



Configuring QoS TCAM Carving

- [About QoS TCAM Carving, on page 1](#)
- [Guidelines and Limitations for QoS TCAM Carving, on page 5](#)
- [Configuring QoS TCAM Carving, on page 7](#)

About QoS TCAM Carving

You can change the size of the access control list (ACL) ternary content addressable memory (TCAM) regions in the hardware.

On Cisco Nexus 9300 and 9500 platform switches and Cisco Nexus 3164Q, 31128PQ, 3232C, and 3264Q switches, the egress TCAM size is 1K, divided into four 256 entries. On Cisco Nexus NFE2-enabled devices (such as the Cisco Nexus 3232C and 3264Q switches), the ingress TCAM size is 6K, divided into twelve 512 slices. Three slices are in one group. On other Cisco Nexus 9300 and 9500 platform switches and Cisco Nexus 3164Q and 31128PQ switches, the ingress TCAM size is 4K, divided into eight 256 slices and four 512 slices. A slice is the unit of allocation. A slice can be allocated to one region only. For example, a 512-size slice cannot be used to configure two features of size 256 each. Similarly, a 256-size slice cannot be used to configure two features of size 128 each. The IPv4 TCAM regions are single wide. The IPv6, QoS, MAC, CoPP, and system TCAM regions are double wide and consume double the physical TCAM entries. For example, a logical region size of 256 entries actually consumes 512 physical TCAM entries.

On Cisco Nexus 9200 Series switches, the egress TCAM size is 2K, and the ingress TCAM size is 4K. The concepts of TCAM slices and single- and double-wide regions do not apply to these switches. For example, the ing-ifacl region can host IPv4, IPv6, or MAC type entries. IPv4 and MAC types occupy one TCAM entry whereas IPv6 types occupy two TCAM entries.

The number of default entries for QoS TCAM carving are:

- The default QoS TCAM carving for the Cisco Nexus 9504, Cisco Nexus 9508, and Cisco Nexus 9516 is for Layer 3 QoS (IPV4) with 256 entries. For these switches, all of the QoS TCAM entries are double wide.
- The default QoS TCAM carving for ALE (Application Leaf Engine) enabled devices is for Layer 2 port QoS (IPV4) with 256 entries. For these switches, all of the QoS TCAM entries are double wide.



Note In addition to the above TCAM, for ALE enabled devices, a separate TCAM in the Cisco Nexus C9396PX (uplink ports) and Cisco Nexus C93128TX (uplink ports) ASIC is used for the QoS classification policies applied on 40G uplink ports. By default, this separate TCAM is carved for Layer 3 QoS (IPV4), Layer 2 Port QoS (IPV4), and VLAN QoS (IPV4) with 256 entries each.

Table 1: QoS TCAM Regions (Cisco NX-OS Release 7.1(3)I6(1))

| Feature | Purpose | Region Name |
|------------|---|---|
| Egress QoS | QoS policy applied on interfaces in output direction. | IPV4: e-qos Cisco Nexus 922 series switch: egr-l2-qos, egr-l3-vlan-qos IPV6: e-ipv6-qos MAC: e-mac-qos See notes following table. |

Table 2: QoS TCAM Regions (Cisco NX-OS Release 6.1(2)I3(4) and earlier)

| Feature | Purpose | Region Name |
|-------------|---|--|
| Layer 3 QoS | QoS policy applied on Layer 3 interfaces. | IPV4: l3qos*, ns-l3qos* IPV6: ipv6-l3qos*, ns-ipv6-l3qos* See notes following table. |
| Port QoS | QoS policy applied on Layer 2 interfaces. | IPV4: qos*, ns-qos* IPV6: ipv6-qos*, ns-ipv6-qos* MAC: mac-qos*, ns-mac-qos* See notes following table. |
| VLAN QoS | QoS policy applied on VLAN. | IPV4: vqos, ns-vqos IPV6: ipv6-vqos*, ns-ipv6-vqos* MAC: mac-vqos*, ns-mac-vqos* See notes following table. |
| FEX QoS | QoS policy applied on FEX interfaces. | IPV4: fex-qos* IPV6: fex-ipv6-qos* MAC: fex-mac-qos* See notes following table. |

