Troubleshoot Packet Drops with ACLs on Nexus Platform

Contents

Introduction

This document describes how to troubleshoot packet loss using Access Control Lists (ACLs) on Nexus platform.

Prerequisites

Requirements

Cisco recommendsrecomends that you have knowledge of these topics:

- NXOS Platform
- Access Control Lists

Component Used

N9K1	N9K-C93108TC-EX	9.3(10)
N9K2	N9K-C93108TC-EX	9.3(10)
N9K3	N9K-C93108TC-EX	9.3(10)

The information in this document was created from Nexus devices in a lab environment. All of the devices used in this document started without any pre-existing configuration. If you are using a live network, make sure that you understand the potential impact of any command.

Topology

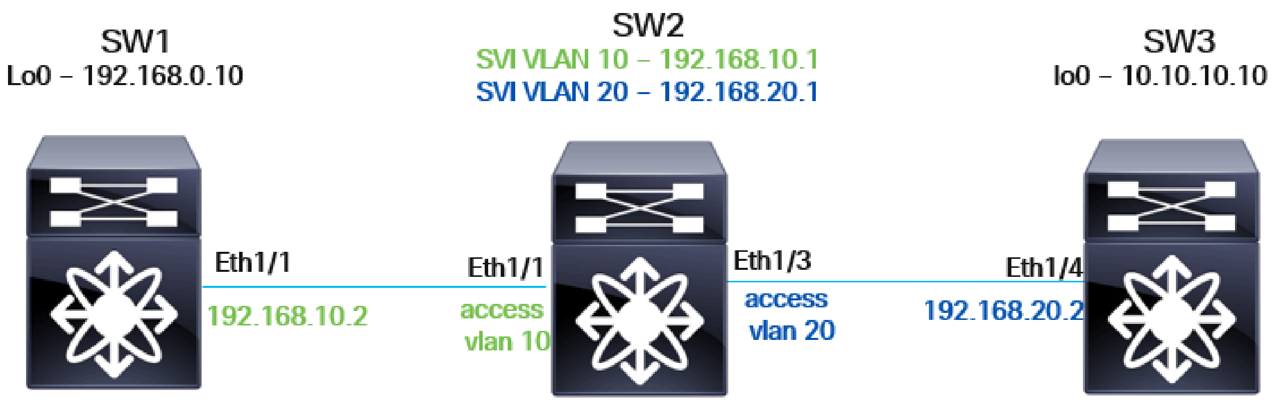

Brief Overview of Access Control Lists and their Functionality

An ACL is essentially used to filter traffic based on a series of ordered rules and criteria (for instance, filtering based on source/destination IP addresses.) These rules determine whether packets match specific conditions in order to decide if they shall be permitted or denied. In simpler terms, the ACL defines whether network packets can be allowed to pass through or be denied based on the rules set within it. If the packets meet the conditions of the permit rules, they shall be processed by the Nexus switch. Conversely, if the packets match the deny conditions, they shall be discarded.

One key feature of ACLs is their ability to provide statistical counters for the packet flow. These counters track the number of packets that match the ACL rules, which can be very useful when troubleshooting packet loss scenarios.

For instance, if a device is sending a certain number of packets, but receiving fewer than expected, the statistical counters from the ACL can assist in isolating the point at which packets are being dropped within the network.

PACL and RACL

The implementation of ACLs can vary depending on whether they are applied to Layer 2 interfaces (PACL), Layer 3 interfaces (RACL), or VLANs (VACL). Here is a brief comparison of these methods:

- Port Access Control List (PACL): The ACL is applied to a Layer 2 (L2) switchport interface.
- Router Access Control List (RACL): The ACL is applied to a Layer 3 (L3) routed interface.

ACL Type	Interface	Action	Applied Direction
PACL	L2	Switchport interfaces If the ACL is applied to a trunk interface, it filters traffic for all VLANs allowed on the trunk.	Inbound only - traffic coming into the interface.
RACL	L3	SVI, Physical L3, and L3 Subinterfaces	Both inbound and outbound - Inbound filters traffic coming into the interface, while outbound filters traffic leaving the interface.

Objective

It is necessary to confirm that all the packets being sent are received properly.

