

CoS for Tunnels on EX9200 Switches



Published: 2015-05-15

Juniper Networks, Inc.
1133 Innovation Way
Sunnyvale, California 94089
USA
408-745-2000
www.juniper.net

Juniper Networks, Junos, Steel-Belted Radius, NetScreen, and ScreenOS are registered trademarks of Juniper Networks, Inc. in the United States and other countries. The Juniper Networks Logo, the Junos logo, and JunosE are trademarks of Juniper Networks, Inc. All other trademarks, service marks, registered trademarks, or registered service marks are the property of their respective owners.

Juniper Networks assumes no responsibility for any inaccuracies in this document. Juniper Networks reserves the right to change, modify, transfer, or otherwise revise this publication without notice.

CoS for Tunnels on EX9200 Switches
Copyright © 2015, Juniper Networks, Inc.
All rights reserved.

The information in this document is current as of the date on the title page.

YEAR 2000 NOTICE

Juniper Networks hardware and software products are Year 2000 compliant. Junos OS has no known time-related limitations through the year 2038. However, the NTP application is known to have some difficulty in the year 2036.

END USER LICENSE AGREEMENT

The Juniper Networks product that is the subject of this technical documentation consists of (or is intended for use with) Juniper Networks software. Use of such software is subject to the terms and conditions of the End User License Agreement ("EULA") posted at <http://www.juniper.net/support/eula.html>. By downloading, installing or using such software, you agree to the terms and conditions of that EULA.

Table of Contents

	About the Documentation	ix
	Documentation and Release Notes	ix
	Supported Platforms	ix
	Using the Examples in This Manual	ix
	Merging a Full Example	x
	Merging a Snippet	x
	Documentation Conventions	xi
	Documentation Feedback	xiii
	Requesting Technical Support	xiii
	Self-Help Online Tools and Resources	xiii
	Opening a Case with JTAC	xiv
Part 1	Overview	
Chapter 1	CoS for Tunnels	3
	CoS for Tunnels Overview	3
Part 2	Configuration	
Chapter 2	Configuration Task	7
	Configuring CoS for Tunnels	7
Chapter 3	Examples	9
	Example: Configuring CoS for Tunnels	9
	Example: Configuring a GRE Tunnel to Copy ToS Bits to the Outer IP Header	12
Chapter 4	Configuration Statements	13
	[edit class-of-service] Hierarchy Level	13
	code-point	17
	dscp (Rewrite Rules)	18
	dscp-ipv6 (CoS Rewrite Rules)	19
	exp	20
	exp-push-push-push	21
	exp-swap-push-push	22
	forwarding-class (BA Classifiers)	23
	ieee-802.1 (Rewrite Rules on Logical Interface)	24
	import (Rewrite Rules)	25
	inet-precedence (CoS Rewrite Rules)	25
	interfaces (CoS)	26
	loss-priority (BA Classifiers)	28
	protocol (Rewrite Rules)	29
	rewrite-rules (Definition)	30

rewrite-rules (Interfaces)	31
unit	33
[edit interfaces] Hierarchy Level	33
copy-tos-to-outer-ip-header	45

List of Figures

Part 2	Configuration	
Chapter 3	Examples	9
	Figure 1: CoS with a Tunnel Configuration	9

List of Tables

About the Documentation	ix
Table 1: Notice Icons	xi
Table 2: Text and Syntax Conventions	xi

About the Documentation

- Documentation and Release Notes on page ix
- Supported Platforms on page ix
- Using the Examples in This Manual on page ix
- Documentation Conventions on page xi
- Documentation Feedback on page xiii
- Requesting Technical Support on page xiii

Documentation and Release Notes

To obtain the most current version of all Juniper Networks® technical documentation, see the product documentation page on the Juniper Networks website at <http://www.juniper.net/techpubs/>.

If the information in the latest release notes differs from the information in the documentation, follow the product Release Notes.

Juniper Networks Books publishes books by Juniper Networks engineers and subject matter experts. These books go beyond the technical documentation to explore the nuances of network architecture, deployment, and administration. The current list can be viewed at <http://www.juniper.net/books>.

Supported Platforms

For the features described in this document, the following platforms are supported:

- EX Series

Using the Examples in This Manual

If you want to use the examples in this manual, you can use the **load merge** or the **load merge relative** command. These commands cause the software to merge the incoming configuration into the current candidate configuration. The example does not become active until you commit the candidate configuration.

If the example configuration contains the top level of the hierarchy (or multiple hierarchies), the example is a *full example*. In this case, use the **load merge** command.

If the example configuration does not start at the top level of the hierarchy, the example is a *snippet*. In this case, use the **load merge relative** command. These procedures are described in the following sections.

Merging a Full Example

To merge a full example, follow these steps:

1. From the HTML or PDF version of the manual, copy a configuration example into a text file, save the file with a name, and copy the file to a directory on your routing platform.

For example, copy the following configuration to a file and name the file **ex-script.conf**. Copy the **ex-script.conf** file to the **/var/tmp** directory on your routing platform.

```
system {
  scripts {
    commit {
      file ex-script.xml;
    }
  }
}
interfaces {
  fxp0 {
    disable;
    unit 0 {
      family inet {
        address 10.0.0.1/24;
      }
    }
  }
}
```

2. Merge the contents of the file into your routing platform configuration by issuing the **load merge** configuration mode command:

```
[edit]
user@host# load merge /var/tmp/ex-script.conf
load complete
```

Merging a Snippet

To merge a snippet, follow these steps:

1. From the HTML or PDF version of the manual, copy a configuration snippet into a text file, save the file with a name, and copy the file to a directory on your routing platform.

For example, copy the following snippet to a file and name the file **ex-script-snippet.conf**. Copy the **ex-script-snippet.conf** file to the **/var/tmp** directory on your routing platform.

```
commit {
  file ex-script-snippet.xml; }
```

2. Move to the hierarchy level that is relevant for this snippet by issuing the following configuration mode command:

```
[edit]
user@host# edit system scripts
[edit system scripts]
```

3. Merge the contents of the file into your routing platform configuration by issuing the **load merge relative** configuration mode command:

```
[edit system scripts]
user@host# load merge relative /var/tmp/ex-script-snippet.conf
load complete
```

For more information about the **load** command, see the *CLI User Guide*.

Documentation Conventions

Table 1 on page xi defines notice icons used in this guide.

Table 1: Notice Icons

Icon	Meaning	Description
	Informational note	Indicates important features or instructions.
	Caution	Indicates a situation that might result in loss of data or hardware damage.
	Warning	Alerts you to the risk of personal injury or death.
	Laser warning	Alerts you to the risk of personal injury from a laser.
	Tip	Indicates helpful information.
	Best practice	Alerts you to a recommended use or implementation.

Table 2 on page xi defines the text and syntax conventions used in this guide.

Table 2: Text and Syntax Conventions

Convention	Description	Examples
Bold text like this	Represents text that you type.	To enter configuration mode, type the configure command: user@host> configure

Table 2: Text and Syntax Conventions (*continued*)

Convention	Description	Examples
Fixed-width text like this	Represents output that appears on the terminal screen.	<code>user@host> show chassis alarms</code> <code>No alarms currently active</code>
<i>Italic text like this</i>	<ul style="list-style-type: none">Introduces or emphasizes important new terms.Identifies guide names.Identifies RFC and Internet draft titles.	<ul style="list-style-type: none">A policy <i>term</i> is a named structure that defines match conditions and actions.<i>Junos OS CLI User Guide</i>RFC 1997, <i>BGP Communities Attribute</i>
<i>Italic text like this</i>	Represents variables (options for which you substitute a value) in commands or configuration statements.	Configure the machine's domain name: [edit] root@# set system domain-name <i>domain-name</i>
Text like this	Represents names of configuration statements, commands, files, and directories; configuration hierarchy levels; or labels on routing platform components.	<ul style="list-style-type: none">To configure a stub area, include the stub statement at the [edit protocols ospf area area-id] hierarchy level.The console port is labeled CONSOLE.
< > (angle brackets)	Encloses optional keywords or variables.	stub <default-metric <i>metric</i>>;
(pipe symbol)	Indicates a choice between the mutually exclusive keywords or variables on either side of the symbol. The set of choices is often enclosed in parentheses for clarity.	broadcast multicast (<i>string1</i> <i>string2</i> <i>string3</i>)
# (pound sign)	Indicates a comment specified on the same line as the configuration statement to which it applies.	rsvp { # Required for dynamic MPLS only
[] (square brackets)	Encloses a variable for which you can substitute one or more values.	community name members [<i>community-ids</i>]
Indentation and braces ({ })	Identifies a level in the configuration hierarchy.	[edit] routing-options { static { route default { nexthop <i>address</i> ; retain; } } }
;(semicolon)	Identifies a leaf statement at a configuration hierarchy level.	}
GUI Conventions		
Bold text like this	Represents graphical user interface (GUI) items you click or select.	<ul style="list-style-type: none">In the Logical Interfaces box, select All Interfaces.To cancel the configuration, click Cancel.