Table 3: QoS TCAM Regions (Cisco NX-OS Release 7.0(3)I1(1))

| Feature | Purpose | Region Name |
|-------------|---|--|
| Layer 3 QoS | QoS policy applied on Layer 3 interfaces. | IPV4: l3qos*, ns-l3qos* IPV6: ipv6-l3qos*, ns-ipv6-l3qos* See notes following table. |
| Port QoS | QoS policy applied on Layer 2 interfaces. | IPV4: qos*, ns-qos* IPV6: ipv6-qos*, ns-ipv6-qos* MAC: mac-qos*, ns-mac-qos* See notes following table. |
| VLAN QoS | QoS policy applied on VLAN. | IPV4: vqos, ns-vqos IPV6: ipv6-vqos*, ns-ipv6-vqos* MAC: mac-vqos*, ns-mac-vqos* See notes following table. |
| FEX QoS | QoS policy applied on FEX interfaces. | IPV4: fex-qos* IPV6: fex-ipv6-qos* MAC: fex-mac-qos* See notes following table. |

Table 4: QoS TCAM Regions (Cisco NX-OS Release 7.0(3)I1(2) and later)

| Feature | Purpose | Region Name |
|-------------|---|--|
| Layer 3 QoS | QoS policy applied on Layer 3 interfaces. | IPV4: l3qos*, ns-l3qos*, rp-qos** Cisco Nexus 9200 Series switch: ing-l3-vlan-qos IPV6: ipv6-l3qos*, ns-ipv6-l3qos*, rp-ipv6-qos** See notes following table. |
| Port QoS | QoS policy applied on Layer 2 interfaces. | IPV4: qos*, ns-qos*, rp-qos** Cisco Nexus 9200 Series switch: ing-l2-qos IPV6: ipv6-qos*, ns-ipv6-qos*, rp-ipv6-qos** MAC: mac-qos*, ns-mac-qos*, rp-mac-qos** See notes following table. |

| Feature | Purpose | Region Name |
|----------|---------------------------------------|---|
| VLAN QoS | QoS policy applied on VLAN. | IPV4: vqos, ns-vqos, rp-qos** Cisco Nexus 9200 Series switch: ing-l3-vlan-qos IPV6: ipv6-vqos*, ns-ipv6-vqos*, rp-ipv6-qos** MAC: mac-vqos*, ns-mac-vqos*, rp-mac-qos** See notes following table. |
| FEX QoS | QoS policy applied on FEX interfaces. | IPV4: fex-qos* IPv6: fex-ipv6-qos* MAC: fex-mac-qos* See notes following table. |



Note * The region is applicable only for ALE enabled devices and are required for classification policies applied on 40G uplink ports.



Note ** The region is applicable only for 100G enabled devices (such as the Cisco Nexus 9300 platform switch with the N9K-M4PC-CFP2 GEM or the Cisco Nexus 9500 platform switch with the Cisco Nexus 9408PC-CFP2 line card) and are required for classification policies and QoS scheduling applied on 100G uplink ports.

You need to save the configuration and reload the system for the region configuration to become effective.

About QoS TCAM Lite Regions

IPV4 requires QoS TCAM regions to be double wide TCAMs to support conform/violate policer statistics. If conform/violate statistics are not required, the size of the QoS TCAM entries can be reduced to single wide TCAMs by using QoS TCAM lite regions. Policing is supported by these regions, however only violate packets/bytes statistics are supported.

Table 5: QoS TCAM Regions (Release 7.1(3)I6(1))

| Feature | Purpose | Region Name |
|------------|---|--|
| Egress QoS | QoS policy applied on interfaces in output direction. | IPV4: e-qos-lite See notes following table. |

Table 6: QoS TCAM Lite Regions

| Feature | Purpose | Region Name |
|-------------|---|--------------------|
| Layer 3 QoS | QoS policy applied on Layer 3 interfaces. | IPV4: l3qos-lite |
| Port QoS | QoS policy applied on Layer 2 interfaces. | IPV4: qos-lite |
| VLAN QoS | QoS policy applied on VLAN. | IPV4: vqos-lite |
| FEX QoS | QoS policy applied on FEX interfaces. | IPV4: fex-qos-lite |



Note Cisco Nexus 9200 Series switches do not support QoS TCAM lite regions.



Note The region is applicable only for ALE enabled devices and are required for classification policies applied on 40G uplink ports.

You need to save the configuration and reload the system for the region configuration to become effective.



Note Either the regular version or the lite version of the QoS TCAM can be enabled. Both cannot be enabled at the same time. For example, either the IPv4 Port QoS or the IPv4 Port QoS lite version can be enabled at any one time.