Topology Explanation

- N9K-1 has L3 connectivity with N9K-2. The Eth1/1 interface on N9K-1 is configured as a L3 routed interface, while N9K-2's Eth1/1 is a L2 switchport interface, tagged with VLAN 10.
- N9K-2 also has L3 connectivity with N9K-3. The Eth1/3 interface on N9K-2 is a L2 switchport interface tagged with VLAN 20, and N9K-3's Eth1/4 is configured as a L3 routed interface.
- Loopback Configuration: Both N9K-1 and N9K-2 have the Lo0 interface configured. These Lo0 interfaces shall be used to send ICMP ping packets between the two devices.

Troubleshooting

Please find the detailed process steps for configuring and verifying RACL and PACL on the N9K devices. During this process,the Port Access Control Lists and Router Access Control Lists are reviewed to analyze the packet flow and determine whether all packets are being transmitted and received correctly.

Step 1. Configure the RACL on L3 Interfaces of N9K-1 (Eth1/1), N9K-2 (SVI 10, SVI 20), and N9K-3 (Eth1/14)

Note: To observe the outbound packet flow, an additional ACL configuration is needed on N9K-2. Since N9K-2 lacks L3 physical routed interfaces (instead, it has SVI and L2 switchport interfaces), PACL only supports inbound traffic.

To capture outbound packet matches, a new ACL can be created and applied to the L3 interfaces.

The ACL shall be applied to N9K-1, N9K-2 and N9K-3.

```
ip access-list TAC-IN
statistics per-entry
10 permit ip 192.168.20.2/32 192.168.0.10/32
20 permit ip 192.168.0.10/32 192.168.20.2/32
30 permit ip any any


ip access-list TAC-OUT
statistics per-entry
10 permit ip 192.168.20.2/32 192.168.0.10/32
20 permit ip 192.168.0.10/32 192.168.20.2/32
30 permit ip any any
```

```
***N9K-1***
interface Ethernet1/1
description ***Link-to-N9K-2***
ip access-group TAC-IN in
ip access-group TAC-OUT out
ip address 192.168.10.2/30
no shutdown



***N9K-2***

interface Vlan10
no shutdown
ip access-group TAC-IN-SVI in
ip access-group TAC-OUT-SVI out
ip address 192.168.10.1/30

interface Vlan20
no shutdown
ip access-group TAC-IN-SVI in
ip access-group TAC-OUT-SVI out
ip address 192.168.20.1/30

***N9K-3***

interface Ethernet1/4
description ***Link-to-N9K-2***
ip access-group TAC-IN in
ip access-group TAC-OUT out
ip address 192.168.20.2/30
no shutdown
```

Step 2. Configure PACL on L2 Switchport Interfaces of N9K-2

TCAM Carving

TCAM carving can be required depending on the ACL type, for more information please refer to:

Understand how to Carve Nexus 9000 TCAM Space

To apply the PACL to L2 physical interfaces, it is necessary to configure an ip port access-group .... However, configuring the TCAM region is also required.

Note: Certain rows have been removed to keep the output clean.

```
N9K-C93180YC-2# conf
Enter configuration commands, one per line. End with CNTL/Z.
N9K-C93180YC-2(config)# int e1/2
N9K-C93180YC-2(config-if)# ip port access-group TAC-IN in
ERROR: TCAM region is not configured. Please configure TCAM region Ingress PACL [ing-ifacl] and retry t
N9K-C93180YC-2(config-if)#
```