Table 2: Text and Syntax Conventions (*continued*)

Convention	Description	Examples
> (bold right angle bracket)	Separates levels in a hierarchy of menu selections.	In the configuration editor hierarchy, select Protocols>Ospf .

Documentation Feedback

We encourage you to provide feedback, comments, and suggestions so that we can improve the documentation. You can provide feedback by using either of the following methods:

- Online feedback rating system—On any page at the Juniper Networks Technical Documentation site at <http://www.juniper.net/techpubs/index.html>, simply click the stars to rate the content, and use the pop-up form to provide us with information about your experience. Alternately, you can use the online feedback form at <https://www.juniper.net/cgi-bin/docbugreport/>.
- E-mail—Send your comments to techpubs-comments@juniper.net. Include the document or topic name, URL or page number, and software version (if applicable).

Requesting Technical Support

Technical product support is available through the Juniper Networks Technical Assistance Center (JTAC). If you are a customer with an active J-Care or JNASC support contract, or are covered under warranty, and need post-sales technical support, you can access our tools and resources online or open a case with JTAC.

- JTAC policies—For a complete understanding of our JTAC procedures and policies, review the *JTAC User Guide* located at <http://www.juniper.net/us/en/local/pdf/resource-guides/7100059-en.pdf>.
- Product warranties—For product warranty information, visit <http://www.juniper.net/support/warranty/>.
- JTAC hours of operation—The JTAC centers have resources available 24 hours a day, 7 days a week, 365 days a year.

Self-Help Online Tools and Resources

For quick and easy problem resolution, Juniper Networks has designed an online self-service portal called the Customer Support Center (CSC) that provides you with the following features:

- Find CSC offerings: <http://www.juniper.net/customers/support/>
- Search for known bugs: <http://www2.juniper.net/kb/>
- Find product documentation: <http://www.juniper.net/techpubs/>
- Find solutions and answer questions using our Knowledge Base: <http://kb.juniper.net/>

- Download the latest versions of software and review release notes:
<http://www.juniper.net/customers/csc/software/>
- Search technical bulletins for relevant hardware and software notifications:
<http://kb.juniper.net/InfoCenter/>
- Join and participate in the Juniper Networks Community Forum:
<http://www.juniper.net/company/communities/>
- Open a case online in the CSC Case Management tool: <http://www.juniper.net/cm/>

To verify service entitlement by product serial number, use our Serial Number Entitlement (SNE) Tool: <https://tools.juniper.net/SerialNumberEntitlementSearch/>

Opening a Case with JTAC

You can open a case with JTAC on the Web or by telephone.

- Use the Case Management tool in the CSC at <http://www.juniper.net/cm/>.
- Call 1-888-314-JTAC (1-888-314-5822 toll-free in the USA, Canada, and Mexico).

For international or direct-dial options in countries without toll-free numbers, see <http://www.juniper.net/support/requesting-support.html>.

PART 1

Overview

- [CoS for Tunnels on page 3](#)

CHAPTER 1

CoS for Tunnels

- [CoS for Tunnels Overview on page 3](#)

CoS for Tunnels Overview

For Adaptive Services, Link Services, and Tunnel PICs installed on Juniper Networks M Series Multiservice Edge Routers and T Series Core Routers with enhanced Flexible PIC Concentrators (FPCs), class-of-service (CoS) information is preserved inside generic routing encapsulation (GRE) and IP-IP tunnels.

For the ES PIC installed on M Series and T Series routers with enhanced FPCs, class-of-service information is preserved inside IP Security (IPsec) tunnels. For IPsec tunnels, you do not need to configure CoS, because the ES PIC copies the type-of-service (ToS) byte from the inner IP header to the GRE or IP-IP header.

For IPsec tunnels, the IP header type-of-service (ToS) bits are copied to the outer IPsec header at encryption side of the tunnel. You can rewrite the outer ToS bits in the IPsec header using a rewrite rule. On the decryption side of the IPsec tunnel, the ToS bits in the IPsec header are not written back to the original IP header field. You can still apply a firewall filter to the ToS bits to apply a packet action on egress. For more information about ToS bits and the Multiservices PICs, see *Multiservices PIC ToS Translation*. For more information about IPsec and Multiservices PICs, see the *Junos OS Services Interfaces Library for Routing Devices*.

To configure CoS for tunnels, include the following statements at the **[edit class-of-service]** and **[edit interfaces]** hierarchy level:

```
[edit class-of-service]
interfaces {
  interface-name {
    unit logical-unit-number {
      rewrite-rules {
        dscp (rewrite-name | default);
        dscp-ipv6 (rewrite-name | default);
        exp (rewrite-name | default) protocol protocol-types;
        exp-push-push-push default;
        exp-swap-push-push default;
        ieee-802.1 (rewrite-name | default);
        inet-precedence (rewrite-name | default);
      }
    }
  }
}
```

```
    }  
  }  
  rewrite-rules {  
    (dscp | dscp-ipv6 | exp | ieee-802.1 | inet-precedence) rewrite-name {  
      import (rewrite-name | default);  
      forwarding-class class-name {  
        loss-priority level code-point (alias | bits);  
      }  
    }  
  }  
}  
[edit interfaces]  
gre-interface-name {  
  unit logical-unit-number;  
  copy-tos-to-outer-ip-header;  
}
```

PART 2

Configuration

- [Configuration Task on page 7](#)
- [Examples on page 9](#)
- [Configuration Statements on page 13](#)

CHAPTER 2

Configuration Task

- [Configuring CoS for Tunnels on page 7](#)

Configuring CoS for Tunnels

To configure CoS for GRE and IP-IP tunnels, perform the following configuration tasks:

1. To configure the tunnel, include the **tunnel** statement at the **[edit interfaces *ip-fpc/pic/port* unit *logical-unit-number*]** or **[edit interfaces *gr-fpc/pic/port* unit *logical-unit-number*]** hierarchy level.
2. To rewrite traffic on the outbound interface, include the **rewrite-rules** statement at the **[edit class-of-service]** and **[edit class-of-service interfaces *interface-name* unit *logical-unit-number*]** hierarchy levels. For GRE and IP-IP tunnels, you can configure IP precedence and DSCP rewrite rules.
3. To classify traffic on the inbound interface, you can configure a behavior aggregate (BA) classifier or firewall filter. Include the **loss-priority** and **forwarding-class** statements at the **[edit firewall filter *filter-name* term *term-name* then]** hierarchy level, or the **classifiers** statement at the **[edit class-of-service]** hierarchy level.
4. For a GRE tunnel, the default is to set the ToS bits in the outer IP header to all 0s. To copy the ToS bits from the inner IP header to the outer, include the **copy-tos-to-outer-ip-header** statement at the **[edit interfaces *gr-fpc/pic/port* unit *logical-unit-number*]** hierarchy level. (This inner-to-outer ToS bits copying is already the default behavior for IP-IP tunnels.)

CHAPTER 3


Examples

- [Example: Configuring CoS for Tunnels on page 9](#)
- [Example: Configuring a GRE Tunnel to Copy ToS Bits to the Outer IP Header on page 12](#)

Example: Configuring CoS for Tunnels

In [Figure 1 on page 9](#), Router A acts as a tunnel ingress device. The link between interfaces **ge-1/0/0** in Router A and **ge-1/3/0** in Router B is the GRE or IP-IP tunnel. Router A monitors the traffic received from interface **ge-1/3/0**. By way of interface **ge-1/0/0**, Router C generates traffic to Router B.