Guidelines and Limitations for QoS TCAM Carving

TCAM region sizes have the following configuration guidelines and limitations:

- Configure egress QoS TCAM before upgrading from Release 7.0(3)I7(5) to Release 9.3(x) or Release 10.1(x) if a service-policy is attached. For more information on enabling egress QoS, see the [Enabling Egress QoS \(IPv4\), on page 13](#) section.
- **hardware access-list tcam label ing-qos optimize** is used to give separate label space for **ing-ifacl** ACLs and **ing-qos** service policies. Three labels are available for QoS policies. Some features, like VxLAN, add a QoS policy to an NVE interface by default and this reduces the available labels. **ing-ifacl-ipv4/ipv6-lite** commands move IPv4/IPv6 ACEs respectively to PT TCAM and are supported only on the following switches:
 - Cisco Nexus 9336C-FX2
 - Cisco Nexus 93240YC-FX2
 - Cisco Nexus 93240YC-FX2Z

- TCAM must be carved for the vQoS region if the QoS policy is configured within a VLAN. This will avoid traffic failure as shown in the syslog message in this example:

```
switch(config-vlan-config)# vlan configuration 3
switch(config-vlan-config)# service-policy type qos input INPUT_PREC
switch(config-vlan-config)# 2019 Jan 2 17:56:49 switch %$ VDC-1 %$
%ACLQOS-SLOT2-2-ACLQOS_FAILED: ACLQOS failure: VLAN QoS policy not
supported without TCAM carving for vQoS, traffic will fail please carve
TCAM for vQoS and IPV6-vQoS reload the module configure vlan qos policy
after module is up
```

- **show** commands with the **internal** keyword are not supported.
- After TCAM carving, you must save the configuration and reload the switch.
- Cisco Nexus 9200 platform switches and Cisco Nexus 9300-EX platform switches are of the same type and therefore, they have the same TCAM regions.
- By default, all IPv6 TCAMs are disabled (the TCAM size is set to 0).
- Use the **show hardware access-list tcam region** command to view the configured TCAM region size.
- The global CLI **hardware qos classify ns-only** command is introduced to enable configuration of the QoS policy on the NS ports without carving the T2 QoS region, for example, qos and l3-qos regions. This command removes the TCAM restrictions that are associated with the QoS classifications on the Application Leaf Engine (ALE) ports and it is only supported on Cisco Nexus 9000 series switches with ALE.

For example, for Layer 2 ALE port with IPv4 traffic, qos, and ns-qos TCAM carving is required for the QoS classification to work. With the **hardware qos classify ns-only** CLI command, ns-QoS TCAM alone is sufficient.

See the following example for applying the CLI **hardware qos classify ns-only** command:

```
switch(config)# hardware qos classify ns-only
Warning: This knob removes the restriction of carving qos as well as ns-qos TCAM region
for NS port QoS classification policies.
Warning: Only NS TCAM will be used, as a result policy-map statistics, marking and
policing is not supported on NS ports
```

See the following example for removing the CLI **hardware qos classify ns-only** command:

```
switch(config)# no hardware qos classify ns-only
Warning: Special knob removed. Please remove and apply QoS policies on NS ports to get
default behavior
```



Note Policing, policy-map statistics, and marking are not supported on the NS ports if the **hardware qos classify ns-only** CLI command is used. The **show policy-map interface ethernet x/y** does not return QoS statistics. The NS TCAM does not have some of the Network Forwarding Engine (NFE) TCAM resources, for example, range and so on. Therefore, the policies may need more TCAM entries.

- By default, the TCAM region for CoPP is 95% utilized on the Nexus 9300/Nexus 9500 platform switch. If you modify the CoPP policy, it is likely that you will need to modify other TCAM region sizes to allow for more space to be applied to the CoPP TCAM region.
- When any of the following classification criteria are used for IPv4 and IPv6, you must carve the IPv4 based QoS TCAM region. It is not necessary to carve an IPv6 based QoS TCAM region.

