Nokia Service Router Linux

Release 23.7

ACL and Policy-Based Routing Guide

3HE 19575 AAAA TQZZA
Edition: 01
August 2023

Nokia is committed to diversity and inclusion. We are continuously reviewing our customer documentation and consulting with standards bodies to ensure that terminology is inclusive and aligned with the industry. Our future customer documentation will be updated accordingly.

Table of contents

1 About this guide

This document describes ACLs and policy-based routing for the Nokia Service Router Linux (SR Linux). Examples of commonly used commands are provided.

This document is intended for network technicians, administrators, operators, service providers, and others who need to understand how the router is configured.

Note:
This manual covers the current release and may also contain some content that will be released in later maintenance loads. See the *SR Linux Release Notes* for information on features supported in each load.

Configuration and command outputs shown in this guide are examples only; actual displays may differ depending on supported functionality and user configuration.

1.1 Precautionary and information messages

The following are information symbols used in the documentation.

DANGER: Danger warns that the described activity or situation may result in serious personal injury or death. An electric shock hazard could exist. Before you begin work on this equipment, be aware of hazards involving electrical circuitry, be familiar with networking environments, and implement accident prevention procedures.

WARNING: Warning indicates that the described activity or situation may, or will, cause equipment damage, serious performance problems, or loss of data.

Caution: Caution indicates that the described activity or situation may reduce your component or system performance.

Note: Note provides additional operational information.

Tip: Tip provides suggestions for use or best practices.

1.2 Conventions

Nokia SR Linux documentation uses the following command conventions.

- **Bold** type indicates a command that the user must enter.
- Input and output examples are displayed in `Courier` text.
- An open right-angle bracket indicates a progression of menu choices or simple command sequence (often selected from a user interface). Example: **start** > **connect to**.
- A vertical bar (|) indicates a mutually exclusive argument.

- Square brackets ([ ]) indicate optional elements.
- Braces ({ }) indicate a required choice. When braces are contained within square brackets, they indicate a required choice within an optional element.
- *Italic* type indicates a variable.

Generic IP addresses are used in examples. Replace these with the appropriate IP addresses used in the system.

2 What's new

There have been no updates in this document since it was last released.

3 Access control lists

An Access Control List, or ACL, is an ordered set of rules that are evaluated on a packet-by-packet basis to determine whether access should be provided to a specific resource. ACLs can be used to drop unauthorized or suspicious packets from entering or leaving a routing device via specified interfaces.

An ACL is applied to a selected set of traffic contexts. A traffic context could be all the IPv4 or IPv6 packets arriving or leaving on a specific subinterface, or the out-of-band IP traffic arriving on a management interface, or all the in-band IPv4 or IPv6 packets that are locally terminating on the CPM of the router.

Each ACL rule, or entry, has a sequence ID. The ACL evaluates packets starting with the entry with the lowest sequence ID, progressing to the entry with the highest sequence ID. Evaluation stops at the first matching entry (that is, when the packet matches all of the conditions specified by the ACL entry).

IPv4/IPv6 ACLs

For IPv4 and IPv6 traffic, SR Linux supports the following types of ACLs:

- Interface filters

  An interface filter is an IPv4 or IPv6 ACL that is applied to a routed or bridged subinterface to restrict traffic allowed to enter or exit the subinterface. An interface ACL can be applied to input or output traffic for one or more subinterfaces.

  See Interface filters for more information.

- Packet capture filters

  A packet capture filter is an IPv4 or IPv6 ACL used to extract packets to the control plane for inspection by packet capture tools. When an ingress IP packet on any line card transits through the router, and it matches a rule in a packet capture filter policy, it is copied and extracted toward the CPM and delivered to a Linux virtual Ethernet interface so that it can be displayed by a packet capture utility, or encapsulated and forwarded to a remote monitoring system.

  See Packet capture filters for more information.

- CPM filters

  A CPM filter is an IPv4 or IPv6 ACL used for control plane protection. There is one CPM filter for IPv4 traffic and another CPM filter for IPv6 traffic. When an ingress IP packet is matched by a CPM filter rule, and it is a terminating packet (that is, it must be extracted to the CPM), then it is processed according to the matching CPM filter rule.

  See Control plane module (CPM) filters for more information.

- System filters (7220 IXR-D1, D2, and D3 systems only)

  A system filter ACL is an IPv4 or IPv6 ACL that is evaluated early in the ingress pipeline, at a stage before tunnel termination occurs and before interface filters are run. For VXLAN traffic, system filters can match and drop unauthorized VXLAN tunnel packets before they are decapsulated, based on information in the outer header.

  See System filters for more information.

MAC ACLs

For Layer 2 traffic, SR Linux supports the following types of MAC ACLs:

- Interface MAC filters

  An interface MAC filter is a Layer 2 ACL that is applied to a routed or bridged subinterface to restrict the Ethernet frames allowed to enter or exit from that subinterface. An interface MAC filter can match Ethernet frames carrying an IP or non-IP payload and can be applied to the input or output traffic of one or more subinterfaces.

  See Interface filters for more information.

- CPM filter MAC ACLs

  A CPM filter MAC ACL is used for control plane protection. When a CPM filter MAC ACL is created, its rules are automatically evaluated against all non-IP traffic that is extracted to the CPM as a result of a match with an extraction rule (for example, a match of a specific well-known MAC DA value). This is regardless of whether the traffic entered from of network instance, subinterface, TD3 pipeline, and so on.

  See Control plane module (CPM) filters for more information.

3.1 ACL actions

When a packet matches an ACL entry, an action specified by the ACL entry is applied to the packet. An ACL entry has a primary action and an optional secondary action. The secondary action extends the primary action with additional packet handling operations.

For traffic transiting through the router, ACL entries support the following primary actions:

- accept – Allow the packet through to the next processing function.
- accept and log – Allow the packet through to the next processing function and send information about the accepted packet to the log application.
- drop – Discard the packet without ICMP generation.
- drop and log – Discard the packet without ICMP generation and send information about the dropped packet to the log application.

For traffic transiting through the router, the following secondary action is supported:

- log – Send information about the packet to the log application.

For traffic terminating on the CPM of the router, the preceding primary and secondary actions are supported, as well as the following secondary actions for ACL entries where accept is the primary action:

- distributed-policer – If the packet is extracted to the CPM, feed the packet to a hardware-based policer, which determines if the packet should be queued by the line card toward the CPM or dropped because a bit-per-second rate is exceeded.
- system-cpu-policer – If the packet has been extracted to the CPM, feed the packet to a software-based policer, which determines if the packet should be delivered to the CPM application or dropped because a packet-per-second rate is exceeded.

If a packet matches an ACL entry, no further evaluation is done for the packet. If the packet does not match any ACL entry, the default action is accept. To drop traffic that does not match any ACL entry, you can optionally configure an entry with the highest sequence ID in the ACL to drop all traffic. This causes traffic that does not match any of the lower-sequence ACL entries to be dropped.

The supported actions for each type of ACL differ based on the hardware platform where the ACL is configured. The following tables indicate which actions are supported for each ACL filter type for each hardware platform.

3.1.1 Supported ACL actions for 7250 IXR systems

The following table lists the supported actions for each ACL filter type on 7250 IXR systems.

Table 1: Supported actions for each ACL filter type (7250 IXR)

ACL filter type	Action	Supported?
IPv4/IPv6 Interface filter (input)	accept	Yes
	accept and log	Yes
	drop	Yes
	drop and log	Yes
IPv4/IPv6 Interface filter (output)	accept	Yes
	accept and log	Yes
	drop	Yes
	drop and log	Yes
IPv4/IPv6 CPM filter	accept	Yes
	accept and log	Yes
	drop	Yes
	drop and log	Yes
	accept and distributed-policer	Yes
	accept and distributed-policer and log	No log generated
	accept and system-cpu-policer	Yes
	accept and system-cpu-policer and log	No log generated
Packet capture filter	accept	Yes
	copy	Yes
System filter	accept	No
	drop	No
	drop and log	No

ACL filter type	Action	Supported?
Interface MAC filter (input)	accept	No
	accept and log	No
	drop	No
	drop and log	No
Interface MAC filter (output)	accept	No
	accept and log	No
	drop	No
	drop and log	No
CPM filter MAC ACL	accept	No
	accept and log	No
	drop	No
	drop and log	No
	accept and distributed-policer	No
	accept and distributed-policer and log	No
	accept and system-cpu-policer	No
	accept and system-cpu-policer and log	No

3.1.2 Supported ACL actions for 7220 IXR-D1/D2/D3 systems

The following table lists the supported actions for each ACL filter type on 7220 IXR-D1, D2, and D3 systems.

Table 2: Supported actions for each ACL filter type (7220 IXR-D1, D2, and D3)

ACL filter type	Action	Supported?
IPv4/IPv6 Interface filter (input)	accept	Yes
	accept and log	No log generated
	drop	Yes
	drop and log	Yes, using separate CPU queue

ACL filter type	Action	Supported?
IPv4/IPv6 Interface filter (output)	accept	Yes
	accept and log	No log generated
	drop	Yes
	drop and log	No log generated
IPv4/IPv6 CPM filter	accept	Yes
	accept and log	No log generated
	drop	Yes
	drop and log	Yes, using shared CPU queue
	accept and distributed-policer	Yes
	accept and distributed-policer and log	No log generated[1]
	accept and system-cpu-policer	Yes
	accept and system-cpu-policer and log	No log generated
Packet capture filter	accept	Yes
	copy	Yes
System filter	accept	Yes
	drop	Yes
	drop and log	Yes
MAC Interface filter (input)	accept	Yes
	accept and log	No log generated
	drop	Yes
	drop and log	Yes, using separate CPU queue
MAC Interface filter (output)	accept	Yes
	accept and log	No log generated
	drop	Yes
	drop and log	No log generated

[1] Log not supported on 7220 IXR-D2 and D3, D4, and D5.