Figure 1: CoS with a Tunnel Configuration



```
Router A [edit interfaces]
ge-1/0/0 {
  unit 0 {
    family inet {
      address 10.80.0.2/24;
    }
  }
}
ge-1/0/1 {
  unit 0 {
    family inet {
      filter {
        input zf-catch-all;
      }
      address 10.90.0.2/24;
    }
  }
}
gr-2/1/0 {
```

```
unit 0 {
  tunnel {
    source 11.11.11.11;
    destination 10.255.245.46;
  }
  family inet {
    address 21.21.21.21/24;
  }
}
}
ip-2/1/0 {
  unit 0 {
    tunnel {
      source 12.12.12.12;
      destination 10.255.245.46;
    }
    family inet {
      address 22.22.22.22/24;
    }
  }
}

[edit routing-options]
static {
  route 1.1.1.1/32 next-hop gr-2/1/0.0;
  route 2.2.2.2/32 next-hop ip-2/1/0.0;
}

[edit class-of-service]
interfaces {
  ge-1/0/0 {
    unit 0 {
      rewrite-rules {
        inet-precedence zf-tun-rw-ipprec-00;
      }
    }
  }
}
rewrite-rules {
  inet-precedence zf-tun-rw-ipprec-00 {
    forwarding-class best-effort {
      loss-priority low code-point 000;
      loss-priority high code-point 001;
    }
    forwarding-class expedited-forwarding {
      loss-priority low code-point 010;
      loss-priority high code-point 011;
    }
    forwarding-class assured-forwarding {
      loss-priority low code-point 100;
      loss-priority high code-point 101;
    }
    forwarding-class network-control {
      loss-priority low code-point 110;
      loss-priority high code-point 111;
    }
  }
}
```



```

    }
  }
  dscp zf-tun-rw-dscp-00 {
    forwarding-class best-effort {
      loss-priority low code-point 000000;
      loss-priority high code-point 001001;
    }
    forwarding-class expedited-forwarding {
      loss-priority low code-point 010010;
      loss-priority high code-point 011011;
    }
    forwarding-class assured-forwarding {
      loss-priority low code-point 100100;
      loss-priority high code-point 101101;
    }
    forwarding-class network-control {
      loss-priority low code-point 110110;
      loss-priority high code-point 111111;
    }
  }
}

[edit firewall]
filter zf-catch-all {
  term term1 {
    then {
      loss-priority high;
      forwarding-class network-control;
    }
  }
}

```

```

Router B [edit interfaces]
ge-1/3/0 {
  unit 0 {
    family inet {
      address 10.80.0.1/24;
    }
  }
}
lo0 {
  unit 0 {
    family inet {
      address 10.255.245.46/32;
    }
  }
}

```

```

Router C [edit interfaces]
ge-1/0/0 {
  unit 0 {
    family inet {
      address 10.90.0.1/24;
    }
  }
}

```

```
[edit routing-options]
static {
  route 1.1.1.1/32 next-hop 10.90.0.2;
  route 2.2.2.2/32 next-hop 10.90.0.2;
}
```

Example: Configuring a GRE Tunnel to Copy ToS Bits to the Outer IP Header

Unlike IP-IP tunnels, GRE tunnels do not copy the ToS bits to the outer IP header by default. To copy the inner ToS bits to the outer IP header (which is required for some tunneled routing protocols) on packets sent by the Routing Engine, include the **copy-tos-to-outer-ip-header** statement at the logical unit hierarchy level of a GRE interface. This example copies the inner ToS bits to the outer IP header on a GRE tunnel:

```
[edit interfaces]
gr-0/0/0 {
  unit 0 {
    copy-tos-to-outer-ip-header;
    family inet;
  }
}
```

CHAPTER 4

Configuration Statements

- [\[edit class-of-service\] Hierarchy Level](#) on page 13
- [code-point](#) on page 17
- [dscp \(Rewrite Rules\)](#) on page 18
- [dscp-ipv6 \(CoS Rewrite Rules\)](#) on page 19
- [exp](#) on page 20
- [exp-push-push-push](#) on page 21
- [exp-swap-push-push](#) on page 22
- [forwarding-class \(BA Classifiers\)](#) on page 23
- [ieee-802.1 \(Rewrite Rules on Logical Interface\)](#) on page 24
- [import \(Rewrite Rules\)](#) on page 25
- [inet-precedence \(CoS Rewrite Rules\)](#) on page 25
- [interfaces \(CoS\)](#) on page 26
- [loss-priority \(BA Classifiers\)](#) on page 28
- [protocol \(Rewrite Rules\)](#) on page 29
- [rewrite-rules \(Definition\)](#) on page 30
- [rewrite-rules \(Interfaces\)](#) on page 31
- [unit](#) on page 33
- [\[edit interfaces\] Hierarchy Level](#) on page 33
- [copy-tos-to-outer-ip-header](#) on page 45

[\[edit class-of-service\] Hierarchy Level](#)

```
class-of-service {
  classifiers {
    type classifier-name {
      forwarding-class class-name {
        loss-priority (high | low | medium-high | medium-low) code-points [ aliases bits ];
      }
      import (classifier-name | default);
    }
  }
  code-point-aliases {
```

```

(dscp | dscp-ipv6 | exp | ieee-802.1 | ieee-802.1ad | inet-precedence) {
    alias-name bits;
}
}
drop-profiles {
    profile-name {
        fill-level percentage drop-probability percentage;
        interpolate {
            drop-probability value;
            fill-level value;
        }
    }
}
fabric {
    scheduler-map {
        priority (high | low) scheduler scheduler-name;
    }
}
forwarding-class-map {
    map-name {
        class class-name queue-num queue-number <restricted-queue queue-number>;
    }
}
forwarding-classes {
    class class-name policing-priority (normal | premium) queue-num queue-number
        priority (high | low);
    queue queue-number class-name policing-priority (normal | premium) priority (high |
        low);
}
forwarding-policy {
    class class-name {
        classification-override {
            forwarding-class class-name;
        }
    }
    next-hop-map map-name {
        forwarding-class class-name {
            discard;
            lsp-next-hop [ lsp-regular-expressions ];
            next-hop [ next-hop-names ];
            non-lsp-next-hop;
        }
    }
}
fragmentation-maps {
    map-name {
        forwarding-class class-name {
            drop-timeout milliseconds;
            fragment-threshold bytes;
            multilink-class number;
            no-fragmentation;
        }
    }
}
host-outbound-traffic {
    dscp-code-point value;
}

```

```

forwarding-class class-name;
ieee-802.1 {
    default value;
    rewrite-rules;
}
tcp {
    raise-internet-control-priority;
}
}
interfaces {
    ... the interfaces subhierarchy appears after the main [edit class-of-service] hierarchy
    ...
}
}
restricted-queues {
    forwarding-class class-name queue-number;
}
rewrite-rules {
    (dscp | dscp-ipv6 | exp | frame-relay-de | ieee-802.1 | ieee-802.1ad | inet-precedence)
    rewrite-rule {
        forwarding-class class-name {
            loss-priority level code-point (alias | bits);
        }
        import (rewrite-rule | default);
    }
}
}
routing-instances routing-instance-name {
    classifiers {
        dscp (classifier-name | default);
        dscp-ipv6 (classifier-name | default);
        exp (classifier-name | default);
        ieee-208.1 (classifier-name | default | encapsulated | vlan-tag (inner | outer));
    }
}
scheduler-maps {
    map-name {
        forwarding-class class-name scheduler scheduler-name;
    }
}
schedulers {
    scheduler-name {
        adjust-minimum value;
        adjust-percent value;
        buffer-size (exact | percent percentage | remainder);
        drop-profile-map loss-priority (any | high | low | medium-high | medium-low)
            protocol any;
        excess-priority (high | low | medium-high | medium-low);
        excess-rate (percent percentage | proportion proportion);
        priority (high | low | medium-high | medium-low | strict-high);
        shaping-rate (bps | percent percentage | burst-size size);
        transmit-rate (bps | percent percentage | remainder) <exact | rate-limit>;
    }
}
}
traceoptions {
    file <files number> <match regular-expression> <size maximum-file-size>
        <world-readable | no-world-readable>;
}