Procedure for Configuring the TCAM Region

Step 1. TCAM Region Modifications

Please evaluate which region can provide free space, as this can differ for each environment.

```
N9K-C93180YC-2# show system internal access-list globals

slot 1
=======
<snip>

LOU Threshold Value : 5

--------------------------------------------------------------------------------------
INSTANCE 0 TCAM Region Information:
--------------------------------------------------------------------------------------
Ingress:
--------
Region TID Base Size Width
--------------------------------------------------------------------------------------
NAT 13 0 0 1
Ingress PACL 1 0 0 1 >>>>>>> Size of 0
Ingress VACL 2 0 0 1
Ingress RACL 3 0 1792 1
Ingress RBACL 4 0 0 1
Ingress L2 QOS 5 1792 256 1
Ingress L3/VLAN QOS 6 2048 512 1 >>>>>> Size of 512
Ingress SUP 7 2560 512 1
Ingress L2 SPAN ACL 8 3072 256 1
Ingress L3/VLAN SPAN ACL 9 3328 256 1
Ingress FSTAT 10 0 0 1
SPAN 12 3584 512 1
Ingress REDIRECT 14 0 0 1
Ingress NBM 30 0 0 1
Ingress Flow-redirect 39 0 0 1
Ingress RACL Lite 42 0 0 1
Ingress PACL IPv4 Lite 41 0 0 1
Ingress PACL IPv6 Lite 43 0 0 1
Ingress CNTACL 44 0 0 1
Mcast NAT 46 0 0 1
Ingress DACL 47 0 0 1
Ingress PACL Super Bridge 49 0 0 1
Ingress Storm Control 50 0 0 1
Ingress VACL Redirect 51 0 0 1
Egress Netflow L3 56 0 0 1
55 0 0 1
-------------------------------------------------------------------------------------
Total configured size: 4096
Remaining free size: 0
Note: Ingress SUP region includes Redirect region

<snip>
```

An Alternative Method for Verification.

```
N9K-C93180YC-2# sh hardware access-list tcam region
NAT ACL[nat] size = 0
Ingress PACL [ing-ifacl] size = 0 >>>>>>> Size of 0
VACL [vacl] size = 0
Ingress RACL [ing-racl] size = 1792
Ingress L2 QOS [ing-l2-qos] size = 256
Ingress L3/VLAN QOS [ing-l3-vlan-qos] size = 512 >>>>>> Size of 512
Ingress SUP [ing-sup] size = 512
Ingress L2 SPAN filter [ing-l2-span-filter] size = 256
Ingress L3 SPAN filter [ing-l3-span-filter] size = 256
Ingress FSTAT [ing-fstat] size = 0
span [span] size = 512
Egress RACL [egr-racl] size = 1792
Egress SUP [egr-sup] size = 256
Ingress Redirect [ing-redirect] size = 0
Egress L2 QOS [egr-l2-qos] size = 0
Egress L3/VLAN QOS [egr-l3-vlan-qos] size = 0
Ingress Netflow/Analytics [ing-netflow] size = 0
Ingress NBM [ing-nbm] size = 0
TCP NAT ACL[tcp-nat] size = 0
Egress sup control plane[egr-copp] size = 0
Ingress Flow Redirect [ing-flow-redirect] size = 0
Ingress CNTACL [ing-cntacl] size = 0
Egress CNTACL [egr-cntacl] size = 0
MCAST NAT ACL[mcast-nat] size = 0
Ingress DACL [ing-dacl] size = 0
Ingress PACL Super Bridge [ing-pacl-sb] size = 0
Ingress Storm Control [ing-storm-control] size = 0
Ingress VACL redirect [ing-vacl-nh] size = 0
Egress PACL [egr-ifacl] size = 0
Egress Netflow [egr-netflow] size = 0
N9K-C93180YC-2#
```

Step 2. Reduce the Size of the Region

Reduce the size of the region allocated for ing-l3-vlan-qos. (This differs for each environment.)

N9K-C93180YC-2(config)# hardware access-list tcam region ing-l3-vlan-qos 256 >>> Reduce the allocation from 512 to 256.
Please save config and reload the system for the configuration to take effect.

Step 3. Increase the TCAM Region for ing-ifacl

N9K-C93180YC-2(config)# hardware access-list tcam region ing-ifacl 256

Save config and reload the system for the configuration to take effect.

