista.com/eos-4.17.1f/gtsm-for-bgp/ *

https://eos.arista.com/eos-4-22-0f/support-for-cpu-traffic-policy/ *

https://eos.arista.com/troubleshooting-based-on-control-plane-policing-copp-for-sand-platform/ *

EOS Hardening guide:

https://eos.arista.com/arista-eos-hardening-guide/ *

Peering resources:

https://eos.arista.com/inet-edge-config/ *

https://eos.arista.com/bgp-primer-for-internet-peering *