ACL filter type	Action	Supported?
MAC CPM filter	accept	Yes
	accept and log	No log generated
	drop	Yes
	drop and log	Yes, using shared CPU queue
	accept and distributed-policer	Yes
	accept and distributed-policer and log	No log generated[2]
	accept and system-cpu-policer	Yes
	accept and system-cpu-policer and log	No log generated

3.1.3 Supported ACL actions for 7220 IXR-D4/D5 systems

The following table lists the supported actions for each ACL filter type on 7220 IXR-D4 and D5 systems.

Table 3: Supported actions for each ACL filter type (7220 IXR-D4 and D5)

ACL filter type	Action	Supported?
IPv4/IPv6 Interface filter (input)	accept	Yes
	accept and log	No log generated
	drop	Yes
	drop and log	Yes
IPv4/IPv6 Interface filter (output)	accept	Yes
	accept and log	No log generated
	drop	Yes
	drop and log	No log generated
IPv4/IPv6 CPM filter	accept	Yes
	accept and log	No log generated
	drop	Yes
	drop and log	Yes, using shared CPU queue

[2] Log not supported on 7220 IXR-D2 and D3, D4, and D5.

ACL filter type	Action	Supported?
	accept and distributed-policer	No
	accept and distributed-policer and log	No
	accept and system-cpu-policer	Yes
	accept and system-cpu-policer and log	No log generated
Packet capture filter	accept	Yes
	copy	Yes
System filter	accept	Yes
	drop	Yes
	drop and log	Yes
MAC Interface filter (input)	accept	Yes
	accept and log	No log generated
	drop	Yes
	drop and log	Yes
MAC Interface filter (output)	accept	Yes
	accept and log	No log generated
	drop	Yes
	drop and log	No log generated
MAC CPM filter	accept	Yes
	accept and log	No log generated
	drop	Yes
	drop and log	Yes
	accept and distributed-policer	Yes
	accept and distributed-policer and log	No log generated
	accept and system-cpu-policer	Yes
	accept and system-cpu-policer and log	No log generated

3.1.4 Supported ACL actions for 7220 IXR-H systems

The following table lists the supported actions for each ACL filter type on 7220 IXR-H2, H3, and H4 systems.

Table 4: Supported actions for each ACL filter type (7220 IXR-H systems)

ACL filter type	Action	Supported?
IPv4/IPv6 Interface filter (input)	accept	Yes
	accept and log	No log generated
	drop	Yes
	drop and log	Yes, but logged packet is also processed by CPM filter
IPv4/IPv6 Interface filter (output)	accept	Yes
	accept and log	No log generated
	drop	Yes
	drop and log	No log generated
IPv4/IPv6 CPM filter	accept	Yes
	accept and log	No log generated
	drop	Yes
	drop and log	Yes
	accept and distributed-policer	No policing
	accept and distributed-policer and log	No policing and no log generated
	accept and system-cpu-policer	Yes
	accept and system-cpu-policer and log	No log generated
Packet capture filter	accept	Yes
	copy	Yes
System filter	accept	No
	drop	No
	drop and log	No

ACL filter type	Action	Supported?
Interface MAC filter (input)	accept	No
	accept and log	No
	drop	No
	drop and log	No
Interface MAC filter (output)	accept	No
	accept and log	No
	drop	No
	drop and log	No
CPM filter MAC ACL	accept	No
	accept and log	No
	drop	No
	drop and log	No
	accept and distributed-policer	No
	accept and distributed-policer and log	No
	accept and system-cpu-policer	No
	accept and system-cpu-policer and log	No

3.2 ACL match conditions

You can specify the following match conditions in IPv4, IPv6, and MAC ACLs.

Match conditions for IPv4 ACLs

IPv4 ACLs analyze IPv4 packets. The following match criteria are supported by IPv4 ACLs:

- IPv4 destination prefix and prefix-length
- IPv4 destination address and address-mask
- TCP/UDP destination port (range)
- ICMP type/code
- IP protocol number
- IPv4 source prefix and prefix-length
- IPv4 source address and address-mask
- TCP/UDP source port (range)

- TCP flags: RST, SYN, and ACK
- Packet fragmentation: whether the packet is a fragment
- Packet fragmentation: whether the packet is a first-fragment (`fragment-offset=0` and `more-fragments=1`)
- IP DSCP

Match conditions for IPv6 ACLs

IPv6 ACLs analyze IPv6 packets. The following match criteria are supported by IPv6 ACLs:

- IPv6 destination prefix and prefix-length
- IPv6 destination address and address-mask
- TCP/UDP destination port (range)
- ICMPv6 type/code
- IPv6 next-header value. This is the value in the very first next-header field, in the fixed header.
- IPv6 source prefix and prefix-length
- IPv6 source address and address-mask
- TCP/UDP source port (range)
- TCP flags: RST, SYN, and ACK
- IP DSCP

Match conditions for MAC ACLs

MAC ACLs analyze Ethernet frames. The following match criteria are supported by MAC ACLs:

- MAC destination address with configurable address mask
- MAC source address match with configurable address mask
- outermost VLAN ID (single VLAN ID value or one contiguous VLAN ID range)
- Ethertype number after the last 802.1Q VLAN tag (if any), specified as a number (0x0600-0xFFFF) or a well-known name

3.3 Interface filters

An interface filter is an IPv4/IPv6 or MAC ACL that restricts traffic allowed to enter or exit a subinterface.

IPv4/IPv6 interface filters

IPv4 and IPv6 ACLs can be applied to a subinterface to restrict IP traffic entering or exiting that subinterface, as follows:

- Input IPv4 ACLs – When an IPv4 filter is used as an input ACL on a subinterface that carries IPv4 traffic, the rules apply to native IPv4 packets (ethertype 0x0800) that enter the subinterface and would normally terminate locally (control/management plane packets) or transit through the router.

  The rules also apply to MPLS-encapsulated IPv4 packets (ethertype 0x8847) that are terminating or transit.

- Input IPv6 ACLs – When an IPv6 filter is used as an input ACL on a subinterface that carries IPv6 traffic, the rules apply to native IPv6 packets (ethertype 0x86DD) that enter the subinterface and would normally terminate locally (control/management plane packets) or transit through the router.

  The rules also apply to MPLS-encapsulated IPv6 packets (ethertype 0x8847) that are terminating or transit.
- Output IPv4 ACLs – When an IPv4 filter is used as an output ACL on a subinterface that carries IPv4 traffic, the rules apply to native IPv4 packets (ethertype 0x0800) that exit the subinterface, including packets that originated locally (control/management plane packets) and packets that transited through the router.

  The rules do not apply to MPLS-encapsulated IPv4 packets (ethertype 0x8847) exiting the subinterface.
- Output IPv6 ACLs – When an IPv6 filter is used as an output ACL on a subinterface that carries IPv6 traffic, the rules apply to native IPv6 packets (ethertype 0x86DD) that exit the subinterface, including packets that originated locally (control/management plane packets) and packets that transited through the router.

  The rules do not apply to MPLS-encapsulated IPv6 packets (ethertype 0x8847) exiting the subinterface.

MAC interface filters

A MAC interface filter is a Layer 2 ACL that can be applied to a routed or bridged subinterface to restrict the Ethernet frames allowed to enter or exit that subinterface. An interface MAC ACL can match Ethernet frames carrying an IP or non-IP payload.

For a specific direction of traffic (input or output), a routed or bridged subinterface of an Ethernet port or LAG can have either a MAC interface ACL or an IPv4/IPv6 interface ACL applied, but not both at the same time. This restriction also applies to subinterfaces of the `mgmt0` management interface.

Table 5: MAC ACL filter rules

Subinterface type	MAC ACL traffic type	Rules
Routed	Input	Rules apply to all Ethernet frames matched by the subinterface encapsulation, even if they contain an IP or MPLS-IP payload.
	Output	Rules apply to all Ethernet frames egressing the subinterface except VXLAN-encapsulated packets and MPLS-encapsulated packets.
Bridged	Input	Rules apply to all Ethernet frames matched by the subinterface encapsulation, even if they contain an IP or MPLS-IP payload, and regardless of whether they are switched or forwarded to the IRB. However, if the IRB also has an input IPv4/IPv6 ACL applied to it, the IRB action takes priority over the

Subinterface type	MAC ACL traffic type	Rules
		bridged subinterface MAC ACL action.
	Output	Rules apply to all Ethernet frames egressing the subinterface except frames routed from the IRB subinterface.

MAC ACLs are not supported on IRB subinterfaces or loopback or `system0` subinterfaces.

3.3.1 Creating an IPv4 ACL

Procedure

To configure an IPv4 ACL filter, you specify one or more entries consisting of match conditions and the action to take for IPv4 traffic that matches the conditions.

Example: Accept IPv4 traffic matching specified IP address and source/destination port

The following is an example IPv4 ACL that has one entry. This example creates an IPv4 ACL named `ip_tcp`. Within the `ip_tcp` ACL, an entry with sequence ID 1000 is configured. The action is specified as `accept`, with logging set to `true`.

The filter matches packets with IP destination address 10.1.3.1/32; for TCP traffic, if the source address 10.1.5.1/32, destination port 6789, and source port 6722 matches the filter, the traffic stream is accepted.

```
--{ * candidate shared default }--[  ]--
# info acl
    ipv4-filter ip_tcp {
        entry 1000 {
            description Match_IP_Address_TCP_Protocol_Ports
            action {
                accept {
                    log true
                }
            }
            match {
                destination-address 10.1.3.1/32
                protocol tcp
                source-address 10.1.5.1/32
                destination-port {
                    value 6789
                }
                source-port {
                    value 6722
                }
            }
        }
    }
```

Example: Drop all matching traffic

The following is an example of an entry added to the `ip_tcp` ACL that causes all traffic to be dropped. Because it has the highest sequence ID, traffic that matches any of the lower-sequenced

ACL entries would have been accepted before being evaluated by this entry. Only traffic that did not match any of the other ACL entries would be dropped by this entry.

```
--{ * candidate shared default }--[  ]--
# info acl
    ipv4-filter ip_tcp {
        entry 65535 {
            action {
                drop {
                    log true
                }
            }
        }
    }
```

Note that the `drop` action with logging set to `true` is not supported on 7220 IXR-D1, D2, D3, and D5 systems when it is attached as an egress filter.