- Differentiated Services Code Point (DSCP) based classification
 - Class of service (CoS) based classification
 - IP precedence-based classification
- When a QoS policy is applied on multiple interfaces or multiple VLANs, the label is not shared since the statistics option is enabled.
To share the label for the same QoS policy that is applied on multiple interfaces or multiple VLANs, you must configure the QoS policy with no-stats option using the **service-policy type qos input my-policy no-stats** command.
 - On Cisco Nexus 9300 platform switches, the Cisco Nexus 9536PQ, 9564PX, and 9564TX line cards are used to enforce the QoS classification policies that are applied on 40G ports. It has 768 TCAM entries available for carving in 256-entry granularity. These region names are prefixed with "ns-".
 - For the Cisco Nexus 9536PQ, 9564PX, and 9564TX line cards, only the IPv6 TCAM regions consume double-wide entries. The rest of the TCAM regions consume single-wide entries.
 - When a VACL region is configured, it is configured with the same size in both the ingress and egress directions. If the region size cannot fit in either direction, the configuration is rejected.
 - On Cisco Nexus 9200 platform switches, the ing-sup region occupies a minimum size of 512 entries, and the egr-sup region occupies a minimum size of 256 entries. These regions cannot be configured to lesser values. Any region size can be carved with a value only in multiples of 256 entries (except for the span region, which can be carved only in multiples of 512 entries).
 - VLAN QoS is only supported on the Cisco Nexus 9508 switch with the -R series line card.
 - QoS has default TCAM sizes and these TCAM sizes must be nonzero on specific line cards to avoid failure of the line card during a reload.

Cisco Nexus 9504 and Cisco Nexus 9508 switches with the following line cards are affected:

- Cisco Nexus 96136YC-R
- Cisco Nexus 9636C-RX
- Cisco Nexus 9636Q-R
- Cisco Nexus 9636C-R

Configuring QoS TCAM Carving

You can change the default QoS TCAM carving to accommodate your network requirements. The following sections contain examples of how to change the default QoS TCAM carving.



Note You can use this procedure for all Cisco Nexus 9200, 9300, and 9500 Series switches and the Cisco Nexus 9500 Series switches. The examples do not apply to NFE2-enabled devices (such as the X9432C-S 100G line card and the C9508-FM-S fabric module), which must use TCAM templates to configure TCAM region sizes. For more information on using TCAM templates, see "Using Templates to Configure TCAM Region Sizes."

Once you apply a TCAM template, the **hardware access-list tcam region** command will not work. You must uncommit the template to use the command.

Enabling Layer 3 QoS (IPv6)

The default TCAM region configuration does not accommodate Layer 3 QoS (IPv6). To enable Layer 3 QoS (IPv6), you must decrease the TCAM size of another region and then increase the TCAM size to enable the new Layer 3 QoS (IPv6) region.

Table 7: Default TCAM Region Configuration (Ingress) for the Cisco Nexus 9504, Cisco Nexus 9508, and Cisco Nexus 9516 devices

| Region Name | Size | Width | Total Size |
|-----------------|------|-------|------------|
| IPV4 RACL | 1536 | 1 | 1536 |
| L3 QoS(IPV4) | 256 | 2 | 512 |
| COPP | 256 | 2 | 512 |
| System | 256 | 2 | 512 |
| Redirect | 256 | 1 | 256 |
| SPAN | 256 | 1 | 256 |
| VPC Convergence | 512 | 1 | 512 |
| | | | 4K |

Table 8: Default TCAM Region Configuration (Ingress) - For Layer 2-to-Layer 3 Configurations on Cisco Nexus 9200 Series Switches

| Region Name | Size | Width | Total Size |
|--------------------------|------|-------|------------|
| Ingress NAT | 0 | 1 | 0 |
| Ingress port ACL | 256 | 1 | 256 |
| Ingress VACL | 256 | 1 | 256 |
| Ingress RACL | 1536 | 1 | 1536 |
| Ingress Layer 2 QoS | 256 | 1 | 256 |
| Ingress Layer 3 VLAN QoS | 256 | 1 | 256 |
| Ingress supervisor | 512 | 1 | 512 |

| Region Name | Size | Width | Total Size |
|--------------------------|------|-------|------------|
| Ingress Layer 2 ACL SPAN | 256 | 1 | 256 |
| Ingress Layer 3 ACL SPAN | 256 | 1 | 256 |
| Port-based SPAN | 512 | 1 | 512 |
| | | | 4096 |

Table 9: Default TCAM Region Configuration (Ingress) - For Layer 3 Configurations on Cisco Nexus 9200 Series Switches

| Region Name | Size | Width | Total Size |
|--------------------------|------|-------|------------|
| Ingress NAT | 0 | 1 | 0 |
| Ingress port ACL | 0 | 1 | 0 |
| Ingress VACL | 0 | 1 | 0 |
| Ingress RACL | 1792 | 1 | 1792 |
| Ingress Layer 2 QoS | 256 | 1 | 256 |
| Ingress Layer 3 VLAN QoS | 512 | 1 | 512 |
| Ingress supervisor | 512 | 1 | 512 |
| Ingress Layer 2 ACL SPAN | 256 | 1 | 256 |
| Ingress Layer 3 ACL SPAN | 256 | 1 | 256 |
| Port-based SPAN | 512 | 1 | 512 |
| | | | 4096 |

