



Configuring QoS TCAM Carving

- [About QoS TCAM Carving, on page 1](#)
- [Guidelines and Limitations, on page 3](#)
- [Configuring QoS TCAM Carving, on page 4](#)

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 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.

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.



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.

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 3: 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 4: 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 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

TCAM region sizes have the following configuration guidelines and limitations:

- **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.
- 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 need to 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 need to 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 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 the Cisco Nexus 9508 switch with the -R series line card, VLAN QoS is only supported with Cisco NX-OS Release 7.0(3)F3(3) and later releases.

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.

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 5: 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

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 6: 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 7: Default TCAM Region Configuration (Ingress)

Region Name	Size	Width	Total Size
PAACL (IPV4)	512	1	512
Port 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

Procedure

	Command or Action	Purpose
Step 1	hardware access-list tcam region <i>region tcam-size</i>	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	hardware access-list tcam region <i>region tcam-size</i>	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 8: Updated TCAM Region Configuration After Reducing the IPv4 Port QoS Ingress

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

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
```

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.



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.
```