N9K-C93180YC-2(config)#

Step 4. Save Configuration

```
N9K-C93180YC-2(config)# copy running-config startup-config
[########################################] 100%
Copy complete, now saving to disk (please wait)...
Copy complete.
N9K-C93180YC-2(config)#
```

Step 5. Reload

```
N9K-C93180YC-2(config)# reload
This command will reboot the system. (y/n)? [n] y
```

Post-Reload Verification

After reloading, check if the changes have taken effect.

```
N9K-C93180YC-2# sh system internal access-list globals

slot 1
=======
<snip>

--------------------------------------------------------------------------------------
INSTANCE 0 TCAM Region Information:
--------------------------------------------------------------------------------------
Ingress:
--------
Region TID Base Size Width
--------------------------------------------------------------------------------------
NAT 13 0 0 1
Ingress PACL 1 0 256 1 >>> The size value is now 256.
Ingress VACL 2 0 0 1
Ingress RACL 3 256 1792 1
Ingress RBACL 4 0 0 1
Ingress L2 QOS 5 2048 256 1
Ingress L3/VLAN QOS 6 2304 256 1 >>> The size value is now 256.
Ingress SUP 7 2560 512 1
Ingress L2 SPAN ACL 8 3072 256 1
Ingress L3/VLAN SPAN ACL 9 3328 256 1
Ingress FSTAT 10 0 0 1
SPAN 12 3584 512 1
Ingress REDIRECT 14 0 0 1
Ingress NBM 30 0 0 1
Ingress Flow-redirect 39 0 0 1
Ingress RACL Lite 42 0 0 1
Ingress PACL IPv4 Lite 41 0 0 1
Ingress PACL IPv6 Lite 43 0 0 1
Ingress CNTACL 44 0 0 1
Mcast NAT 46 0 0 1
Ingress DACL 47 0 0 1
Ingress PACL Super Bridge 49 0 0 1
Ingress Storm Control 50 0 0 1
Ingress VACL Redirect 51 0 0 1
Egress Netflow L3 56 0 0 1
55 0 0 1
-------------------------------------------------------------------------------------
Total configured size: 4096
```

```
Remaining free size: 0
Note: Ingress SUP region includes Redirect region

<snip>
```

An Alternative Method for Verification.

```
N9K-C93180YC-2# sh hardware access-list tcam region
NAT ACL[nat] size = 0
Ingress PACL [ing-ifacl] size = 256 >>> The size value is now 256.
VACL [vacl] size = 0
Ingress RACL [ing-racl] size = 1792
Ingress L2 QOS [ing-l2-qos] size = 256
Ingress L3/VLAN QOS [ing-l3-vlan-qos] size = 256 >>> The size value is now 256.
Ingress SUP [ing-sup] size = 512
Ingress L2 SPAN filter [ing-l2-span-filter] size = 256
Ingress L3 SPAN filter [ing-l3-span-filter] size = 256
Ingress FSTAT [ing-fstat] size = 0
span [span] size = 512
Egress RACL [egr-racl] size = 1792
Egress SUP [egr-sup] size = 256
Ingress Redirect [ing-redirect] size = 0
Egress L2 QOS [egr-l2-qos] size = 0
Egress L3/VLAN QOS [egr-l3-vlan-qos] size = 0
Ingress Netflow/Analytics [ing-netflow] size = 0
Ingress NBM [ing-nbm] size = 0
TCP NAT ACL[tcp-nat] size = 0
Egress sup control plane[egr-copp] size = 0
Ingress Flow Redirect [ing-flow-redirect] size = 0
Ingress CNTACL [ing-cntacl] size = 0
Egress CNTACL [egr-cntacl] size = 0
MCAST NAT ACL[mcast-nat] size = 0
Ingress DACL [ing-dacl] size = 0
Ingress PACL Super Bridge [ing-pacl-sb] size = 0
Ingress Storm Control [ing-storm-control] size = 0
Ingress VACL redirect [ing-vacl-nh] size = 0
Egress PACL [egr-ifacl] size = 0
Egress Netflow [egr-netflow] size = 0
N9K-C93180YC-2#
```

Configuration of IP Port Access Group

Configure the ip port access-group on L2 physical interfaces.

```
N9K-C93180YC-2# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
N9K-C93180YC-2(config)# int e1/2,e1/51
N9K-C93180YC-2(config-if-range)# ip port access-group TAC-IN in
```

```
N9K-C93180YC-2(config-if-range)# ip port access-group TAC-OUT out
Port ACL is only supported on ingress direction >>>>>>>>
N9K-C93180YC-2(config-if-range)#
```

```
interface Ethernet1/1
description ***Link-to-N9K-1***
switchport
switchport access vlan 10
ip port access-group TAC-IN in >>> Inboud only
no shutdown
```

```
interface Ethernet1/3
description ***Link-to-N9K-3***
switchport
switchport access vlan 20
ip port access-group TAC-IN in >>> Inboud only
no shutdown
```

Step 3. Loopback

N9K-1 shall utilize its Loopback0 (Lo0) as the source, while N9K-3 can use its Loopback0 (Lo0) as the destination.
The running configuration of the Loopback interfaces that you employ for testing purposes is detailed as follows.