```

```

    flag flag;
    no-remote-trace;
}
traffic-control-profiles {
    profile-name {
        adjust-minimum rate;
        delay-buffer-rate (bps | cps cps | percent percentage);
        excess-rate (percent percentage | proportion value);
        guaranteed-rate (bps | percent percentage) <burst-size bytes>;
        overhead-accounting (frame-mode | cell-mode) <bytes byte-value>;
        scheduler-map map-name;
        shaping-rate (bps | percent percentage) <burst-size bytes>;
    }
}
tri-color;
}

class-of-service {
    interfaces {
        interface-name {
            excess-bandwidth-share (equal | proportional value);
            input-excess-bandwidth-share (equal | proportional value);
            input-scheduler-map map-name;
            input-shaping-rate bps;
            input-traffic-control-profile profile-name;
            output-forwarding-class-map map-name;
            output-traffic-control-profile profile-name;
            scheduler-map map-name;
            scheduler-map-chassis (map-name | derived);
            shaping-rate bps;
            unit (logical-unit-number | *) {
                classifiers {
                    dscp (classifier-name | default) {
                        family [ inet mpls ];
                    }
                    dscp-ipv6 (classifier-name | default) {
                        family [ inet mpls ];
                    }
                    exp (classifier-name | default);
                    ieee-208.1 (classifier-name | default) <vlan-tag (inner | outer)>;
                    ieee-208.1ad (classifier-name | default);
                    inet-precedence (classifier-name | default);
                }
                forwarding-class class-name;
                input-scheduler-map map-name;
                input-shaping-rate bps;
                input-traffic-control-profile profile-name shared-instance instance-name;
                loss-priority-maps {
                    (map-name | default);
                }
                loss-priority-rewrites {
                    (map-name | default);
                }
                output-forwarding-class-map map-name;
                output-traffic-control-profile profile-name shared-instance instance-name;
                rewrite-rules {

```

```

dscp (rule-name | default) <protocol mpls>;
dscp-ipv6 (rule-name | default);
exp (rule-name | default) <protocol [ mpls-any | mpls-inet-both |
    mpls-inet-both-non-vpn ]>;
exp-push-push-push default;
exp-swap-push-push default;
ieee-802.1 (rewrite-name | default) <vlan-tag (outer | outer-and-inner)>;
ieee-802.1ad (rewrite-name | default) <vlan-tag (outer | outer-and-inner)>;
inet-precedence (rewrite-name | default) <protocol mpls>;
}
scheduler-map map-name;
shaping-rate bps;
translation-table (to-dscp-from-dscp | to-dscp-ipv6-from-dscp-ipv6 |
    to-exp-from-exp | to-inet-precedence-from-inet-precedence) table-name;
}
}
interface-set interface-set-name {
    excess-bandwidth-share (equal | proportional value);
    input-excess-bandwidth-share (equal | proportional value);
    input-traffic-control-profile profile-name;
    input-traffic-control-profile-remaining profile-name;
    internal-node;
    output-traffic-control-profile profile-name;
    output-traffic-control-profile-remaining profile-name;
}
}
}

```

**Related
Documentation**

- *Notational Conventions Used in Junos OS Configuration Hierarchies*

code-point

Syntax	<code>code-point [<i>aliases</i>] [<i>bit-patterns</i>];</code>
Hierarchy Level	<code>[edit class-of-service rewrite-rules <i>type</i> <i>rewrite-name</i> forwarding-class <i>class-name</i>]</code>
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Specify one or more code-point aliases or bit sets for association with a forwarding class.
Options	<p><i>aliases</i>—Name of each alias.</p> <p><i>bit-patterns</i>—Value of the code-point bits, in decimal form.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring Rewrite Rules</i>

dscp (Rewrite Rules)

Syntax	<code>dscp (rewrite-name default) protocol mpls;</code>
Hierarchy Level	[edit class-of-service interfaces <i>interface-name</i> unit <i>logical-unit-number</i> rewrite-rules]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	<p>For IPv4 traffic, apply a Differentiated Services (DiffServ) code point (DSCP) rewrite rule.</p> <p>Logical interfaces do not support multiple dscp rewrite rules for the same protocol.</p> <p>DSCP and DSCP IPv6 rewrite rules are supported on M Series and T Series routers when non-queuing PICs are installed, but are disabled when queuing PICs are installed with the following exceptions:</p> <ul style="list-style-type: none"> On M320 routers, DSCP rewrite is supported on IQ, IQ2, IQE, and IQ2E PICs when used with the Enhanced III FPC. On M120 routers, DSCP rewrite is supported on IQ, IQ2, IQE, and IQ2E PICs. <p>DSCP and DCSP IPv6 rewrite rules are supported on MIC and MPC interfaces on MX Series routers.</p> <p>DSCP rewrite rules are not supported on T Series routers when IQ, IQ2, IQE, IQ2E, SONET/SDH OC48/STM16 IQE, or PD-5-10XGE-SFPP PICs are installed.</p>
Options	<p>rewrite-name—Name of a rewrite-rules mapping configured at the [edit class-of-service rewrite-rules dscp] hierarchy level.</p> <p>default—The default mapping.</p> <p>protocol mpls—(Optional for ingress MPLS tunnel nodes) For interfaces on MX Series routers or hosted on Enhanced III FPCs in M120 or M320 routers only, rewrite the MPLS EXP bits in the MPLS header independently of the IPv4 DSCP value for IPv4 packets entering an MPLS tunnel.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> <i>Configuring Rewrite Rules</i> <i>Applying Rewrite Rules to Output Logical Interfaces</i> protocol (Rewrite Rules) on page 29 <i>Rewriting MPLS and IPv4 Packet Headers</i> rewrite-rules (Definition) on page 30

dscp-ipv6 (CoS Rewrite Rules)

Syntax	<code>dscp-ipv6 (<i>rewrite-name</i> <default>) protocol mpls;</code>
Hierarchy Level	[edit class-of-service interfaces <i>interface-name</i> unit <i>logical-unit-number</i> rewrite-rules]
Release Information	Statement introduced before Junos OS Release 7.4. Support for protocol mpls option introduced in Junos OS Release 10.4R2.
Description	<p>For IPv6 traffic, apply a DSCP rewrite rule.</p> <p>Logical interfaces do not support multiple dscp-ipv6 rewrite rules for the same protocol.</p> <p>DSCP and DSCP IPv6 rewrite rules are supported on M Series and T Series routers when non-queuing PICs are installed, but are disabled when queuing PICs are installed with the following exceptions:</p> <ul style="list-style-type: none"> On M320 routers, DSCP rewrite is supported on IQ, IQ2, IQE, and IQ2E PICs when used with the Enhanced III FPC. On M120 routers, DSCP rewrite is supported on IQ, IQ2, IQE, and IQ2E PICs. <p>DSCP and DCSP IPv6 rewrite rules are supported on MIC and MPC interfaces on MX Series routers.</p> <p>DSCP rewrite rules are not supported on T Series routers when IQ, IQ2, IQE, IQ2E, SONET/SDH OC48/STM16 IQE, or PD-5-10XGE-SFPP PICs are installed.</p>
Options	<p>rewrite-name—Name of a rewrite-rules mapping configured at the [edit class-of-service rewrite-rules dscp-ipv6] hierarchy level.</p> <p>default—Default mapping.</p> <p>protocol mpls—(Optional for ingress MPLS tunnel nodes) For interfaces on MX Series routers or hosted on Enhanced III FPCs in M120 or M320 routers only, rewrite the MPLS EXP bits in the MPLS header independently of the IPv6 DSCP value for IPv6 packets entering an MPLS tunnel.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> Configuring Rewrite Rules protocol on page 29 Setting IPv6 DSCP and MPLS EXP Values Independently Configuring DSCP Values for IPv6 Packets Entering the MPLS Tunnel Applying Rewrite Rules to Output Logical Interfaces rewrite-rules (Definition) on page 30