3.3.2 Creating an IPv6 ACL

Procedure

To configure an IPv6 ACL filter, you specify one or more entries consisting of match conditions and the action to take for IPv6 traffic that matches the conditions.

Example

The following is an example IPv6 ACL that has one entry. This example creates an IPv6 ACL named `ipv6_tcp`. Within the `ipv6_tcp` ACL, an entry with sequence ID 100 is configured. The action is specified as `accept`, with logging set to `true`.

The filter matches packets with IPv6 destination address 2001:db8:1:3::1/120; for TCP traffic, if the source address 2001:db8:1:5::1/120, destination port 6789, and source port 6722 matches the filter, the traffic stream is accepted.

```
--{ * candidate shared default }--[  ]--
# info acl
    ipv6-filter ipv6_tcp {
        entry 100 {
            description Match_Dest_Address_TCP_Src_Address_DP_SP
            action {
                accept {
                    log true
                }
            }
            match {
                destination-address 2001:db8:1:3::1/120
                next-header tcp
                source-address 2001:db8:1:5::1/120
                destination-port {
                    value 6789
                }
                source-port {
                    value 6722
                }
            }
        }
    }
```

3.3.3 Creating a MAC ACL

Procedure

To create a MAC ACL specify statements consisting of match criteria (see Match conditions for MAC ACLs) and actions (see Supported ACL actions for 7220 IXR-D1/D2/D3 systems).

Example: Drop traffic from a source MAC

This example creates a MAC ACL named `mac01`. Within the `mac01` ACL, an entry with sequence ID 100 is configured. The filter matches Ethernet frames with a source MAC address of AA:BB:19:DD:EE:FF. The action is specified as `drop`, with logging set to `true`.

```
--{ * candidate shared default }--[  ]--
# info acl
    acl {
        mac-filter mac01 {
            entry 100 {
                action {
                    drop {
                        log true
                    }
                }
                match {
                    source-mac {
                        address AA:BB:19:DD:EE:FF
                    }
                }
            }
        }
    }
```

Example: Accept traffic using MAC address mask

The following example uses an address mask to specify the source MAC address. The address mask, entered in MAC address format such as ff:ff:ff:ff:00:00, indicates the bit values that must be matched exactly (FF value in the mask) and the bit values that can be 0 or 1 (`00` value in the mask).

```
--{ * candidate shared default }--[  ]--
# info acl
    acl {
        mac-filter mac01 {
            entry 200 {
                action {
                    accept {
                    }
                }
                match {
                    source-mac {
                        address AA:BB:19:DD:EE:FF
                        mask FF:FF:FF:FF:00:00
                    }
                }
            }
        }
    }
```

Example: Drop traffic of specific Ethertype

The following example drops Ethernet frames with ARP Ethertype (`0x0806`).

```
--{ * candidate shared default }--[  ]--
# info acl
    acl {
        mac-filter mac01 {
            entry 300 {
                action {
                    drop {
                        log true
                    }
                }
                match {
                    ethertype 0x0806
                }
            }
        }
    }
```

Example: Drop traffic with specific outermost VLAN ID

You can configure a specific VLAN ID or range of VLAN IDs as match criteria. The match is based on the outermost VLAN ID of the Ethernet frame. The VLAN ID can be specified as an integer from 0 to 4095 or **none**. Specifying 0 as the VLAN ID matches priority-tagged frames; specifying **none** matches untagged frames.

The following example drops Ethernet frames with VLAN ID 4095.

```
--{ * candidate shared default }--[  ]--
# info acl
    acl {
        mac-filter mac01 {
            entry 400 {
                action {
                    drop {
                        log true
                    }
                }
                match {
                    vlan {
                        outermost-vlan-id {
                            value 4095
                        }
                    }
                }
            }
        }
    }
```

The following example configures a range of VLAN IDs, 100 to 999 inclusive, as match criteria.

```
--{ * candidate shared default }--[  ]--
# info acl
    acl {
        mac-filter mac01 {
            entry 500 {
                action {
                    drop {
                        log true
```

```
                    }
                }
                match {
                    vlan {
                        outermost-vlan-id {
                           range {
                                start 100
                                end 999
                            }
                        }
                    }
                }
            }
        }
    }
```

Note the following when specifying VLAN IDs as match criteria:

- When used in an ingress ACL applied to a routed or bridged subinterface, the VLAN ID is matched before the subinterface-defining VLAN tag (of a single-tagged subinterface) has been removed.
- When used in an egress ACL applied to a routed subinterface, the VLAN ID is matched after the subinterface-defining VLAN tag (of a single-tagged subinterface) has been added.
- When used in an egress ACL applied to a bridged subinterface the VLAN ID is matched before the subinterface defining VLAN tag (of a single-tagged subinterface) has been added.

3.3.4 Attaching an ACL to a subinterface

Procedure

After an ACL is configured, you can attach it to a subinterface so that traffic entering or exiting the subinterface is subject to the rules defined in the ACL. For an interface ACL to have an effect on the system, it must be explicitly applied to the input and, or output traffic of at least one subinterface.

Example: Attach IPv4/IPv6 ACLs to a subinterface

The following example applies IPv4 and IPv6 ACLs to inbound traffic on a subinterface:

```
--{ * candidate shared default }--[  ]--
# info interface ethernet-1/16 subinterface 1 acl
    interface ethernet-1/16 {
        subinterface 1 {
            acl {
                input {
                    ipv4-filter ip_tcp
                    ipv6-filter ipv6_tcp
                }
            }
        }
    }
```

Example: Attach a MAC ACL to a subinterface

The following example applies a MAC ACL to outbound traffic on a subinterface. To apply a MAC ACL to traffic in the outbound direction on any subinterface, you must set the **acl egress-mac-filtering** command to **true**.

```
--{ * candidate shared default }--[  ]--
# info interface ethernet-1/1 subinterface 1 acl
    interface ethernet-1/1 {
        subinterface 1 {
            acl {
                output {
                    mac-filter mac01
                }
            }
        }
    }
```

```
--{ * candidate shared default }--[  ]--
# info acl
    acl {
        egress-mac-filtering true
    }
```

3.3.5 Attaching an ACL to the management interface

Procedure

To modulate traffic for the management interface, navigate to the subinterface of interface `mgmt0`. Under the `acl` context, attach the IPv4 or IPv6 ACL for input/output traffic.

Example: Attach an ACL to the mgmt0 interface

```
--{ * candidate shared default }--[  ]--
# interface mgmt0
--{ * candidate shared default }--[ interface mgmt0 ]—
# subinterface 0
--{ * candidate shared default }--[ interface mgmt0 subinterface 0 ]--
# acl
--{ * candidate shared default }--[ interface mgmt0 subinterface 0 acl ]—
# input ipv4-filter ip_tcp
--{ * candidate shared default }--[ interface mgmt0 subinterface 0 acl ]—
# input ipv6-filter ipv6_tcp
```

Example: Verify the configuration

To verify the configuration for the management interface ACL:

```
--{ * candidate shared default }--[  ]--
# info interface mgmt0
    interface mgmt0     {
        admin-state enable
            subinterface 0 {
                admin-state enable
                ipv4 {
                    dhcp-client true
                }
                ipv6 {
```

```
                dhcp-client true
            }
            acl {
                input {
                    ipv4-filter ip_tcp
                    ipv6-filter ipv6_tcp
                }
            }
        }
    }
```

3.3.6 Detaching an ACL from an interface

Procedure

To detach an ACL from an interface, enter the subinterface context and delete the ACL from the configuration.

Example: Detach an ACL from a subinterface

```
--{ * candidate shared default }--[  ]--
# interface ethernet-1/16
--{ * candidate shared default }--[ interface ethernet-1/16 ]–
# subinterface 1
--{ * candidate shared default }--[ interface ethernet-1/16 subinterface 1 ]–
# delete acl input ipv4-filter
--{ * candidate shared default }--[ interface ethernet-1/16 subinterface 1 ]--
```

Example: Verify the configuration

Use the **info interface** command to verify that the ACL is no longer part of the subinterface configuration. For example:

```
--{ * candidate shared default }--[  ]--
# info interface ethernet-1/16
    interface ethernet-1/16 {
        description dut1-dut-2
        subinterface 1 {
            ipv4 {
                address 10.1.5.2/24 {
                }
                arp {
                    neighbor 10.1.5.1 {
                        link-layer-address 00:00:64:01:05:05
                    }
                }
            }
            ipv6 {
                address 2001:db8:1:5::2/120 {
                }
                neighbor-discovery {
                    neighbor 2001:db8:1:5::1 {
                        link-layer-address 00:00:64:01:05:05
                    }
                }
            }
            acl {
                input {
                    ipv6-filter ipv6_tcp
                }
```

```
            }
        }
    }
```

3.3.7 Detaching an ACL from the management interface

Procedure

To detach an ACL from the management interface, enter the mgmt0 subinterface 0 context and delete the ACL from the configuration.

Example: Detach an ACL from the management interface

```
--{ * candidate shared default }--[  ]--
# interface mgmt0
--{ * candidate shared default }--[ interface mgmt0 ]—
# subinterface 0
--{ * candidate shared default }--[ interface mgmt0 subinterface 0 ]--
# delete acl input ipv4-filter
```

Example: Verify the configuration

To verify that the ACL was detached from the management interface:

```
--{ * candidate shared default }--[  ]--
# info interface mgmt0
    interface mgmt0     {
        admin-state enable
            subinterface 0 {
                admin-state enable
                ipv4 {
                    dhcp-client true
                }
                ipv6 {
                    dhcp-client true
                }
                acl {
                    input {
                        ipv6-filter ipv6_tcp
                    }
                }
            }
        }
```

3.3.8 Modifying ACLs

Procedure

You can add entries to an ACL, delete entries from an ACL, and delete the entire ACL from the configuration.