Note: Layer 3 connectivity with a routing protocol has been previously configured.

```
***N9K-1***
interface loopback0
ip address 192.168.0.10/32
```

```
***N9K-3***
interface loopback0
ip address 10.10.10.10/30
```

Step 4. Generate Traffic and Send a Ping from N9K-3 Using the Source IP 192.168.20.2 to Lo0 192.168.0.10 of N9K-1

```
N9K-3# ping 192.168.0.10 source 192.168.20.2
PING 192.168.0.10 (192.168.0.10) from 192.168.20.2: 56 data bytes
```

```
64 bytes from 192.168.0.10: icmp_seq=0 ttl=253 time=1.163 ms
64 bytes from 192.168.0.10: icmp_seq=1 ttl=253 time=0.738 ms
64 bytes from 192.168.0.10: icmp_seq=2 ttl=253 time=0.706 ms
64 bytes from 192.168.0.10: icmp_seq=3 ttl=253 time=0.668 ms
64 bytes from 192.168.0.10: icmp_seq=4 ttl=253 time=0.692 ms

--- 192.168.0.10 ping statistics ---
5 packets transmitted, 5 packets received, 0.00% packet loss
round-trip min/avg/max = 0.668/0.793/1.163 ms
N9K-3#
```

Step 5. Verify PACL and RACL Statistics Information on N9K-1, N9K-2, and N9K-3

- Since the ICMP packets are originating from N9K-3, it is necessary to verify that the five ICMP request packets have been received by N9K-2.
- PACL Verification on N9K-2: It is expected that five packets originating from 192.168.20.2 (Eth1/4 of N9K-3) are received, with the destination being N9K-1’s Lo0 (192.168.0.10).

```
N9K-2# show ip access-lists TAC-IN
IP access list TAC-IN
statistics per-entry
10 permit ip 192.168.20.2/32 192.168.0.10/32 [match=5] >>>
20 permit ip 192.168.0.10/32 192.168.20.2/32 [match=0]
30 permit ip any any [match=0]
```

Related configuration on Eth1/3 of N9K-2.

```
interface Ethernet1/3
description ***Link-to-N9K-3***
switchport
switchport access vlan 20
ip port access-group TAC-IN in >>> PACL
no shutdown
```

- On N9K-2, the RACL reports 5 ICMP request packets leaving N9K-2 and being forwarded to N9K-1.
- Since PACL does not support the outbound direction, it is essential to verify the other ACL (TAC-OUT-SVI) configured on the SVI for VLAN 10, which is configured as a RACL (since outbound direction is supported on RACLs). VLAN 10 provides the connectivity between N9K-2 and N9K-1.

```
N9K-2# show ip access-lists TAC-OUT-SVI

IP access list TAC-OUT-SVI
statistics per-entry
10 permit ip 192.168.20.2/32 192.168.0.10/32 [match=5] >>>
20 permit ip 192.168.0.10/32 192.168.20.2/32 [match=0]
30 permit ip any any [match=0]
```

```
configuration associated:


interface Vlan10
no shutdown
ip access-group TAC-IN-SVI in
ip access-group TAC-OUT-SVI out >>>>
ip address 192.168.10.1/30
```

Based on the previous results, it is confirmed that there is no packet loss with the ICMP request packets sent from N9K-3.