exp

Syntax	<code>exp (rewrite-name default) protocol protocol-types;</code>
Hierarchy Level	<code>[edit class-of-service interfaces interface-name unit logical-unit-number rewrite-rules]</code>
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced before Junos OS Release 12.2. for ACX series
Description	Apply an MPLS experimental (EXP) rewrite rule.
Options	<p>rewrite-name—Name of a rewrite-rules mapping configured at the <code>[edit class-of-service rewrite-rules exp]</code> hierarchy level.</p> <p>default—The default mapping.</p> <p>By default, IP precedence rewrite rules alter the first three bits on the type-of-service (ToS) byte while leaving the last three bits unchanged. This default behavior applies to rewrite rules you configure for MPLS packets with IPv4 payloads. You configure these types of rewrite rules by including the mpls-inet-both or mpls-inet-both-non-vpn option at the <code>[edit class-of-service interfaces interface interface-name unit logical-unit-number rewrite-rules exp rewrite-rule-name protocol]</code> hierarchy level. The IP precedence rewrite rules explanation does not apply to ACX Series Universal Access routers.</p> <p>On interfaces configured on Modular Port Concentrators (MPCs) and Modular Interface Cards (MICs) on MX Series 3D Universal Edge Routers and EX Series switches, we highly recommend that you configure the default option when you configure a behavior aggregate (BA) classifier that does not include a specific rewrite rule for MPLS packets. Doing so ensures that MPLS exp value is rewritten according to the BA classifier rules configured for forwarding or packet loss priority. This does not apply to ACX Series Universal Access routers.</p> <p>protocol-types—Specify one or more protocol matching criteria:</p> <ul style="list-style-type: none"> • mpls-any—Apply to MPLS packets, write MPLS header only. • mpls-inet-both—Apply to IPv4 MPLS packets, write MPLS and IPv4 header. • mpls-inet-both-non-vpn—Apply to IPv4 MPLS packets, write MPLS and IPv4 header for only non VPN traffic.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring Rewrite Rules</i> • <i>Rewriting the EXP Bits of All Three Labels of an Outgoing Packet</i> • <i>Applying Rewrite Rules to Output Logical Interfaces</i> • protocol (Rewrite Rules) on page 29

- [rewrite-rules \(Definition\) on page 30](#)

exp-push-push-push

Syntax	exp-push-push-push default;
Hierarchy Level	[edit class-of-service interfaces <i>interface-name</i> unit <i>logical-unit-number</i> rewrite-rules]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	For M Series routers, rewrite the EXP bits of all three labels of an outgoing packet, thereby maintaining CoS of an incoming non-MPLS packet.
Options	default —Apply the default MPLS EXP rewrite table.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Rewriting the EXP Bits of All Three Labels of an Outgoing Packet</i> • dscp (Rewrite Rules) on page 18 • dscp-ipv6 (CoS Rewrite Rules) on page 19 • exp on page 20 • exp-swap-push-push on page 22 • ieee-802.1 (Rewrite Rules on Logical Interface) on page 24 • <i>ieee-802.1ad</i> • inet-precedence (CoS Rewrite Rules) on page 25 • rewrite-rules (Definition) on page 30

exp-swap-push-push

Syntax	exp-swap-push-push default;
Hierarchy Level	[edit class-of-service interfaces <i>interface-name</i> unit <i>logical-unit-number</i> rewrite-rules]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	For M Series routers, rewrite the EXP bits of all three labels of an outgoing packet, thereby maintaining CoS of an incoming MPLS packet.
Options	default —Apply the default MPLS EXP rewrite table.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Rewriting the EXP Bits of All Three Labels of an Outgoing Packet</i>• dscp (Rewrite Rules) on page 18• dscp-ipv6 (CoS Rewrite Rules) on page 19• exp on page 20• exp-push-push-push on page 21• ieee-802.1 (Rewrite Rules on Logical Interface) on page 24• <i>ieee-802.1ad</i>• inet-precedence (CoS Rewrite Rules) on page 25• rewrite-rules (Definition) on page 30

forwarding-class (BA Classifiers)

Syntax	<code>forwarding-class <i>class-name</i> { <i>loss-priority level</i> code-points [<i>aliases</i>] [<i>bit-patterns</i>]; }</code>
Hierarchy Level	[edit class-of-service classifiers <i>type classifier-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 14.2 for PTX Series Packet Transport Routers.
Description	Define forwarding class name and option values.
Options	<i>class-name</i> —Name of the forwarding class. The remaining statements are explained separately.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Defining Classifiers</i>

ieee-802.1 (Rewrite Rules on Logical Interface)

Syntax	<code>ieee-802.1 (<i>rewrite-name</i> default) vlan-tag (outer outer-and-inner);</code>
Hierarchy Level	<code>[edit class-of-service interfaces <i>interface-name</i> unit <i>logical-unit-number</i> rewrite-rules]</code>
Release Information	Statement introduced before Junos OS Release 7.4. vlan-tag statement introduced in Junos OS Release 8.1.
Description	Apply an IEEE-802.1 rewrite rule. For IQ PICs, you can only configure one IEEE 802.1 rewrite rule on a physical port. All logical ports (units) on that physical port should apply the same IEEE 802.1 rewrite rule.
Options	<i>rewrite-name</i> —Name of a rewrite-rules mapping configured at the <code>[edit class-of-service rewrite-rules ieee-802.1]</code> hierarchy level. default —The default mapping.
Required Privilege Level	interface —To view this statement in the configuration. interface-control —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Rewrite Rules• dscp (Rewrite Rules) on page 18• dscp-ipv6 (CoS Rewrite Rules) on page 19• exp on page 20• exp-push-push-push on page 21• exp-swap-push-push on page 22• ieee-802.1ad• inet-precedence (CoS Rewrite Rules) on page 25• rewrite-rules (Definition) on page 30

import (Rewrite Rules)

Syntax	<code>import (rewrite-name default);</code>
Hierarchy Level	<code>[edit class-of-service rewrite-rules type rewrite-name]</code>
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Specify a default or previously defined rewrite-rules mapping to import.
Options	<p>rewrite-name—Name of a rewrite-rules mapping configured at the <code>[edit class-of-service rewrite-rules]</code> hierarchy level.</p> <p>default—The default rewrite-rules mapping.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring Rewrite Rules

inet-precedence (CoS Rewrite Rules)

Syntax	<code>inet-precedence (rewrite-name default);</code>
Hierarchy Level	<code>[edit class-of-service interfaces interface-name unit logical-unit-number rewrite-rules]</code>
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Apply a IPv4 precedence rewrite rule.
Options	<p>rewrite-name—Name of a rewrite-rules mapping configured at the <code>[edit class-of-service rewrite-rules inet-precedence]</code> hierarchy level.</p> <p>default—The default mapping. By default, IP precedence rewrite rules alter the first three bits on the type of service (ToS) byte while leaving the last three bits unchanged.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring Rewrite Rules • Applying Rewrite Rules to Output Logical Interfaces • protocol (Rewrite Rules) on page 29 • rewrite-rules (Definition) on page 30

interfaces (CoS)

```
Syntax interfaces {
    interface-name {
        classifiers{
            dscp(classifier-name | default) {
            }
            ieee-802.1 (classifier-name | default) vlan-tag (inner | outer | classifier-name);
            inet-precedence (rewrite-name | default);
        }
        input-scheduler-map map-name;
        input-shaping-rate rate;
        irb {
            unit logical-unit-number {
                classifiers {
                    type (classifier-name | default);
                }
                rewrite-rules {
                    dscp (rewrite-name | default);
                    dscp-ipv6 (rewrite-name | default);
                    exp (rewrite-name | default) protocol protocol-types;
                    ieee-802.1 (rewrite-name | default) vlan-tag (outer | outer-and-inner);
                    inet-precedence (rewrite-name | default);
                }
            }
        }
        member-link-scheduler (replicate | scale);
        rewrite-rules {
            dscp (rewrite-name | default);
            ieee-802.1 (rewrite-name | default) vlan-tag (outer);
            inet-precedence (rewrite-name | default);
        }
        scheduler-map map-name;
        scheduler-map-chassis map-name;
        shaping-rate rate;
        unit logical-unit-number {
            classifiers {
                type (classifier-name | default) family (mpls | inet);
            }
            forwarding-class class-name;
            fragmentation-map map-name;
            input-shaping-rate (percent percentage | rate);
            input-traffic-control-profile profile-name shared-instance instance-name;
            output-traffic-control-profile profile-name shared-instance instance-name;
            per-session-scheduler;
            rewrite-rules {
                dscp (rewrite-name | default);
                dscp-ipv6 (rewrite-name | default);
                exp (rewrite-name | default) protocol protocol-types;
                exp-push-push-push default;
                exp-swap-push-push default;
                ieee-802.1 (rewrite-name | default) vlan-tag (outer | outer-and-inner);
                inet-precedence (rewrite-name | default);
            }
        }
    }
}
```



```

    }
    scheduler-map map-name;
    shaping-rate rate;
    translation-table (to-dscp-from-dscp | to-dscp-ipv6-from-dscp-ipv6 | to-exp-from-exp
    | to-inet-precedence-from-inet-precedence) table-name;
  }
}
interface-set interface-set-name {
  excess-bandwidth-share;
  internal-node;
  output-traffic-control-profile profile-name;
  output-traffic-control-profile-remaining profile-name;
}
}

```

Hierarchy Level [edit class-of-service]

Release Information Statement introduced before Junos OS Release 7.4.
Interface-set level added in Junos OS Release 8.5.