Example: Add an entry to an ACL

To add an entry to an ACL, enter the context for the ACL, then add the entry. For example, the following commands add an entry to IPv4 ACL `ip_tcp`:

```
--{ * candidate shared default }--[  ]--
# acl ipv4-filter ip_tcp
--{ candidate shared default }--[ acl ipv4-filter ip_tcp ]--
# entry 65535
--{ candidate shared default }--[ acl ipv4-filter ip_tcp entry 65535 ]--
# action drop
--{ candidate shared default }--[ acl ipv4-filter ip_tcp entry 65535 action drop ]--
# log true
--{ candidate shared default }--[ acl ipv4-filter ip_tcp entry 65535 action drop ]--
```

Example: Delete an entry in an ACL

To delete an entry in an ACL, use the **delete** command under the context for the ACL and specify the sequence ID of the entry to be deleted. For example, the following commands delete the entry in IPv4 ACL `ip_tcp` with sequence ID 65535:

```
--{ candidate shared default }--[  ]--
# acl ipv4-filter ip_tcp
--{ candidate shared default }--[ acl ipv4-filter ip_tcp ]--
# delete entry 65535
```

Example: Delete an entire ACL

To delete the entire ACL, use the **delete** command under the `acl` context. For example, the following commands delete the `ip_tcp` ACL:

```
--{ * candidate shared default }--[  ]--
# acl
--{ candidate shared default }--[ acl ]--
# delete ipv4-filter ip_tcp
```

3.3.9 Resequencing ACL entries

Procedure

To aid in managing complex ACLs that have many entries, you can resequence the ACL entries to set a sequence ID number for the first entry and a constant increment for the sequence ID for subsequent entries.

For example, if you have an ACL with three entries, sequence IDs 123, 124, and 301, you can resequence the entries so that the initial entry has sequence ID 100, and the other two entries have sequence ID 110 and 120.

Example

The following is an example of an ACL with three entries:

```
--{ candidate shared default }--[ acl ]–
# info ipv4-filter ip_tcp
    ipv4-filter ip_tcp {
        entry 123 {
            action {
```

```
                drop {
                    log true
                }
            }
            match {
                destination-address 10.1.2.1/24
            }
        }
        entry 124 {
            action {
                accept {
                    log true
                }
            }
            match {
                destination-address 10.1.3.1/24
            }
        }
        entry 301 {
            action {
                drop {
                }
            }
            match {
                destination-address 10.1.4.1/24
            }
        }
```

To resequence the entries in the ACL so that the first entry has sequence ID 100, and the next two entries are incremented by 10, enter the context for the ACL, issue the **tools resequence** command, then specify the initial sequence ID and the increment for the subsequent entries. For example:

```
--{ candidate shared default }--[  ]--
# acl ipv4-filter ip_tcp
--{ candidate shared default }--[ acl ipv4-filter ip_tcp ]--
# tools resequence start 100 increment 10
```

After you enter the command, the ACL entries are renumbered. For example:

```
--{ candidate shared default }--[ acl ]—
# info ipv4-filter ip_tcp
    ipv4-filter ip_tcp {
        entry 100 {
            action {
                drop {
                    log true
                }
            }
            match {
                destination-address 10.1.2.1/24
            }
        }
        entry 110 {
            action {
                accept {
                    log true
                }
            }
            match {
                destination-address 10.1.3.1/24
            }
        }
```

```
        entry 120 {
            action {
                drop {
                }
            }
            match {
                destination-address 10.1.4.1/24
            }
        }
```

The **resequence** command is only available inside an ACL configuration context, and it only applies to the entries of the ACL associated with that context.

3.4 Packet capture filters

For troubleshooting purposes, the SR Linux supports ACL policies called packet capture filters. When an ingress IP packet on any line card transits through the router, and it matches a rule in a capture-filter policy, it is copied and extracted toward the CPM (using the capture-filter extraction queue) and delivered to a Linux virtual Ethernet interface, so that it can be displayed by **tcpdump** (or similar packet capture utility), or encapsulated and forwarded to a remote monitoring system.

Similarly, when an ingress IP packet on any line card terminates locally, and it matches a rule of a capture-filter policy, it is extracted toward the CPM (using the normal protocol-based extraction queue), and a header field indicates to the CPM to replicate it (after running the CPM-filter rules) toward the Linux virtual Ethernet interface.

There is one capture-filter for IPv4 traffic and another capture-filter for IPv6 traffic. The default IPv4 capture-filter policy copies no IPv4 packets and the default IPv6 capture-filter copies no IPv6 packets.

The entries for each capture-filter are installed on every line card. On the line card, the entries are evaluated after the input subinterface ACLs and before the CPM-filter ACLs. On the CPM, the entries in the capture-filter policy are evaluated after the CPM-filter entries.

When a capture-filter ACL is created, its rules are evaluated against all transit and terminating IPv4 or IPv6 traffic that is arriving on any subinterface of the router, regardless of where that traffic entered in terms of network instance, subinterface, linecard, pipeline, and so on. Note that capture-filter ACL rules cannot override interface filter or system-filter ACL drop outcomes; packets dropped by interface filter ACLs or a system filter ACL cannot be mirrored to the control plane.

Each capture-filter entry has a set of zero or more match conditions, and one of two possible actions: accept and copy. The match conditions are the same as the other filter types. The accept action passes the matching packet to the next stage of processing, without creating a copy. The copy action creates a copy of the matching packet, extracts it toward the CPM and delivers it to the designated virtual Ethernet interface.

3.5 Control plane module (CPM) filters

For control plane protection, SR Linux supports ACL policies called CPM filters. You can configure one CPM filter that applies to IPv4 traffic, one that applies to IPv6 traffic, and one that applies to non-IP traffic. When ingress traffic is matched by a CPM filter rule, and it is a terminating packet (that is, it must be extracted to the CPM), then it is processed according to the matching CPM filter rule.

The entries for each CPM filter are installed on every line card. They are evaluated after the input subinterface ACLs and after the capture-filter ACLs. CPM filter rules have no effect on locally originating traffic or transit traffic, and they have no interaction with output subinterface ACLs.

When a CPM filter ACL is created, its rules are evaluated against all IPv4 or IPv6 traffic that is locally terminating on the router, regardless of where that traffic entered in terms of network instance, subinterface, linecard, pipeline, and so on.

On 7250 IXR systems, for traffic terminating on the CPM of the router, the following secondary actions are supported for ACL entries where accept is the primary action:

- distributed-policer – If the packet is extracted to the CPM, feed the packet to a hardware-based policer, which determines if the packet should be queued by the line card toward the CPM or dropped because a bit-per-second rate is exceeded.
- system-cpu-policer – If the packet has been extracted to the CPM, feed the packet to a software-based policer, which determines if the packet should be delivered to the CPM application or dropped because a packet-per-second rate is exceeded.

On 7220 IXR-D1, D2, D3, and D5 systems, CPM filter ACLs support the following actions:

- accept – Allow the packet through to the next processing function.
- accept and distributed-policer – The packet is allowed through to the next processing function and rate limited by a policer instance implemented by the 7220 IXR-D2 and D3.
- accept and system-cpu-policer – The packet is allowed through to the next processing function and rate limited by a policer instance implemented by XDP-CPM.
- drop – Discard the packet without ICMP generation.
- drop and log – Discard the packet without ICMP generation and send information about the dropped packet to the log application.

The system-cpu-policer and distributed-policer actions police terminating traffic to ensure that the rate does not exceed a safe limit.

The system-cpu-policer action applies an aggregate rate limit, regardless of ingress line card, while the distributed-policer action applies a rate limit to the extracted traffic from each core (7250 IXR) or complex (7220 IXR) associated with the ingress port. You can have both types of policer actions in the same CPM filter entry, or only one of them.

CPM filter rules that apply a system-cpu-policer or distributed-policer action do not directly specify the policer parameters; they refer to a generically defined policer. This allows different CPM filter entries, even across multiple ACLs, to use the same policer. Optionally, each policer can be configured as entry-specific, which means a different policer instance is used by each referring filter entry, even if they are part of the same ACL.

3.5.1 Creating CPM filters

Procedure

To create CPM filters for IPv4 traffic, IPv6 traffic, or non-IP traffic, you specify statements consisting of match criteria and actions. CPM-bound traffic that matches the match criteria is processed according to the specified action.

Example: Creating an IPv4 CPM filter

The following example creates a CPM filter for IPv4 traffic. IPv4 traffic extracted to the CPM that matches the rule in the filter is processed according to the action configured in the rule. In this example, the matching CPM-bound IPv4 traffic is accepted and rate-limited according to the limit specified by the `sp1` system-cpu-policer.

```
--{ * candidate shared default }--[  ]--
# info acl cpm-filter
    acl {
        cpm-filter {
            ipv4-filter {
                entry 1000 {
                    action {
                        accept {
                            rate-limit {
                                system-cpu-policer 1000
                            }
                        }
                    }
                    match {
                        protocol tcp
                        destination-ip {
                            address 10.1.3.1
                            mask 255.255.255.255
                        }
                        source-ip {
                            address 10.1.5.1
                            mask 255.255.255.255
                        }
                        destination-port {
                            value 6789
                        }
                        source-port {
                            value 6722
                        }
                    }
                }
            }
        }
    }
```

```
--{ * candidate shared default }--[  ]--
# info acl policers system-cpu-policer sp1
    acl {
        policers {
            system-cpu-policer sp1 {
                peak-packet-rate 1000000
                max-packet-burst 16
            }
        }
    }
```

Example: Creating a MAC CPM filter

The following example creates a MAC CPM filter. Non-IP traffic extracted to the CPM that matches the filter rule is processed according to the configured action. In this example, the matching CPM-

bound Ethernet frames are accepted and rate-limited according to the limit specified by the `dp1` distributed-cpu-policer setting.

```
--{ * candidate shared default }--[  ]--
# info acl cpm-filter
    acl {
        cpm-filter {
            mac-filter {
                entry 100 {
                    action {
                        accept {
                            rate-limit {
                                distributed-policer dp1
                            }
                        }
                    }
                    match {
                        source-mac {
                            address 00:00:5e:00:53:FF
                            mask 00:00:5e:00:53:00
                        }
                    }
                }
            }
        }
    }
```

```
--{ * candidate shared default }--[  ]--
# info acl policers policer dp1
    acl {
        policers {
            policer dp1 {
                peak-rate 500000
                max-burst 100000
            }
        }
    }
```

3.6 System filters

A system filter ACL is an IPv4 or IPv6 ACL that evaluates ingress traffic before all other ACL rules. If an IP packet is dropped by a system filter rule, it is the final disposition of the packet; neither a capture-filter copy/accept action, nor an ingress interface ACL accept action, nor a CPM-filter accept action can override the drop action of a system filter.