Procedure

| | Command or Action | Purpose |
|--------|---|--|
| Step 1 | hardware access-list tcam region <i>region tcam-size</i> | To enable carving your Layer 3 QoS (IPv6) TCAM region, specify another region to free up resources. Also specify the reduced TCAM size for the region. Note Repeat this step for as many regions as necessary to free up sufficient resources to carve the new Layer 3 QoS (IPv6) TCAM region. |
| Step 2 | hardware access-list tcam region <i>region tcam-size</i> | Carve the new Layer 3 QoS (IPv6) TCAM region including the TCAM size (number of double wide entries). |

Example

This example sets the ingress Layer 3 QoS (IPv6) TCAM region size to 256. A Layer 3 QoS (IPv6) of size 256 takes 512 entries because IPv6 is double wide.

- Reduce the span and redirect regions to 0. This creates 512 entry spaces that are used to carve Layer 3 QoS (IPv6) with 256 entries (double wide).

```
switch(config)# hardware access-list tcam region redirect 0
Warning: Please reload the linecard for the configuration to take effect
Warning: BFD, DHCPv4 and DHCPv6 features will NOT be supported after this configuration
change.
switch(config)# hardware access-list tcam region span 0
Warning: Please reload the linecard for the configuration to take effect
switch(config)# hardware access-list tcam region ipv6-l3qos 256
Warning: Please reload the linecard for the configuration to take effect
```

Table 10: Updated TCAM Region Configuration After Reducing the IPv4 RACL (Ingress)

| Region Name | Size | Width | Total Size |
|--------------------|------|-------|------------|
| IPv4 RACL | 1536 | 1 | 1536 |
| Layer 3 QoS (IPv6) | 256 | 2 | 512 |
| Layer 3 QoS (IPv4) | 256 | 2 | 512 |
| CoPP | 256 | 2 | 512 |
| System | 256 | 2 | 512 |
| Redirect | 0 | 1 | 0 |
| SPAN | 0 | 1 | 0 |
| VPC Convergence | 512 | 1 | 512 |
| | | | 4K |

Enabling VLAN QoS (IPv4)

To enable VLAN QoS (IPv4), you must decrease the TCAM size of another region and then increase the TCAM size to enable the new VLAN QoS (IPv4) region.

The following table list the default sizes for the ingress TCAM regions for ALE enabled devices.

Table 11: Default TCAM Region Configuration (Ingress)

| Region Name | Size | Width | Total Size |
|-----------------|------|-------|------------|
| PACL (IPv4) | 512 | 1 | 512 |
| Port QoS (IPv4) | 256 | 2 | 512 |

| Region Name | Size | Width | Total Size |
|--------------|------|-------|------------|
| VACL (IPV4) | 512 | 1 | 512 |
| RACL(IPV4) | 512 | 1 | 512 |
| System | 256 | 2 | 512 |
| COPP | 256 | 2 | 512 |
| Redirect | 512 | 1 | 512 |
| SPAN | 256 | 1 | 256 |
| VPC Converge | 256 | 1 | 256 |
| | | | 4K |

Procedure

| | Command or Action | Purpose |
|---------------|--|--|
| Step 1 | <code>hardware access-list tcam region region tcam-size</code> | To enable carving for your VLAN QoS (IPv4) TCAM region, specify another region to free up resources. Also specify the reduced TCAM size for the region. Note Repeat this step for as many regions as necessary to free up sufficient resources to carve the new VLAN QoS (IPv4) TCAM region. |
| Step 2 | <code>hardware access-list tcam region region tcam-size</code> | Carve the new VLAN QoS (IPv4) TCAM region including the TCAM size (number of double wide entries). |

Example

This example sets the VLAN QoS (IPv4) TCAM size to 256. A VLAN QoS (IPv4) of size 256 takes 512 entries because QoS TCAM is double wide.