- The next step is to proceed to the next device (destination N9K-1) and verify that the same number of ICMP request packets are being received from N9K-3.
- The RACL statistics indicate that N9K-2 is sending out 5 ICMP request packets originating from N9K-3.

```
N9K-1# show ip access-lists TAC-IN
IP access list TAC-IN
statistics per-entry
10 permit ip 192.168.20.2/32 192.168.0.10/32 [match=5] >>>>
20 permit ip 192.168.0.10/32 192.168.20.2/32 [match=0]
30 permit ip any any [match=0]
```

Related configuration on Eth1/1 of N9K-1.

```
interface Ethernet1/1
description ***Link-to-N9K-2***
ip access-group TAC-IN in >>> RACL
ip access-group TAC-OUT out
ip address 192.168.10.2/30
no shutdown
```

- Based on the information, it is confirmed that there is no packet loss (ICMP request) from N9K-3 to Lo0 192.168.0.10 on N9K-2.
- The next step is to track the ICMP reply packets originating from N9K-1 Lo0 192.168.0.10 and destined for N9K-3 at 192.168.20.2.
- Then, it is then necessary to proceed to N9K-2 and verify whether it is receiving the five ICMP reply packets from 192.168.0.10 to 192.168.20.2.
- To track the ICMP reply packets from N9K-1, it is required to verify the PACL (TAC-IN) configured on Eth1/1.

```
N9K-2# show ip access-lists TAC-IN
IP access list TAC-IN
statistics per-entry
```

```
10 permit ip 192.168.20.2/32 192.168.0.10/32 [match=0]
20 permit ip 192.168.0.10/32 192.168.20.2/32 [match=5] >>> 5 icmp reply comming from 192.168.0.10 to 19
30 permit ip any any [match=0]


interface Ethernet1/1
description ***Link-to-N9K-1***
switchport
switchport access vlan 10
ip port access-group TAC-IN in >>> PACL (Inboud direction only)
no shutdown
```

- Based on the information previously provided, it is confirmed that there is no packet loss on the traffic from N9K-1 to N9K-2.
- The next step is to confirm that N9K-2 is properly sending the ICMP reply packets to N9K-3. Since PACL does not support the outbound direction, it is necessary to verify the other ACL (TAC-OUT-SVI) configured on the SVI for VLAN 20, which is configured as an RACL (as outbound direction is supported on RACLs). VLAN 20 provides the connectivity between N9K-2 and N9K-3.

```
N9K-2# show ip access-lists TAC-OUT-SVI
IP access list TAC-OUT-SVI
statistics per-entry
10 permit ip 192.168.20.2/32 192.168.0.10/32 [match=0]
20 permit ip 192.168.0.10/32 192.168.20.2/32 [match=5] >>> 5 ICMP reply packets are being sent out to N
```

Related configuration:

```
interface Vlan10
no shutdown
ip access-group TAC-IN-SVI in
ip access-group TAC-OUT-SVI out >>> RACL outboud direccion
ip address 192.168.20.1/30
```

Based on the ACL counters from above outputs, it is confirmed that N9K-1 is properly sending the five ICMP reply packets to N9K-2.

- There's no packet loss occurring from N9K-2 to N9K-3.

- The final step is to proceed to the source of the traffic, N9K-3, and verify whether it is receiving the five ICMP reply packets.
- It is confirmed that the five ICMP packets are hitting the ACL TAC-IN for the ICMP replies coming from N9K-1 Lo0 (192.168.0.10).
  To investigate further, it is necessary to review the RACL (TAC-IN) configured on Eth1/4..

```
N9K-3# sh ip access-lists TAC-IN

IP access list TAC-IN
statistics per-entry
10 permit ip 192.168.20.2/32 192.168.0.10/32 [match=0]
```

```
20 permit ip 192.168.0.10/32 192.168.20.2/32 [match=5] >>> 5 icmp replies comming from Lo0 N9K-1
30 permit ip any any [match=0]
```

Related configuration:

```
interface Ethernet1/4
description ***Link-to-N9K-2***
ip access-group TAC-IN in >>>
ip access-group TAC-OUT out
ip address 192.168.20.2/30
no shutdown
```

- Using troubleshooting steps previously outlined, the incoming and outgoing path of the packet was validated on a hop by hop basis between source and destination.

For this example, it was confirmed that there is no packet loss since all 5 ICMP packets were received and forwarded properly on each device.