At most one system filter can be defined for IPv4 traffic, and at most one system filter can be defined for IPv6 traffic. System filter ACLs are supported on 7220 IXR-D1, D2, and D3 systems only. They can be applied only at ingress, not egress.

When a system-filter ACL is created, its rules are automatically installed everywhere, meaning they are evaluated against all transit and terminating IPv4 or IPv6 traffic arriving on any subinterface of the router, regardless of where that traffic entered in terms of network instance, subinterface, pipeline, and so on.

A system filter is the only type of filter that can match the outer header of tunneled packets. For VXLAN traffic, this allows you to configure a system filter that matches and drops unauthorized VXLAN tunnel packets before they are decapsulated. The system filter matches the outer header of tunneled packets; they do not filter the payload of VXLAN tunnels.

3.6.1 Creating a system filter

Procedure

When you apply a system filter, it evaluates all transit and terminating IPv4 arriving on any subinterface of the router. The system filter ACL evaluates the traffic before any other ACL filters. System filter ACLs can be configured on 7220 IXR-D1, D2, and D3 systems only.

Example

The following is an example of a system filter ACL that filters IPv4 traffic.

```
--{ * candidate shared default }--[  ]--
# info acl system-filter
    acl
       system-filter {
            ipv4-filter {
                entry 44 {
                    action {
                        drop {
                            log true
                        }
                    }
                    match {
                        source-ip {
                            address 10.1.5.1
                            mask 10.0.0.255
                        }
                    }
                }
            }
        }
```

3.7 Configuring logging for ACLs

You can configure the SR Linux to log information about packets that match an ACL entry in the system log.

You can set thresholds for ACL or TCAM resource usage. When utilization of a specified resource reaches the threshold in either the rising or falling direction, it can trigger a log message.

3.7.1 Enabling syslog for the ACL subsystem

Procedure

If you set the **log** parameter to **true** for the accept or drop action in an ACL entry, information about packets that match the ACL entry is recorded in the system log. You can specify settings for the log file for the ACL subsystem, including the location of the log file, maximum log file size, and the number of log files to keep.

Example

The following configuration specifies that the log file for the ACL subsystem be stored in the file `dut1_file`, located in the `/opt/srlinux/bin/logs/srbase` directory. The log file can be a

maximum of 1 Mb. When the log file reaches this size, it is renamed using `dut1_file` as its base name. The five most recent log files are kept.

```
--{ candidate shared default }--[  ]--
# system
--{ candidate shared default }--[ system ]--
# info logging file dut1_file
    logging {
        file dut1_file {
            directory /opt/srlinux/bin/logs/srbase/
            rotate 5
            size 1M
            subsystem acl {
            }
        }
    }
```

Ensure that write permission is set for the specified directory path.

3.7.1.1 Syslog entry examples

The following are examples of syslog entries for ACLs.

IPv4 Accept:

```
acl||I Type: Ingress IPv4 Filter: testing Sequence Id: 100 Action: Accept Interface: ethernet-
1/16:1 Packet length: 56 IP Source: 10.1.5.1 Destination: 100.1.3.1 Protocol: 6 TCP Source
 port: 6722 Destination Port: 6789 Flags: SYN
```

IPv4 Drop:

```
acl||I Type: Ingress IPv4 Filter: test Sequence Id: 65535 Action: Drop Interface: ethernet-1/
16:1 Packet length: 44 IP Source: 10.3.2.3 Destination: 100.1.3.1 Protocol: 17 UDP Source port:
 6722 Destination Port: 6789
```

IPv6 Accept:

```
acl||I Type: Ingress IPv6 Filter: tests Sequence Id: 1000 Action: Accept Interface: ethernet-
1/16:1 Packet length: 76 IP Source: 2001:db8:1:5::1 Destination: 2001:10:1:3::1 Protocol: 6 TCP
 Source port: 6722 Destination Port: 6789 Flags: SYN
```

3.7.2 Logging ACL resource usage

Procedure

You can set thresholds for ACL resource usage. When utilization of a specified ACL resource, such as input IPv4 filter instances, reaches the threshold in either the rising or falling direction, it can trigger a log message.

Example

The following example sets thresholds for resource usage by input IPv4 filter instances. If the resource usage percentage falls below the `falling-threshold-log` value, a log message

of priority `notice` is generated. If the resource usage percentage falls below the `rising-threshold-log` value, a log message of priority `warning` is generated.

```
--{ * candidate shared default }--[  ]--
# info platform resource-monitoring
    platform {
        resource-monitoring {
            acl {
                resource input-ipv4-filter-instances {
                    rising-threshold-log 90
                    falling-threshold-log 90
                }
            }
        }
    }
```

3.7.3 Logging TCAM resource usage

Procedure

You can set thresholds for Ternary Content Addressable Memory (TCAM) resource usage. When utilization of a specified TCAM resource, such as TCAM used by IPv4 CPM filters, reaches the threshold in either the rising or falling direction, it can trigger a log message.

Example

The following example sets thresholds for TCAM resource usage by IPv4 CPM filters. If the resource usage percentage falls below the `falling-threshold-log` value, a log message of priority `notice` is generated. If the resource usage percentage falls below the `rising-threshold-log` value, a log message of priority `warning` is generated.

```
--{ * candidate shared default }--[  ]--
# info platform resource-monitoring
    platform {
        resource-monitoring {
            tcam {
                resource cpm-capture-ipv4 {
                    rising-threshold-log 90
                    falling-threshold-log 90
                }
            }
        }
    }
```

3.8 Collecting and displaying ACL statistics

The SR Linux can collect statistics for packets matching an ACL and display statistics for the matched packets. You can display the amount of system resources (TCAM) used by each type of ACL on each line card. ACL statistics can also be displayed using **show** commands.

3.8.1 Collecting ACL statistics

Procedure

You can configure an ACL to collect statistics for packets matching the ACL. Statistics can be collected for packets that match each ACL entry, as well as for matching input/output traffic per subinterface.

Example

The following example configures the ACL to record the number of matching packets for each entry:

```
--{ candidate shared default }--[ acl ]--
# ipv4-filter ip_tcp
--{ candidate shared default }--[ acl ipv4-filter ip_tcp ]--
# statistics-per-entry true
```

Example

By default, if two or more subinterfaces on the same line card reference the same ACL for filtering the same direction of traffic, they use a shared instance of the same ACL in hardware. This means that per-entry statistics (including the number of matched packets and the time stamp of the last matching packet), if enabled, reflect the aggregate of the data gathered for the multiple subinterfaces.

To collect per-entry, per-subinterface statistics, instead of the aggregate of the subinterfaces where the ACL is applied, you can configure an ACL to operate in subinterface-specific mode.

If you change an ACL from subinterface-specific mode to shared mode, or the other way around, during the transition from one mode to the next, traffic continues to be subject to the previous mode until the system resources (TCAM) entries are programmed for the new mode.

The following example configures the ACL to collect statistics for matching packets inbound and outbound on each subinterface:

```
--{ candidate shared default }--[ acl ]--
# ipv4-filter ip_tcp
--{ candidate shared default }--[ acl ipv4-filter ip_tcp ]--
# subinterface-specific input-and-output
```

You can configure the following values for the **subinterface-specific** parameter:

- **disabled** (the default) – All subinterfaces on a single line card that reference the ACL as an input ACL use a shared filter instance, and all subinterfaces on a single line card that reference the ACL as an output ACL use a shared filter instance.
- **input-only** – All subinterfaces on a single line card that reference the ACL as an output ACL use a shared filter instance, but each subinterface that references the ACL as an input ACL uses its own separate instance of the filter.
- **output-only** – All subinterfaces on a single line card that reference the ACL as an input ACL use a shared filter instance, but each subinterface that references the ACL as an output ACL uses its own separate instance of the filter.
- **input-and-output** – Each subinterface that references the ACL as either an input ACL or an output ACL uses its own separate instance of the filter.

3.8.2 Displaying ACL statistics

Procedure

Use the **info from state** command to display the matched packet statistics and the time of the last match for the interfaces to which the ACL is attached.

Example

In the following example, the ACL is attached to two interfaces, and statistics are collected for each subinterface:

```
--{ candidate shared default }--[ acl ipv4-filter ip_tcp ]--
# info from state
    subinterface-specific input-and-output
    statistics-per-entry true
    entry 1000 {
        description Match_IP_Address_TCP_Protocol_Ports
        action {
            accept {
                log true
            }
        }
        match {
            destination-address 10.1.3.1/32
            protocol tcp
            source-address 10.1.5.1/32
            destination-port {
                value 6789
            }
            source-port {
                value 6722
            }
        }
        statistics {
            aggregate {
                in-matched-packets 3000
                in-last-match 2019-07-16T10:53:00.1563Z
                out-matched-packets 0
            }
            per-interface {
                subinterface ethernet-1/16.1 {
                    in-matched-packets 3000
                    in-last-match 2019-07-16T10:53:00.1563Z
                }
            }
        }
    }
    entry 65535 {
        action {
            drop {
                log true
            }
        }
        statistics {
            aggregate {
                in-matched-packets 1000
                in-last-match 2019-07-16T10:53:30.1563Z
            }
            per-interface {
                subinterface ethernet-1/16.1 {
                    in-matched-packets 1000
                    in-last-match 2019-07-16T10:53:30.1563Z
```

```
                }
            }
        }
    }
```

If the match criteria changes for an ACL entry, the statistics counter does not reset to zero. To reset the statistics counter for an ACL entry to zero, use the **tools acl clear** command, as described in Clearing ACL statistics.