Description Configure interface-specific CoS properties for incoming packets.

Options The remaining statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Understanding How Behavior Aggregate Classifiers Prioritize Trusted Traffic*
- *Configuring Rewrite Rules*

loss-priority (BA Classifiers)

Syntax	<code>loss-priority <i>level</i>;</code>
Hierarchy Level	[edit class-of-service classifiers <i>type classifier-name</i> forwarding-class <i>class-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 14.2 for PTX Series Packet Transport Routers.
Description	Specify packet loss priority value for a specific set of code-point aliases and bit patterns.
Options	<i>level</i> can be one of the following: <ul style="list-style-type: none">• high—Packet has high loss priority.• medium-high—Packet has medium-high loss priority.• medium-low—Packet has medium-low loss priority.• low—Packet has low loss priority.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Understanding How Behavior Aggregate Classifiers Prioritize Trusted Traffic</i>• <i>Configuring Tricolor Marking</i>

protocol (Rewrite Rules)

Syntax	<code>protocol protocol-types;</code>
Hierarchy Level	<p>[edit class-of-service interfaces <i>interface-name</i> unit <i>logical-unit-number</i> rewrite-rules exp <i>rewrite-name</i>],</p> <p>[edit class-of-service interfaces <i>interface-name</i> unit <i>logical-unit-number</i> rewrite-rules dscp <i>rewrite-name</i>],</p> <p>[edit class-of-service interfaces <i>interface-name</i> unit <i>logical-unit-number</i> rewrite-rules dscp-ipv6 <i>rewrite-name</i>],</p> <p>[edit class-of-service interfaces <i>interface-name</i> unit <i>logical-unit-number</i> rewrite-rules inet-prec <i>rewrite-name</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Option for dscp and inet-prec introduced in Junos OS Release 8.4.</p> <p>Option for dscp-ipv6 introduced in Junos OS Release 10.4R2.</p>
Description	Apply a rewrite rule to MPLS packets only, and write the CoS value to MPLS headers only; or apply a rewrite rule to MPLS and IPv4 packets, and write the CoS value to MPLS and IPv4 headers.
Options	<p>protocol-types can be one of the following:</p> <ul style="list-style-type: none"> • mpls—Apply a rewrite rule to MPLS packets and write the CoS value to MPLS headers. • mpls-inet-both—Apply a rewrite rule to VPN MPLS packets with IPv4 payloads. On M120, M320, MX Series, and T Series routers (except T4000 routers), and EX Series switches, write the CoS value to the MPLS and IPv4 headers. On M Series routers, initialize all ingress MPLS LSP packets with IPv4 payloads with 000 code points for the MPLS EXP value, and the configured rewrite code point for IP precedence. • mpls-inet-both-non-vpn—Apply a rewrite rule to non-VPN MPLS packets with IPv4 payloads. On M120, M320, MX Series, T Series routers, and EX Series switches write the CoS value to the MPLS and IPv4 headers. On M Series routers, initialize all ingress MPLS LSP packets with IPv4 payloads with 000 code points for the MPLS EXP value, and the configured rewrite code point for IP precedence.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Rewriting MPLS and IPv4 Packet Headers</i>

rewrite-rules (Definition)

Syntax	<pre>rewrite-rules { type <i>rewrite-name</i>{ import (<i>rewrite-name</i> default); forwarding-class <i>class-name</i> { loss-priority <i>level</i> <i>code-point</i> [<i>aliases</i>] [<i>bit-patterns</i>]; } } }</pre>
Hierarchy Level	[edit class-of-service]
Release Information	Statement introduced before Junos OS Release 7.4. ieee-802.1ad option introduced in Junos OS Release 9.2.
Description	Specify a rewrite-rules mapping for the traffic that passes through all queues on the interface.
Options	<p><i>rewrite-name</i>—Name of a <i>rewrite-rules</i> mapping.</p> <p><i>type</i>—Traffic type.</p> <p>Values: dscp, dscp-ipv6, exp, ieee-802.1, ieee-802.1ad, inet-precedence</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring Rewrite Rules</i>

rewrite-rules (Interfaces)

Syntax	<pre>rewrite-rules { dscp (rewrite-name default) protocol mpls; dscp-ipv6 (rewrite-name default) protocol mpls; exp (rewrite-name default) protocol protocol-types; exp-push-push-push default; exp-swap-push-push default; ieee-802.1 (rewrite-name default) vlan-tag (outer outer-and-inner); ieee-802.1ad (rewrite-name default) vlan-tag (outer outer-and-inner); inet-precedence (rewrite-name default) protocol mpls; }</pre>
Hierarchy Level	<p>[edit class-of-service interfaces <i>interface-name</i>],</p> <p>[edit class-of-service interfaces <i>interface-name</i> unit <i>logical-unit-number</i>]</p>
Release Information	Statement introduced before Junos OS Release 7.4.
Description	<p>Associate a rewrite-rules configuration or default mapping with a specific interface.</p> <p>The [edit class-of-service interfaces <i>interface-name</i>] hierarchy level is not supported on M Series routers.</p> <p>The [edit class-of-service interfaces <i>interface-name</i> unit <i>logical-unit-number</i>] hierarchy level is not supported on ACX Series routers.</p> <p>On MX Series routers, although you can configure firewall filters and CoS rewrite rules on IRB interfaces, we recommend that you do not configure these functionalities on IRB interfaces because they do not work properly.</p> <p>On an MX Series router and on an EX Series switch, exp-push-push-push, exp-swap-push-push, and frame-relay-de are not supported on an integrated routing and bridging (IRB) interface.</p> <p>On an ACX Series router, only the outer tag is supported for dscp, inet-precedence, and ieee802.1.</p> <p>On M Series routers only, if you include the control-word statement at the [edit protocols l2circuit neighbor address interface <i>interface-name</i>] hierarchy level, the software cannot rewrite MPLS EXP bits.</p> <p>For IQ PICs, you can configure only one IEEE 802.1 rewrite rule on a physical port. All logical ports (units) on that physical port should apply the same IEEE 802.1 rewrite rule.</p> <p>On M320 and T Series routers (except for T4000 routers with Type 5 FPCs), for a single interface, you cannot enable a rewrite rule on a subset of forwarding classes. You must assign a rewrite rule to either none of the forwarding classes or all of the forwarding classes. When you assign a rewrite rule to a subset of forwarding classes, the commit does not fail, and the subset of forwarding classes works as expected. However, the forwarding classes to which the rewrite rule is not assigned are rewritten to all zeros.</p>

For example, if you configure a Differentiated Services code point (DSCP) rewrite rule, the bits in the forwarding classes to which you do not assign the rewrite rule are rewritten to 000000. If you configure an IP precedence rewrite rule, the bits in the forwarding classes to which you do not assign the rewrite rule are rewritten to 000.

Options *rewrite-name*—Name of a **rewrite-rules** mapping configured at the **[edit class-of-service rewrite-rules]** hierarchy level.

default—The default mapping.