- Reduce the ingress Port QoS (IPv4) by 256 bytes (QoS features are double wide, 2 x 256 = 512) and add an ingress VLAN QoS (IPv4) with 256 (2 x 256).

```
switch(config)# hardware access-list tcam region qos 0
Warning: Please reload the linecard for the configuration to take effect
switch(config)# hardware access-list tcam region vqos 256
Warning: Please reload the linecard for the configuration to take effect
```

Table 12: Updated TCAM Region Configuration After Reducing the IPv4 Port QoS Ingress

| Region Name | Size | Width | Total Size |
|-------------|------|-------|------------|
| PACL (IPV4) | 512 | 1 | 512 |

| Region Name | Size | Width | Total Size |
|-----------------|------|-------|------------|
| Port QoS (IPV4) | 0 | 2 | 0 |
| VLAN QoS(IPV4) | 256 | 2 | 512 |
| VACL (IPV4) | 512 | 1 | 512 |
| RACL(IPV4) | 512 | 1 | 512 |
| System | 256 | 2 | 512 |
| COPP | 256 | 2 | 512 |
| Redirect | 512 | 1 | 512 |
| SPAN | 256 | 1 | 256 |
| VPC Converg | 256 | 1 | 256 |
| | | | 4K |

Notes for Enabling VLAN QoS

The VLAN QoS feature enables Layer 2 bridged database lookup for QoS with VLAN as the key instead of the port.

To enable VLAN QoS, you must decrease the TCAM size of another region and increase the TCAM size for the VLAN QoS region.

To configure the size of the VLAN QoS TCAM region:

- Configure the IPv4 vqos to 640 entries.
- Configure the IPv6 ipv6-vqos to 256 entries.
- Decrease the IPv4 qos to 0 entries.
- Decrease the IPv6 ipv6-qos to 0 entries.

```
switch(config)# hardware access-list tcam region vqos 640
switch(config)# hardware access-list tcam region ipv6-vqos 256
switch(config)# hardware access-list tcam region qos 0
switch(config)# hardware access-list tcam region ipv6-qos 0
```



Note After configuring the TCAM size for VLAN QoS, it is necessary to reload the line card.

Enabling FEX QoS (IPv4)



Note The FEX QoS feature is not supported on the Cisco Nexus 9508 switch (NX-OS 7.0(3)F3(3)).

To enable FEX QoS (IPv4), you must decrease the TCAM size of another region and then increase the TCAM size to enable the new FEX QoS (IPv4) region.

Procedure

| | Command or Action | Purpose |
|--------|--|--|
| Step 1 | <code>hardware access-list tcam region region tcam-size</code> | To enable carving your FEX QoS (IPv4) TCAM region, specify another region to free up resources. Also specify the reduced TCAM size for the region. Note Repeat this step for as many regions as necessary to free up sufficient resources to carve the new FEX QoS (IPv4) TCAM region. |
| Step 2 | <code>hardware access-list tcam region region tcam-size</code> | Carve the new FEX QoS (IPv4) TCAM region including the TCAM size (number of double wide entries). |

Example

This example sets the FEX QoS (IPv4) TCAM size to 256. A FEX QoS (IPv4) of size 256 takes 512 entries because QoS TCAM is double wide.

- Reduce the IPv4 FEX IFACL region by 512 entries and add a FEX QoS (IPv4) region with 512 entries.

```
switch(config)# hardware access-list tcam region fex-ifacl 0
Warning: Please reload the linecard for the configuration to take effect
switch(config)# hardware access-list tcam region fex-qos 256
Warning: Please reload the linecard for the configuration to take effect
```

Enabling Egress QoS (IPv4)

To enable QoS (IPv4) TCAM, you must decrease the TCAM size of another region and then increase the TCAM size to enable the new QoS (IPv4) TCAM region.



Note The egress QoS feature is not supported on the Cisco Nexus 9508 switch (Cisco NX-OS 7.0(3)F3(3)).



Note Egress marking and policing is supported on all Network Forwarding Engine (NFE) platforms. Egress classification for egress packet scheduling is supported only on 100G platforms.

Beginning with Cisco NX-OS Release 7.0(3)I6(1), the Cisco Nexus 93108TC-EX, 93180LC-EX, and 93180YC-EX switches, and 97160YC-EX, 9732C-EX, 9736C-EX line cards support the Layer 2 and Layer 3 egress policer.

Beginning with Cisco NX-OS Release 7.0(3)I1(2), to enable egress QoS (IPv4), you must decrease the TCAM size of the **e-racl** region and then increase the TCAM size for the egress QoS (IPv4) region.