3.8.3 Displaying ACL resource usage

Procedure

Use the **info from state platform** command to display the amount of system resources (TCAM) used by each type of ACL on each line card.

Example: Display free-static and free-dynamic resources

The following example shows two different numbers for the remaining (free) TCAM entry resources that are available for input IPv4 ACLs on the line card. The free-static value refers to the available number of resources assuming no additional TCAM banks are dynamically assigned to the ACL type. The free-dynamic value refers to the available number of resources assuming all unallocated TCAM banks are dedicated to the specified ACL type.

```
--{ candidate shared default }--[  ]--
# info from state platform linecard 1 forwarding-complex 0 tcam resource if-input-ipv4
    platform {
        linecard 1 {
            forwarding-complex 0 {
                tcam {
                    resource if-input-ipv4 {
                        free-static 2046
                        free-dynamic 18430
                        reserved 2
                        programmed 2
                    }
                }
            }
        }
    }
```

Example: Display system resources allocated to input IPv4 ACLs

The following example shows the amount of system resources allocated to input IPv4 ACLs on the line card and how much is used and free:

```
--{ candidate shared default }--[  ]--
# info from state platform linecard 1 forwarding-complex 0 acl resource input-ipv4-filter-
instances
    platform {
        linecard 1 {
            forwarding-complex 0 {
                acl {
                    resource input-ipv4-filter-instances {
                        used 1
                        free 254
                    }
                }
```

```
            }
        }
    }
```

3.8.4 Clearing ACL statistics

Procedure

To reset ACL statistics counters to zero, use the **tools acl clear** command. This command can clear statistics at the IPv4 / IPv6 / CPM filter level, ACL entry level, or for an interface or subinterface to which the ACL is attached.

Example: Clear statistics for an IPv4 filter

```
--{ candidate shared default }--[ acl ]--
# tools acl ipv4-filter tcp_ip clear
```

Example

After this command is executed, the **info from state** output for the IPv4 filter includes a timestamp indicating when the statistics were cleared. For example:

```
--{ candidate shared default }--[ acl ipv4-filter tcp_ip ]--
# info from state
    subinterface-specific output-only
    statistics-per-entry true
    entry 1000 {
        description Match_IP_Address_TCP_Protocol_Ports
        action {
            accept {
                log true
            }
        }
        match {
            destination-address 100.1.3.1/32
            protocol tcp
            source-address 10.1.5.1/32
            destination-port {
                value 6789
            }
            source-port {
                value 6722
            }
        }
        statistics {
            last-clear 2021-10-03T13:53:51.000Z
        }
    }
```

Example: Clears statistics for a specific IPv4 filter entry

The following example clears statistics for a specific entry in the IPv4 filter:

```
--{ candidate shared default }--[ acl ]--
# tools acl ipv4-filter tcp_ip entry 1000 statistics clear
```

Example: Clears statistics for a subinterface for a IPv4 filter entry

The following example clears statistics for a specified subinterface for a specified entry in the IPv4 filter:

```
--{ candidate shared default }--[ acl ]--
# tools acl ipv4-filter tcp_ip entry 1000 statistics per-interface subinterface 1 clear
```

3.8.5 Displaying ACL statistics using show commands

Procedure

You can display ACL statistics using relevant **show** commands.

Example: Display all active ACLs

To display information about all active ACLs, use the **show acl summary** command. For example:

```
--{ candidate shared default }--[  ]--
# show acl summary
-------------------------------------------------------------------------------
CPM Filter ACLs
-------------------------------------------------------------------------------
ipv4-entries: 1
ipv6-entries: 0
mac-entries : 1
-------------------------------------------------------------------------------
Capture Filter ACLs
-------------------------------------------------------------------------------
ipv4-entries: 0
ipv6-entries: 0
-------------------------------------------------------------------------------
IPv4 Filter ACLs
-------------------------------------------------------------------------------
Filter : ip_tcp
Active on : 1 subinterfaces (input) and 0 subinterfaces (output)
Entries : 2
-------------------------------------------------------------------------------
IPv6 Filter ACLs
-------------------------------------------------------------------------------
Filter : ipv6_tcp
Active on : 1 subinterfaces (input) and 0 subinterfaces (output)
Entries : 1
-------------------------------------------------------------------------------
MAC Filter ACLs
-------------------------------------------------------------------------------
Filter   : mac01
Active On: 1 subinterfaces (input) and 0 subinterfaces (output)
Entries  : 5
-------------------------------------------------------------------------------
```

Example: Display statistics for a specific ACL

You can display statistics for a specific ACL, including how many times each ACL entry was matched on all subinterfaces to which the ACL was applied. For example:

```
--{ candidate shared default }--[  ]--
# show acl ipv4-filter ip_tcp
===============================================================================
```

```
Filter        : ip_tcp
SubIf-Specific: input-and-output
Entry-stats   : yes
Entries       : 2
-------------------------------------------------------------------------------
 Subinterface     Input   Output
ethernet-1/16.1   yes     no
-------------------------------------------------------------------------------
Entry 1000
  Match               : protocol=tcp, 10.1.5.1/32(6722-6722)->10.1.3.1/32(6789-
6789)
  Action              : accept
  Input Match Packets : 3000
  Input Last Match    : 18 seconds ago
  Output Match Packets: 0
  Output Last Match   : never
Entry 65535
  Match               : protocol=<undefined>, any(*)->any(*)
  Action              : drop
  Input Match Packets : 1000
  Input Last Match    : 6 minutes ago
  Output Match Packets: 0
  Output Last Match   : never
```

Example: Display per-interface statistics for each ACL entry

To display per-interface statistics for packets matching each ACL entry, specify the interface name in addition to the ACL name. For example:

```
--{ candidate shared default }--[  ]--
# show acl ipv4-filter ip_tcp interface ethernet-1/16
===============================================================================
Filter        : ip_tcp
SubIf-Specific: input-and-output
Entry-stats   : yes
Entries       : 2
-------------------------------------------------------------------------------
 Subinterface     Input   Output
ethernet-1/16.1   yes     no
-------------------------------------------------------------------------------
Entry 1000
  Match               : protocol=tcp, 10.1.5.1/32(6722-6722)->10.1.3.1/32(6789-
6789)
  Action              : accept
  Input Match Packets : 3000
  Input Last Match    : 5 minutes ago
  Output Match Packets: 0
  Output Last Match   : never
Entry 65535
  Match               : protocol=<undefined>, any(*)->any(*)
  Action              : drop
  Input Match Packets : 1000
  Input Last Match    : 10 minutes ago
  Output Match Packets: 0
  Output Last Match   : never
```

Example: Display statistics for a CPM filter

To display statistics for packets matching a CPM filter, specify the CPM filter type (IPv4 or IPv6). For example:

```
--{ candidate shared default }--[  ]--
```

```
# show acl cpm-filter ipv4-filter
===============================================================================
Filter     : CPM IPv4-filter
Entry-stats: no
Entries    : 1
-------------------------------------------------------------------------------
Entry 1001
  Match          : protocol=<undefined>, any(*)->10.1.2.1/24(*)
  Action         : none
  Matched Packets: 0
-------------------------------------------------------------------------------
```

4 Policy-based forwarding

Policy-based forwarding (PBF) supports traffic forwarding in a network instane based on match conditions and actions defined in a policy, as an alternative to forwarding based on entries in a routing table.

Each PBF policy is modeled as a sequence of rules, each of which has match conditions and actions. Match conditions specify values for various packet header fields. A packet matches a rule only if all the match conditions evaluate to true. Actions specify the processing to apply to each matching packet.

Each PBF policy is associated with a specific network instane. The PBF rules only apply to the ingress IP packets on selected routed subinterfaces of the network instane. Policy-forwarded packets are classified according to the DSCP policy that is attached to the ingress subinterface.

Match conditions for PBF policies

The following match conditions can be specified in a PBF policy:

- **dscp-set** – list of DSCP values to match for incoming packets; a packet must match one of the DSCP values defined in this list for the rule to apply
- **protocol** – protocol carried in the IP packet, specified either by name or IP protocol value
- **source-ip** – source IP address of the IP packet; for an IP-in-IP packet; this refers to the outer IP header source address

Actions for PBF policies

You can specify the following action in a PBF policy:

- **network-instance** – Look up matching packets in the network instane referenced in the PBR policy instead of the network instane associated with the subinterface.

4.1 Creating a PBF policy

Procedure

To create a PBF policy, configure the match conditions for the policy and the action to take for packets that meet the match conditions.

Example: Match based on IPv4 protocol value

The following example configures a PBF policy that applies to the default network instane. On subinterfaces where this policy is applied, incoming IPv4 packets that have a value of 4 in their IP protocol field are looked up and forwarded in network instane red.

```
--{ candidate shared default }--[  ]--
# info network-instance default policy-forwarding
    network-instance default {
        policy-forwarding {
            policy 100 {
                description "Sample PBF Policy"
                rule 1 {
                    action {
```

```
                        network-instance red
                    }
                    match {
                        ipv4 {
                            protocol 4
                        }
                    }
                }
            }
        }
    }
```

Example: Match based on DSCP values

In the following example, incoming packets matching DSCP values 0, 1, or 2 are looked up and forwarded in network-instance blue:

```
--{ * candidate shared default }--[  ]--
# info network-instance default policy-forwarding
    network-instance default {
        policy-forwarding {
            policy 101 {
                rule 1 {
                    action {
                        network-instance blue
                    }
                    match {
                        ipv4 {
                            dscp-set [
                                0
                                1
                                2
                            ]
                        }
                    }
                }
            }
        }
    }
```

Example: Match based on source IP prefix

In the following example, incoming packets whose source IP address matches prefix 10.10.0.0/16 are looked up and forwarded in network-instance green:

```
--{ candidate shared default }--[  ]--
# info network-instance default policy-forwarding
    network-instance default {
        policy-forwarding {
            policy 100 {
                rule 1 {
                    action {
                        network-instance green
                    }
                    match {
                        ipv4 {
                            source-ip {
                                prefix 10.10.0.0/16
                            }
                        }
                    }
                }
```

```
            }
        }
    }
```

4.2 Applying a PBF policy

Procedure

To activate a PBF policy, apply the policy to one or more routed subinterfaces of the network instane configured in the policy.