The remaining statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Configuring Rewrite Rules*
- [rewrite-rules \(Definition\) on page 30](#)
- *Applying Rewrite Rules to Output Logical Interfaces*

unit

Syntax	<pre> unit <i>logical-unit-number</i> { classifiers { type (<i>classifier-name</i> default) family (mpls all); } forwarding-class <i>class-name</i>; fragmentation-map <i>map-name</i>; input-traffic-control-profile <i>profile-name</i> shared-instance <i>instance-name</i>; output-traffic-control-profile <i>profile-name</i> shared-instance <i>instance-name</i>; per-session-scheduler; rewrite-rules { dscp (<i>rewrite-name</i> default); dscp-ipv6 (<i>rewrite-name</i> default); exp (<i>rewrite-name</i> default) <i>protocol</i> <i>protocol-types</i>; exp-push-push default; exp-swap-push-push default; ieee-802.1 (<i>rewrite-name</i> default) vlan-tag (outer outer-and-inner); inet-precedence (<i>rewrite-name</i> default); } scheduler-map <i>map-name</i>; shaping-rate <i>rate</i>; } </pre>
Hierarchy Level	[edit class-of-service interfaces <i>interface-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Configure a logical interface on the physical device. You must configure a logical interface to be able to use the physical device.
Options	<p><i>logical-unit-number</i>—Number of the logical unit.</p> <p>Range: 0 through 16,384</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> <i>Understanding How Behavior Aggregate Classifiers Prioritize Trusted Traffic</i> <i>Configuring Rewrite Rules</i>

[edit interfaces] Hierarchy Level

The following statement hierarchy can also be included at the [edit logical-systems *logical-system-name*] hierarchy level.

```

interfaces {
  interface-name {
    ... the "interface-name" subhierarchy appears after the main [edit interfaces] hierarchy level ...
  }
}

```

```
}
interface-set interface-set-name {
  interface interface-name {
    (unit unit-number | vlan-tags-outer vlan-tag);
  }
}
irb {
  accounting-profile name;
  description text;

  (gratuitous-arp-reply | no-gratuitous-arp-reply);
  hold-time up milliseconds down milliseconds;
  mtu bytes;
  no-gratuitous-arp-request;

  traceoptions {
    flag flag;
  }
  (traps | no-traps);
  unit logical-unit-number {
    accounting-profile name;
    bandwidth rate;
    description text;
    disable;
    encapsulation type;
    family inet {
      accounting {
        destination-class-usage;
        source-class-usage {
          input;
          output;
        }
      }
    }
  }
  address ipv4-address {
    arp ip-address (mac | multicast-mac) mac-address <publish>;
    broadcast address;
    preferred;
    primary;
    vrrp-group group-id {
      (accept-data | no-accept-data);
      advertise-interval seconds;
      advertisements-threshold number;
      authentication-key key;
      authentication-type authentication;
      fast-interval milliseconds;
      (preempt | no-preempt) {
        hold-time seconds;
      }
      priority number;
      track {
        interface interface-name {
          bandwidth-threshold bits-per-second priority-cost priority;
          priority-cost priority;
        }
        priority-hold-time seconds;
      }
    }
  }
}
```



```

        route prefix/prefix-length routing-instance instance-name priority-cost priority;
    }
    virtual-address [ addresses ];
    vrrp-inherit-from vrrp-group;
}
}
filter {
    input filter-name;
    output filter-name;
}
mtu bytes;
no-neighbor-learn;
no-redirects;
primary;
rpf-check {
    fail-filter filter-name;
    mode {
        loose;
    }
}
targeted-broadcast {
    forward-and-send-to-re;
    forward-only;
}
}
family inet6 {
    accounting {
        destination-class-usage;
        source-class-usage {
            input;
            output;
        }
    }
}
address address {
    eui-64;
    ndp ip-address (mac | multicast-mac) mac-address <publish>;
    preferred;
    primary;
    vrrp-inet6-group group-id {
        accept-data | no-accept-data;
        advertisements-threshold number;
        authentication-key key;
        authentication-type authentication;
        fast-interval milliseconds;
        inet6-advertise-interval milliseconds;
        preempt | no-preempt {
            hold-time seconds;
        }
    }
    priority number;
    track {
        interface interface-name {
            bandwidth-threshold bandwidth priority-cost number;
            priority-cost number;
        }
    }
    priority-hold-time seconds;
    route ip-address/mask routing-instance instance-name priority-cost cost;
}

```

```

    }
    virtual-inet6-address [addresses];
    virtual-link-local-address ipv6-address;
    vrrp-inherit-from {
        active-group group-number;
        active-interface interface-name;
    }
}
}
(dad-disable | no-dad-disable);
filter {
    input filter-name;
    output filter-name;
}
mtu bytes;
nd6-stale-time seconds;
no-neighbor-learn;
no-redirects;
policer {
    input policer-name;
    output policer-name;
}
rpf-check {
    fail-filter filter-name;
    mode {
        loose;
    }
}
}
}
family iso {
    address interface-address;
    mtu bytes;
}
family mpls {
    filter {
        input filter-name;
        output filter-name;
    }
    mtu bytes;
    policer {
        input policer-name;
        output policer-name;
    }
}
native-inner-vlan-id vlan-id;
proxy-arp (restricted | unrestricted);
(traps | no-traps);
vlan-id-list [vlan-id's];
vlan-id-range [vlan-id-range];
}
}
traceoptions {
    file <filename> <files number> <match regular-expression> <size maximum-file-size>
        <world-readable | no-world-readable>;
    flag flag <disable>;
    no-remote-trace;
}

```

```

    }
}

interfaces {
  interface-name {
    disable;
    accounting-profile name;
    aggregated-ether-options {
      ethernet-switch-profile {
        tag-protocol-id [ hexadecimal-identifiers ];
      }
      (flow-control | no-flow-control);
      lacp {
        (active | passive);
        admin-key key;
        fast-failover;
        link-protection {
          disable;
          (revertive | non-revertive);
        }
        periodic (fast | slow);
        system-id mac-address;
        system-priority priority;
      }
      (link-protection | no-link-protection);
      link-speed (100m | 1g | 8g | 10g | 40g | 50g | 80g | 100g | oc192);
      logical-interface-fpc-redundancy;
      (loopback | no-loopback);
      mc-ae {
        chassis-id chassis-id;
        events {
          iccp-peer-down {
            force-icl-down;
            prefer-status-control-active;
          }
        }
        mc-ae-id mc-ae-id;
        mode (active-active | active-standby);
        redundancy-group group-id;
        status-control (active | standby);
      }
      minimum-links number;
      rebalance-periodic {
        start-time time;
        interval number;
      }
      source-address-filter {
        mac-address;
      }
      (source-filtering | no-source-filtering);
    }
  }
  auto-configure {
    remove-when-no-subscribers;
    stacked-vlan-ranges {
      access-profile profile-name;
      authentication {

```

```

    password password-string;
    username-include {
        circuit-type;
        delimiter delimiter-character;
        domain-name domain-name-string;
        interface-name;
        mac-address;
        option-82 ( circuit-id | remote-id);
        radius-realm radius-realm-string;
        user-prefix user-prefix-string;
    }
}
dynamic-profile profile-name {
    accept (any | dhcp-v4 | dhcp-v6 | inet | inet6);
    ranges (any | low-tag-high-tag), (any | low-tag-high-tag);
}
}
vlan-ranges {
    access-profile profile-name;
    authentication {
        password password-string;
        username-include {
            circuit-type;
            delimiter delimiter-character;
            domain-name domain-name-string;
            interface-name;
            mac-address;
            option-82;
            radius-realm radius-realm-string;
            user-prefix user-prefix-string;
        }
    }
}
dynamic-profile profile-name {
    accept (any | dhcp-v4 | dhcp-v6 | inet | inet6);
    ranges (any | low-tag)—(any | high-tag);
}
}
override tag vlan-tag dynamic-profile profile name;
}
encapsulation (ethernet-bridge | ethernet-vpls | extended-vlan-bridge |
    extended-vlan-vpls | flexible-ethernet-services | vlan-vpls);
ether-options {
    802.3ad {
        aex;
        (backup | primary);
        lacp {
            force-up;
            port-priority
        }
    }
}
asynchronous-notification;
(auto-negotiation | no-auto-negotiation);
ethernet-switch-profile {
    ethernet-policer-profile {
        input-priority-map {
            ieee802.1p premium [ values ];