The following are considerations for egress QoS (IPv4) and TCAM regions:

- Egress QoS TCAM is based on packet type, such as **e-qos**. TCAM carving is needed to match IPv4 packets on VLAN, layer 2, and layer 3 port types.
- All egress QoS (IPv4, IPv6, and MAC) TCAM regions are double-wide, except for the **e-qos-lite** region which is single-wide.
- Violated and non-violated statistics are supported for policing action when a double-wide TCAM is configured.
- When a single-wide TCAM (**e-qos-lite**) is configured, only non-violated statistics are reported in the presence of a policing action. The violated statistics are always reported as zero instead of NA for the **qos-lite** region. The policing action (1R2C or 2R3C) is still properly enforced. Only statistics reporting is limited to non-violated statistics. If you want to view violated statistics, regular QoS TCAM should be used instead.
- Statistics are disabled when the optional **no-stats** keyword is used and policies are shared (where applicable).
- Egress QoS policies on ALE uplink ports on top-of-rack (TOR) platforms are not supported.
- The egress QoS policy supports marking, policing, and classification.



Note Egress classification for egress packet scheduling is supported only on 100G platforms.

- Egress qos policies do not support packet-length based matching.
- The **set qos-group** command is not supported for egress QoS policies.

However, the **set qos-group** command is supported for egress QoS policies when applied on a 100G interface.

- Depending on the policy-map match criteria, the relevant egress QoS TCAM regions, such as **e-qos**, **e-mac-qos**, **e-ipv6-qos**, **egr-l2-qos**, and **egr-l3-vlan-qos**, must be carved for end-to-end QoS within the device.
- Set the egress QoS TCAM region size to 0 before downgrading to earlier images. Remove all egress QoS policies before downgrading to earlier images.

Procedure

| | Command or Action | Purpose |
|--------|---|--|
| Step 1 | hardware access-list tcam region e-racl <i>tcam-size</i> | To enable carving your QoS (IPv4) TCAM region, specify the e-racl region to free up resources. Also specify the reduced TCAM size for the e-racl region. |

| | Command or Action | Purpose |
|--------|---|--|
| Step 2 | <p>hardware access-list tcam region [e-qos e-qos-lite e-ip6-qos e-mac-qos egr-l2-qos egr-l3-vlan-qos] <i>tcam-size</i></p> <p>Example:</p> <pre>switch(config)# hardware access-list tcam region egr-l2-vlan-qos 256</pre> <p>Warning: Please reload all linecards for the configuration to take effect</p> <pre>switch(config)#</pre> <p>Example:</p> <pre>switch(config)# hardware access-list tcam region egr-l3-vlan-qos 256</pre> <p>Warning: Please reload all linecards for the configuration to take effect</p> <pre>switch(config)#</pre> | <p>The hardware access-list tcam region [e-qos e-qos-lite e-ip6-qos e-mac-qos egr-l2-qos egr-l3-vlan-qos] <i>tcam-size</i> command specifies the egress QoS (IPv4) TCAM region and the TCAM size. The egr-l2-qos egr-l3-vlan-qos options specify the egress QoS TCAM regions and TCAM size. An egress QoS TCAM of 256 size, takes 512 entries because QoS TCAM is double-wide.</p> <p>Note All egress QoS (IPv4) TCAM regions are double wide, except for the e-qos-lite region which is single wide.</p> |
| Step 3 | <p>Optional: [no]hardware access-list tcam label egr-l2-qos 6</p> <p>Example:</p> <pre>switch(config)# hardware access-list tcam label egr-l2-qos 6</pre> <p>Warning: Please reload all linecards for the configuration to take effect</p> <pre>switch(config)#</pre> | <p>Configures 64 unique egress QoS policies on Layer 2 physical interfaces.</p> <p>To disable the changes, use the no form of this command.</p> <p>Note This command is supported only on Cisco Nexus 9300-FX and FX2 platform switches.</p> <p>Note Beginning with Cisco NX-OS Release 10.3(3)F, this command is supported on Cisco Nexus 9300-FX3 platform switches.</p> |

Using Templates to Configure TCAM Region Sizes



Note Using templates to configure TCAM region sizes is not supported on the Cisco Nexus 9508 switch (NX-OS 7.0(3)F3(3)).

Beginning with Cisco NX-OS Release 7.0(3)I3(1), you can use create and apply custom templates to configure TCAM region sizes.



Note Once you apply a TCAM template, the **hardware access-list tcam region** command will not work. You must uncommit the template in order to use the command.