Example

The following example applies a PBF policy to a subinterface in the default network-instance. The system evaluates ingress packets on the subinterface according to the match conditions in the policy and forwards the matching packets according to the action specified in the policy.

```
--{ candidate shared default }--[  ]--
# info network-instance default policy-forwarding
    network-instance default {
        policy-forwarding {
            interface ethernet-1/1.1 {
                apply-forwarding-policy 100
            }
        }
    }
```

5 TCAM allocation on SR Linux devices

Ternary Content Addressable Memory (TCAM) on SR Linux devices can be allocated to system resources either statically or dynamically, based on configuration requirements.

The following sections describe static and dynamic allocation for TCAM banks (known as TCAM slices) and provide details for how SR Linux allocates TCAM on the following devices:

- 7220 IXR-D1
- 7220 IXR-D2 and D3
- 7220 IXR-D4 and D5
- 7250 IXR-6/10 and IXR-6e/10e

Static TCAM allocation

Static TCAM refers to TCAM slices that are allocated at initialization and remain constantly reserved for their allocated purpose. Each static TCAM slice has a type, which refers to the type or types of entries that it stores. Examples of different entry types are IPv4 ingress ACL, IPv6 ingress ACL, and so on. A static TCAM slice consumes resources even if has no entries.

When the system is initialized (for example, after a restart) static TCAM slices are allocated first, before the configuration is evaluated and dynamic TCAM slices are created.

Dynamic TCAM allocation

Dynamic TCAM refers to TCAM slices that are taken from a free or unused pool of TCAM slices and released back to that pool when configuration determines that they are no longer needed.

Each dynamic TCAM slice has a type, which refers to the type or types of entries that it stores. When the configuration of the device has no entries of the type associated with a particular dynamic TCAM slice, there are no allocated slices of that type. The device does not hold onto empty TCAM banks to guarantee a minimum reservation.

The addition of the first entry associated with a particular dynamic TCAM slice attempts to create the required group of slices (slice group). Depending on the entry type (and resulting TCAM key length) the size of the slice group can be one TCAM slice (single-wide TCAM slice type), two TCAM slices (double-wide TCAM slice type), or three TCAM slices (triple-wide TCAM slice type). If a configuration change adds many entries of a new type, multiple slice groups may be needed.

Subsequent configuration changes may add more entries of a type that already has dynamic TCAM slices allocated. When the net number of new entries of that type (that is, additional entries minus deleted entries) exceeds the limit of the current allocation (after filling all empty holes in existing allocated slice groups), the system attempts to allocate new slice groups. There is no artificially imposed limit on the maximum number of slices that can be consumed for an entry type.

If a configuration change requires additional dynamic TCAM slices, and there are not enough free slices to accommodate the new slice groups, even if existing groups are moved to satisfy system constraints on the placement of groups, then the configuration commit is blocked.

If a configuration change requires additional dynamic TCAM slices, and there are enough free slices to accommodate the new slice groups, but only if existing groups are moved to satisfy system constraints on the placement of groups, then the commit is allowed and the TCAM slice relocation operations are initiated

automatically. Moving a bank to a new location can take two or three minutes, so an entire operation involving multiple banks could delay the programming of the new entries for up to 20 minutes or more.

The deletion of the last entry associated with a particular dynamic slice group deletes the group and makes the freed slices available to other slice types. If a single configuration commit deletes entries of one type and adds entries of another type, the deleted entries (and any reclamation of slices) are processed first, so that there is a greater chance of successfully allocating the additional slices.

The deletion of some (but not all) entries of a specific type that already has dynamic TCAM slices allocated can leave empty entries in these TCAM slices. This could create a situation where there are N allocated slices for a specific entry type, but considering the empty entries in each of these N slices, fewer than N slices would be needed if the used entries could be moved between slices to consolidate them. This process, called slice compaction, is done automatically during the processing of a configuration transaction that results in a net deletion of entries; it should not be service impacting.

The maximum number of statistics resources that is available to dynamic TCAM entry types remains fixed, and does not scale to the maximum possible size for an entry type. If an ACL rule is configured to have per-entry statistics, but a statistics index resource cannot be allocated, this is indicated by `statistics-resource-allocated = false` in the YANG state for the entry.

Displaying static and dynamic TCAM usage

Use the **info from state platform** command to display the number of free static and dynamic TCAM entries available to each type of ACL. For example:

```
--{ running }--[  ]--
# info from state platform linecard 1 forwarding-complex 0 tcam resource *
    platform {
        linecard 1 {
            forwarding-complex 0 {
                tcam {
                    resource if-input-ipv4 {
                        free-static 6912
                        free-dynamic 0
                        reserved 0
                        programmed 0
                    }
                    resource if-input-ipv4-qos {
                        free-static 6904
                        free-dynamic 0
                        reserved 8
                        programmed 8
                    }
                    resource if-input-ipv6 {
                        free-static 2304
                        free-dynamic 0
                        reserved 0
                        programmed 0
                    }
                    resource if-input-ipv6-qos {
                        free-static 2262
                        free-dynamic 0
                        reserved 42
                        programmed 42
                    }
                    resource if-input-mac {
                        free-static 2304
                        free-dynamic 0
                        reserved 0
                        programmed 0
                    }
```

```
                    resource if-input-policer {
                        free-static 1536
                        free-dynamic 0
                        reserved 0
                        programmed 0
                    }
                    resource if-output-cpm-ipv4 {
                        free-static 1536
                        free-dynamic 0
                        reserved 0
                        programmed 0
                    }
                    resource if-output-cpm-ipv6 {
                        free-static 512
                        free-dynamic 0
                        reserved 0
                        programmed 0
                    }
                    resource if-output-cpm-mac {
                        free-static 1536
                        free-dynamic 0
                        reserved 0
                        programmed 0
                    }
                    resource system-capture-ipv4 {
                        free-static 768
                        free-dynamic 0
                        reserved 0
                        programmed 0
                    }
                    resource system-capture-ipv6 {
                        free-static 384
                        free-dynamic 0
                        reserved 0
                        programmed 0
                    }
                }
            }
        }
    }
```

The `free-static` statistic is the number of free entries assuming the number of TCAM slices that are currently allocated to the type of entry remains constant.

The `free-dynamic` statistic is the number of free entries that are possible if all remaining unused TCAM slices are dynamically assigned to this one type of entry.

5.1 TCAM allocation on 7220 IXR-D1

On the 7220 IXR-D1, all TCAM slices are allocated statically. The 7220 IXR-D1 does not support dynamic TCAM allocation. The 7220 IXR-D1 has 3 groups of TCAM slices associated with 3 different stages of the forwarding pipeline. Each of these groups is a separate resource pool.

TCAM allocation details for each stage are summarized in the following table.

Table 6: 7220 IXR-D1 TCAM allocation

Stage	TCAM allocation details
VFP	There are 4 VFP slices. Each VFP slice provides 512 entries. IPv4 capture-filter entries require a single slice of single-width entries, providing a maximum of 512 entries per slice. One bank is statically allocated, providing scaling of 512 entries. IPv6 capture-filter entries require a single slice of double-width entries, providing a maximum of 256 entries per slice. One bank is statically allocated, providing scaling of 256 entries. One bank is unused.
IFP	There are 18 IFP slices. Each IFP slice provides 512 entries. Ingress IPv4 filter entries require a single slice of single-width entries, providing a maximum of 512 entries per slice. Two banks are statically allocated, providing scaling of 1024 entries. Ingress IPv6 filter entries require a triple-wide slice, providing a maximum of 512 entries per triple-wide slice. Six banks are statically allocated, providing scaling of 1024 entries. Ingress MAC filter entries require a double-wide slice, providing a maximum of 512 entries per double-wide slice. Zero banks are statically allocated. Six banks are unused.
EFP	There are 4 EFP slices. Each EFP slice provides 256 entries. Egress IPv4 and IPv4 CPM filter entries require a single slice of single-width entries, providing a maximum of 256 entries per slice. One bank is statically allocated, providing scaling of 256 entries. Egress IPv6 and IPv6 CPM filter entries require a double-wide slice, providing a maximum of 256 entries per double-wide slice. Two banks are statically allocated, providing scaling of 256 entries. Egress MAC filter entries require a single-wide slice, providing a maximum of 256 entries per slice. Zero banks are statically allocated. One bank is unused.

5.2 TCAM allocation on 7220 IXR-D2 and D3

The 7220 IXR-D2 and 7220 IXR-D3 have 3 groups of TCAM slices associated with 3 different stages of the forwarding pipeline. Each of these groups is a separate resource pool. A free slice in one pool is not available to an entry type associated with a different stage of the pipeline. For example, freeing an IFP bank does not provide one more available EFP bank.

The details of each stage in terms of supported dynamic TCAM entry types, total number of TCAM slices, and number of pre-reserved static TCAM slices (with their associated entry types) is summarized in the following table.