```

```

    }
    output-priority-map {
        classifier {
            premium {
                forwarding-class class-name {
                    loss-priority (high | low);
                }
            }
        }
    }
    policer cos-policer-name {
        aggregate {
            bandwidth-limit bps;
            burst-size-limit bytes;
        }
        premium {
            bandwidth-limit bps;
            burst-size-limit bytes;
        }
    }
    tag-protocol-id;
}
(mac-learn-enable | no-mac-learn-enable);
}
(flow-control | no-flow-control);
ignore-l3-incompletes;
link-mode (automatic | full-duplex | half-duplex);
(lloopback | no-loopback);
keepalives <interval seconds> <down-count number> <up-count number>;
speed (1g | 10m | 100m | 10m-100m | auto-negotiation);
source-address-filter {
    mac-address;
}
source-filtering | no-source-filtering;
}
flexible-vlan-tagging;
(gratuitous-arp-reply | no-gratuitous-arp-reply);
hold-time (up milliseconds | down milliseconds);
interface-transmit-statistics;
(keepalives <down-count number> <interval seconds> <up-count number> |
no-keepalives);
layer2-policer {
    apply-groups [ group-names ];
    apply-groups-except [ group-names ];
}
link-mode (automatic | full-duplex);
mac mac-address;
mtu bytes;
multi-chassis-protection peer-ip-address {
    interface interface-name;
}
native-vlan-id number;
no-gratuitous-arp-request;
optics-options {
    alarm low-light-alarm {
        (link-down | syslog);
    }
}

```

```

    }
    warning low-light-warning {
        (link-down | syslog);
    }
    wavelength nm;
}
passive-monitor-mode;
per-unit-scheduler;
speed (10m | 100m | 1g | auto | oc3 | oc12 | oc48);
stacked-vlan-tagging;
traceoptions {
    flag flag;
}
transmit-bucket {
    overflow discard;
    rate percentage;
    threshold bytes;
}
(traps | no-traps);
unidirectional;
vlan-tagging;
}

```

```

interface-name {
    unit logical-unit-number {
        disable;
        accept-source-mac {
            mac-address mac-address {
                policer {
                    input policer-name;
                    output policer-name;
                }
            }
        }
        accounting-profile name;
        advisory-options {
            downstream-rate rate;
            upstream-rate rate;
        }
        arp-resp (restricted|unrestricted);
        bandwidth rate;
        clear-dont-fragment-bit;
        copy-tos-to-outer-ip-header;
        demux-destination family;
        encapsulation (vlan-bridge | vlan-vpls);
        epd-threshold cells plp1 cells;
        filter filter-name;
        inner-vlan-id-range start start-id end end-id;
        input-vlan-map {
            (pop | pop-pop | pop-swap | push | push-push | swap | swap-push | swap-swap);
            inner-tag-protocol-id tpid;
            inner-vlan-id number;
            tag-protocol-id tpid;
            vlan-id number;
        }
    }
}

```

```

interface-shared-with psdnumerical-index;
layer2-policer {
    input-hierarchical-policer policer-name;
    input-policer policer-name;
    input-three-color policer-name;
    output-policer policer-name;
    output-three-color policer-name;
}
multi-chassis-protection peer-ip-address {
    interface interface-name;
}
native-inner-vlan-id number;
output-vlan-map {
    (pop | pop-pop | pop-swap | push | push-push | swap | swap-push | swap-swap);
    inner-tag-protocol-id tpid;
    inner-vlan-id number;
    tag-protocol-id tpid;
    vlan-id number;
}
peer-interface interface-name;
peer-unit unit-number;
plp-to-clp;
proxy-arp <restricted | unrestricted>;
rpm {
    (client | server);
    twamp-server;
}
swap-by-poppush;
vlan-id number;
vlan-id-list [ vlan-id vlan-id-vlan-id ];
vlan-id-range number-number;
vlan-tags (inner <tpid.>vlan-id | inner-list [ vlan-id vlan-id-vlan-id ] |
    inner-range <tpid.>vlan-id-vlan-id) outer <tpid.>vlan-id;
}

unit logical-unit-number {
    family ethernet-switching {
        filter {
            group filter-group-number;
            (input filter-name | input-list [ filter-names ]);
            (output filter-name | output-list [ filter-names ]);
            (inner-vlan-id-list [ vlan-ids ] | vlan-id number | vlan-id-list [ number
                number-number ]);
            interface-mode (access | trunk);
            policer {
                input policer-name;
                output policer-name;
            }
            vlan-rewrite {
                translate old-vlan-id new-vlan-id;
            }
            vlan {
                members [ all vlan-identifiers ];
            }
        }
    }
    family inet {

```

```

filter {
    group filter-group-number;
    (input filter-name | input-list [ filter-names ]);
    (output filter-name | output-list [ filter-names ]);
}
input-hierarchical-policer policer-name;
mac-validate (loose | strict);
mtu bytes;
no-neighbor-learn;
no-redirects;
policer {
    arp policer-template-name;
    input policer-name;
    output policer-name;
}
primary;
receive-options-packets;
receive-ttl-exceeded;
rpf-check {
    fail-filter filter-name;
    mode loose;
}
sampling {
    (input | output | input output);
}
simple-filter {
    input filter-name;
}
targeted-broadcast {
    forward-and-send-to-re;
    forward-only;
}
unnumbered-address interface-name <destination address>
    <destination-profile profile-name> <preferred-source-address address>;
}

family inet6 {
    address ipv6-address {
        destination destination-address;
        eui-64;
        ndp ipv6-address <l2-interface interface-name> <(mac mac-address |
            multicast-mac multicast-mac-address) <publish>>;
        preferred;
        primary;
        vrrp-inet6-group group-number {
            (accept-data | no-accept-data);
            fast-interval milliseconds;
            inet6-advertise-interval seconds;
            (no-preempt; | ... the following preempt statement ...)
            preempt {
                hold-time seconds;
            }
            priority number;
            track {
                interface interface-name {

```



```

        bandwidth-threshold bits-per-second priority-cost priority;
        priority-cost priority;
    }
    priority-hold-time seconds;
    route ip-address-prefix/prefix-length routing-instance instance-name
        priority-cost priority;
    }
    virtual-inet6-address [ addresses ];
    virtual-link-local-address ipv6-address;
    vrrp-inherit-from {
        active-group group-number;
        active-interface interface-name;
    }
}
(dad-disable | no-dad-disable);
filter {
    group filter-group-number;
    (input filter-name | input-list [ filter-names ]);
    (output filter-name | output-list [ filter-names ]);
}
input-hierarchical-policer policer-name;
mtu bytes;
nd6-stale-time seconds;
no-neighbor-learn;
policer {
    input policer-name;
    output policer-name;
}
rpf-check {
    fail-filter filter-name;
    mode loose;
}
sampling {
    (input | output | input output);
}
unnumbered-address interface-name preferred-source-address address;
}

family iso {
    address iso-address;
    mtu bytes;
}

family mlfr-end-to-end {
    bundle logical-interface-name;
}

family mpls {
    filter {
        group filter-group-number;
        (input filter-name | input-list [ filter-names ]);
        (output filter-name | output-list [ filter-names ]);
    }
}

```

```
    }
    input-hierarchical-policer policer-name;
    maximum-labels maximum-labels;
    mtu bytes;
    policer {
        input policer-name;
        output policer-name;
    }
}

family vpls {
    core-facing;
    filter {
        group filter-group-number;
        (input filter-name | input-list [ filter-names ]);
        (output filter-name | output-list [ filter-names ]);
    }
    policer {
        input policer-name;
        output policer-name;
    }
}
}
```

Related Documentation

- *Notational Conventions Used in Junos OS Configuration Hierarchies*

copy-tos-to-outer-ip-header

Syntax	copy-tos-to-outer-ip-header;
Hierarchy Level	[edit interfaces <i>gr-fpc/pic/port</i> unit <i>logical-unit-number</i>], [edit interfaces gre unit <i>logical-unit-number</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>gr-fpc/pic/port</i> unit <i>logical-unit-number</i>], [edit logical-systems <i>logical-system-name</i> interfaces gre unit <i>logical-unit-number</i>]
Release Information	Statement introduced in Junos OS Release 8.2. Support for GRE interfaces for Generalized MPLS (GMPLS) introduced in Junos OS Release 12.3R7.
Description	For GRE tunnel interfaces and GRE interfaces for GMPLS control channels only, enable the inner IP header's ToS bits to be copied to the outer IP packet header.
Default	If you omit this statement, the ToS bits in the outer IP header are set to 0.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring a GRE Tunnel to Copy ToS Bits to the Outer IP Header on page 12