SUMMARY STEPS

1. **configure terminal**
2. **[no] hardware profile tcam resource template** *template-name* **ref-template** {**nfe** | **nfe2** | {**l2-l3** | **l3**}}
3. (Optional) *region tcam-size*

4. **exit**
5. **[no] hardware profile tcam resource service-template** *template-name*
6. (Optional) **show hardware access-list tcam template** {**all** | **nfe** | **nfe2** | **I2-I3** | **I3** | *template-name*}
7. (Optional) **copy running-config startup-config**
8. **reload**

DETAILED STEPS

| | Command or Action | Purpose |
|---------------|--|--|
| Step 1 | configure terminal Example: <pre>switch# configure terminal switch(config)#</pre> | Enters global configuration mode. |
| Step 2 | Required: [no] hardware profile tcam resource template <i>template-name</i> ref-template { nfe nfe2 { I2-I3 I3 }} Example: <pre>switch(config)# hardware profile tcam resource template SR_MPLS_CARVE ref-template nfe2 switch(config-tcam-temp)#</pre> | Creates a template for configuring ACL TCAM region sizes. nfe —The default TCAM template for Network Forwarding Engine (NFE)-enabled Cisco Nexus 9300 and 9500 Series, 3164Q, and 31128PQ devices. nfe2 —The default TCAM template for NFE2-enabled Cisco Nexus 9500 Series, 3232C, and 3264Q devices. I2-I3 —The default TCAM template for Layer 2 and Layer 3 security configurations on Cisco Nexus 9200 Series switches. I3 —The default TCAM template for Layer 3 configurations on Cisco Nexus 9200 Series switches. The Layer 3 TCAM template is the default template for the Cisco Nexus 9200 Series switches. |
| Step 3 | (Optional) <i>region tcam-size</i> Example: <pre>switch(config-tcam-temp)# mpls 256</pre> | Adds any desired TCAM regions and their sizes to the template. Enter this command for each region you want to add to the template. |
| Step 4 | Required: exit Example: <pre>switch(config-tcam-temp)# exit switch(config#)</pre> | Exits the TCAM template configuration mode. |
| Step 5 | Required: [no] hardware profile tcam resource service-template <i>template-name</i> Example: <pre>switch(config)# hardware profile tcam resource service-template SR_MPLS_CARVE</pre> | Applies the custom template to all line cards and fabric modules. |

| | Command or Action | Purpose |
|--------|--|---|
| Step 6 | (Optional) show hardware access-list tcam template { all nfe nfe2 I2-I3 I3 <i>template-name</i> } Example: <pre>switch(config)# show hardware access-list tcam template SR_MPLS_CARVE</pre> | Displays the configuration for all TCAM templates or for a specific template. |
| Step 7 | (Optional) copy running-config startup-config Example: <pre>switch(config)# copy running-config startup-config</pre> | Copies the running configuration to the startup configuration. |
| Step 8 | reload Example: <pre>switch(config)# reload</pre> | Reloads the device. Note The configuration is effective only after you enter copy running-config startup-config + reload . |

Verifying QoS TCAM Carving

After you adjust the TCAM region sizes, enter the **show hardware access-list tcam region** command to display the TCAM sizes that will be applicable on the next reload of the device.

To display the configuration of a TCAM template, use the **show hardware access-list tcam template** {**all** | **nfe** | **nfe2** | **I2-I3** | **I3** | *template-name*} command where:

- **all**—Displays configuration for all TCAM templates.
- **nfe**—The default TCAM template for Network Forwarding Engine (NFE)-enabled Cisco Nexus 9300 and 9500 Series, 3164Q, and 31128PQ devices.
- **nfe2**—The default TCAM template for NFE2-enabled Cisco Nexus 9500, 3232C, and 3264Q devices.
- **I2-I3**—The default TCAM template for Layer 2-to-Layer 3 configurations on Cisco Nexus 9200 Series switches.
- **I3**—The default TCAM template for Layer 3 configurations on Cisco Nexus 9200 Series switches.



Note To keep all modules synchronized, you must reload all line card modules or enter the **copy running-config startup-config** command and the **reload** command to reload the device. Multiple TCAM region configurations require only a single reload. You can wait until you complete all of your TCAM region configurations before you reload the device.

If you exceed the 4K ingress limit for all TCAM regions when you configure a TCAM region, the following message appears:

```
ERROR: Aggregate TCAM region configuration exceeded the available Ingress TCAM space.
```

Please re-configure.

If TCAM for a particular feature is not configured and you try to apply a feature that requires TCAM carving, the following message appears:

```
ERROR: Module x returned status: TCAM region is not configured. Please configure TCAM region and retry the command.
```