Table 7: 7220 IXR-D2/D3 TCAM allocation

Stage	TCAM allocation details
VFP	Lookup happens after `MY_STATION` lookup, before tunnel encapsulation (if any) is removed. Used to assign a virtual port (VP) to packets arriving on an untagged bridged subinterface. Also used for capture and system filters. There are 4 VFP slices. Each VFP slice provides 256 entries indexed by a 234-bit key or 128 entries indexed by a 468-bit key (intra-slice double-wide mode).
	1 slice is allocated statically. Entries serve 2 purposes: • to strip the transport VLAN 1 tag from outbound CPU-originated packets (to support egress mirroring of such traffic). This requires 1 entry per system. • to assign a VP to packets arriving on an untagged bridged subinterface. This requires 1 entry per untagged bridged subinterface.
	3 slices are available for dynamic allocation: • Capture IPv4 TCAM entries and system IPv4 TCAM entries can share a slice supporting up to 256 entries; maximum possible scale is 768 entries. • Capture IPv6 TCAM entries and system IPv6 TCAM entries can share a slice supporting up to 128 entries; maximum possible scale is 384 entries.
IFP	Lookup happens after QoS classification, tunnel decapsulation, and FIB lookup. Used for ingress interface ACLs, ingress subinterface policing, ingress MF QoS classification, VXLAN ES functionality, and CPM extraction (CPU QoS queue assignment). There are 12 IFP slices. Each IFP slice provides 768 entries indexed by a 160-bit key (intra-slice double-wide mode).
	4 slices are allocated statically: • 2 slices for CPU QoS queue assignment • 1 slice for VXLAN ES related functionality • 1 slice for ingress subinterface policing, supporting 1536 entries (intra-slice single-wide mode)
	8 slices are available for dynamic allocation: • 768 ingress IPv4 ACL filter entries are supported per bank (intra-slice double-wide mode). There are no restrictions on the placement of intra-slice double-wide slices. Maximum possible ingress IPv4 TCAM entries = 8*768 • 768 ingress IPv6 ACL filter entries are supported by 3 consecutive banks (triple-wide mode). There is a virtual boundary between every group of 3 slices. The end of a triple-wide group cannot cross any of these virtual boundaries. Maximum possible ingress IPv6 TCAM entries = 2*768 • 768 ingress MAC ACL filter entries are supported by 2 consecutive banks (double-wide mode). There is a virtual boundary between every group of 3 slices. The end of a double-wide group cannot cross any of these virtual

Stage	TCAM allocation details
	boundaries but the start of the double-wide group does not have to align with a 3-slice boundary. Maximum possible ingress MAC TCAM entries = 3*768 • 768 ingress IPv4 MF QoS classification entries are supported per bank (intra-slice double-wide mode). There are no restrictions on the placement of intra-slice double-wide slices. Maximum possible ingress IPv4 TCAM entries = 8*768 • 768 ingress IPv6 MF QoS classification entries are supported by 3 consecutive banks (triple-wide mode). There is a virtual boundary between every group of 3 slices. The end of a triple-wide group cannot cross any of these virtual boundaries. Maximum possible ingress IPv6 TCAM entries = 2*768
EFP	Lookup happens before final packet modification, after CoS rewrite. Used for out-mirror stats, egress interface ACLs and CPM filter ACLs. There are 4 EFP slices. Each EFP slice provides 512 entries indexed by a 272-bit key.
	1 slice is allocated statically. Entries serve 2 purposes: • ES pruning of local-biased traffic. This requires 1 entry per system. • Egress port mirroring stats. This requires 1 entry per outgoing interface.
	3 slices are available for dynamic allocation: • interface egress IPv4 TCAM entries and CPM-filter IPv4 TCAM entries share a single-wide slice supporting up to 512 entries; maximum possible scale is 1536 entries • interface egress IPv6 TCAM entries and CPM-filter IPv6 TCAM entries share a double-wide slice supporting up to 512 entries; maximum possible scale is 512 entries. IPv6 TCAM entries cannot be added unless all 3 dynamic TCAM banks are free at the time of adding the first IPv6 entry. If there is already an IPv4 or MAC TCAM bank that has been allocated, no IPv6 entries are permitted. • interface egress MAC TCAM entries and CPM-filter MAC TCAM entries share a single-wide slice supporting up to 512 entries; maximum possible scale is 1536 entries
	XGS has limitations expanding an EFP slice when the entries have policers; therefore the following restriction is imposed: If a single-wide IPv4 slice has been created, and it has entries with policers (for example. CPM IPv4 filter entries) or entries with a drop and log action, it is not possible to expand the number of IPv4 slices beyond this single slice; conversely, if the number of IPv4 slices was allowed to extend to 2 or more it is not possible to attach a policer or add a drop and log action to any entries in the expanded set of slices

5.3 TCAM allocation on 7220 IXR-D4 and D5

The 7220 IXR-D4 and 7220 IXR-D5 have 3 groups of TCAM slices associated with 3 different stages of the forwarding pipeline.

TCAM allocation details for each stage is summarized in the following table.

Table 8: 7220 IXR-D4/D5 TCAM allocation

Stage	TCAM allocation details
VFP	There are 4 VFP slices, providing a total of 4*256 = 1024 entries. All slices are allocated statically.
	1 slice: • to strip the transport VLAN 1 tag from outbound-CPU-originated packets (to support egress mirroring of such traffic). This requires 1 entry per system. • to assign a virtual port (VP) to packets arriving on an untagged bridged subinterface. This requires 1 entry per untagged bridged subinterface.
	3 slices for capture and system filter entries: • Three physical slices are grouped to support triple-wide entries. • One triple-wide entry is consumed by each IPv4 capture-filter entry, IPv4 system filter entry, IPv6 capture-filter entry, or IPv6 system filter entry (the system supports mapping IPv4 groups and IPv6 groups onto the same physical slices); maximum possible scale is 256 entries.
IFP	There are 12 IFP slices on 7220 IXR-D4/D5, providing a total of 12*2K = 24K entries.
	On 7220 IXR-D5, all slices are allocated statically: • 2 slices for CPU QoS queue assignment • 1 slice for VXLAN ES-related functionality • 2 slices for ingress IPv4 ACL 2K ingress IPv4 ACL entries are supported by group of 2 side-by-side banks (inter-slice double-wide mode). Maximum possible ingress IPv4 TCAM entries = 2048 • 3 slices for ingress IPv6 ACL 2K ingress IPv6 ACL entries are supported by group of 3 side-by-side banks (inter-slice triple-wide mode). Maximum possible ingress IPv6 TCAM entries = 2048 • 2 slices for MAC ACL 2K ingress MAC ACL entries are supported by group of 2 side-by-side banks (inter-slice double-wide mode). Maximum possible ingress MAC TCAM entries = 2048. • 1 slice for ingress subinterface policing
	On 7220 IXR-D5, all slices are allocated statically:

Stage	TCAM allocation details
	• 2 slices for CPU QoS queue assignment, providing 2048 entries • 1 slice for VXLAN ES related functionality, providing 2048 entries • 2 slices for ingress IPv4 ACLs, providing 1024 entries • 3 slices for ingress IPv6 ACLs, providing 1024 entries • 2 slices for MAC ACLs, providing 1024 entries • 1 slice for ingress subinterface policing, providing 2048 entries
EFP	There are 4 EFP slices providing a total of 4*512 = 2048 entries
	1 slice is allocated statically. Entries serve 2 purposes: • ES pruning of local-biased traffic. This requires 1 entry per system. • Egress port mirroring statistics. This requires 1 entry per outgoing interface.
	3 slices are available for dynamic allocation • 2 physical slices are grouped to support double-wide entries • 3 physical slices are grouped to support triple-wide entries • 1 triple-wide entry is consumed by each interface egress IPv4 filter entry, interface egress IPv6 filter entry, IPv4 CPM filter entry or IPv6 CPM filter entry (the system supports mapping IPv4 groups and IPv6 groups onto the same physical slices); maximum possible scale is 512 entries • 1 double-wide entry is consumed by each interface egress MAC filter entry or MAC CPM filter entry

5.4 TCAM allocation on 7250 IXR-6/10 and 7250 IXR-6e/10e

Each forwarding complex on a 7250 IXR-6/10 and IXR-6e/10e IMM has a TCAM with 12 large banks and 4 small banks. Each of the large banks supports 2K entries each, addressable with a 160-bit key. Each of the small banks supports 256 entries each with a 160-bit key size. This TCAM bank allocation is shown in Figure 1: TCAM allocation on 7250 IXR-6/10 and IXR-6e/10e.

Figure 1: TCAM allocation on 7250 IXR-6/10 and IXR-6e/10e

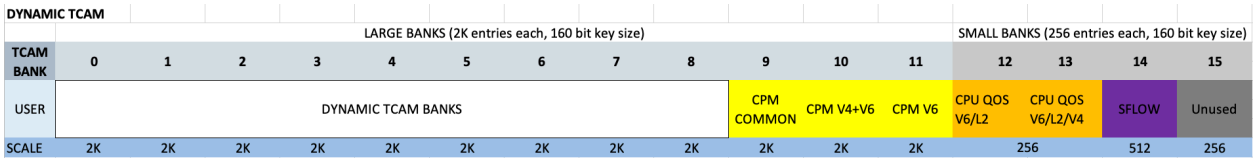

ACLs can be dynamically allocated to banks 0-8. Requirements for each ACL type are as follows:

- Ingress IPv4 ACL: entries require single-wide banks that can start at any bank number (and do not need to be contiguous)
- Egress IPv4 ACL: entries require single-wide banks that can start at any bank number (and do not need to be contiguous)

- Ingress IPv6 ACL: entries require double-wide (side-by-side) banks that must start at an even bank number
- Egress IPv6 ACL: entries require double-wide (side-by-side) banks that must start at an even bank number
- IPv4 policy-forwarding (PBF): entries require single-wide banks that can start at any bank number (and do not need to be contiguous)

Dynamic TCAM allocation works as follows:

- When a new bank needs to be allocated, the system looks for the first available space, progressing in ascending order from bank 0. If space for a double-wide bank cannot be found, the system attempts to make space by moving the fewest number of single-wide banks.
- When the number of entries required by a particular user drops to a level where a single-wide or double-wide bank can be freed up, the system selects the bank that can create the largest space of free banks, moving entries between banks as necessary.

Customer document and product support

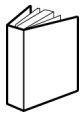

Customer documentation

Customer documentation welcome page

Technical support

Product support portal

Documentation feedback

Customer documentation feedback