JUNIPER NETWORKS | Engineering Simplicity

Junos® OS

Subscriber-Aware and Application-Aware Traffic Treatment User Guide

JUNOS

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

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

Table of Contents

3

Configuring Reporting for Subscriber-Aware Data Sessions

4

Modifying Subscriber-Aware Configuration

5

Monitoring and Troubleshooting

6 **Configuration Statements and Operational Commands**

About This Guide

Use this guide to configure and monitor subscriber-aware and application-aware traffic policies. This lets you identify the mobile or fixed-line subscriber associated with a data session, and enforce traffic treatment for the subscriber based on Layer 7 or Layer 3/Layer 4 application information for the session.

1

PART

Subscriber-Aware and Application-Aware Traffic Treatment Overview

Subscriber-Aware and Application-Aware Traffic Treatment Overview

IN THIS CHAPTER

Subscriber-Aware and Application-Aware Traffic Treatment Overview

IN THIS SECTION

This topic contains an overview of subscriber-aware and application-aware traffic treatment.

Introduction

Junos Subscriber Aware identifies the mobile or fixed-line subscriber associated with a data session, and enforces traffic treatment based on policies assigned to the subscriber. This permits highly customizable differentiated services for subscribers. A subscriber policy can be based on Layer 7 application information for the IP flow (for example, YouTube) or can be based on Layer 3/Layer 4 information for

the IP flow (for example, the source and destination IP address). Junos Subscriber Aware resides on an MX Series router.

Subscriber-aware policies can specify the following actions:

- Redirecting HTTP traffic to another URL or IP address

- Forwarding packets to a routing instance so that packets are directed to external service chains (predefined sequence of services)

- Setting the forwarding class

- Setting the maximum bit rate

- Performing HTTP header enrichment (provided by Junos Web Aware, which resides on the same MX Series router as Junos Subscriber Aware)

- Setting the gating status to blocked or allowed

Subscriber-aware policies can also specify the time of day that the policies are in effect.

Access-Independent Subscriber Traffic Treatment

Subscriber identification for both mobile access and wireline access provides a unified experience for the subscriber, regardless of the connection method.

Junos Subscriber Aware resides on an MX Series router that is located between the gateway of the access network and the public network and network services, as shown in Figure 1 on page 4. Subscribers may be controlled by a broadband network gateway (BNG) in a wireline access network, by a gateway GPRS support node (GGSN) in a 2G or 3G network architecture, or by a Packet Data Network Gateway (PGW) in a 4G/LTE network architecture.

Figure 1: Subscriber-Aware Policy Enforcement on the MX Series
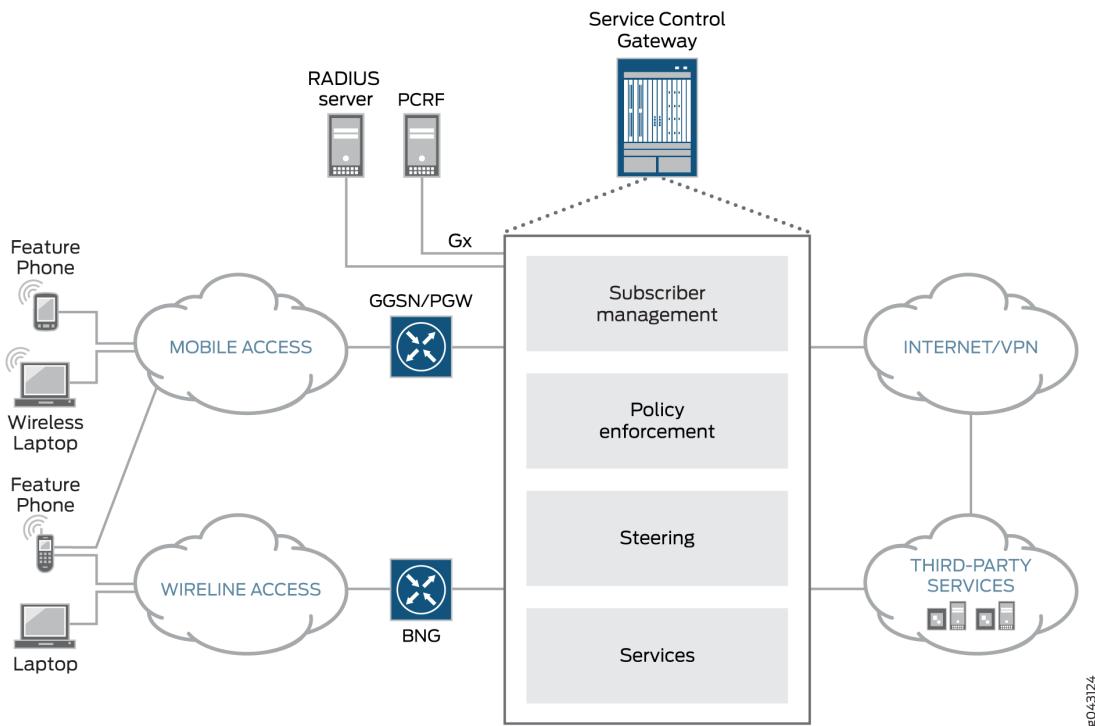
Subscriber Identification Methods

You can use the following methods to identify subscribers:

- IP-based—Processes a RADIUS accounting start request to identify the subscriber. An IP-based subscriber session is for one unique user IP address.

- IFL-based—Requires you to configure a subscriber name and specify a set of MX Series router access interfaces for the subscriber. Junos Subscriber Aware assigns all data sessions received on those interfaces to the configured subscriber.

Application Identification

Layer 7 application identification is provided by Junos Application Aware, which performs deep packet inspection (DPI) to determine whether the subscriber's data packets match an application signature. When an application is identified, the appropriate subscriber policy is applied to the packets. Juniper Networks provides a set of predefined application signatures that you can download and that are periodically updated. You can also configure your own custom application signatures.

Junos Subscriber Aware and Junos Application Aware reside on the same MX Series router, allowing policy control on a single platform.

Policy Control Methods

Subscriber-aware policies can be controlled dynamically by a policy and charging rules function (PCRF) server, can be activated by a RADIUS server, or can be under static control.

Under dynamic control, a PCRF either sends policies to the MX Series router or activates predefined policies that you configured on the MX Series router. Dynamic policy control is provided by Junos Policy Control, which resides on the same MX Series router as Junos Subscriber Aware.

Under RADIUS server control, the RADIUS server controls the activation of your predefined polices but does not send policies to the MX Series router.

Under static control, your predefined policies are not controlled by a PCRF or RADIUS server.

Subscriber-Aware Data Session Logging and Reporting

Junos Subscriber Aware can log data for subscriber-aware data sessions and send that data in an IPFIX format to an external log collector. These logs can include subscriber information, application information, HTTP metadata, data volume, time-of-day information, and source and destination details. You can then use the external collector, which is not a Juniper Networks product, to perform analytics that provide you with insights about subscriber and application usage, enabling you to create packages and policies that increase revenue.

Usage Monitoring

For subscriber data sessions that are under the dynamic policy control of a PCRF, Junos Subscriber Aware can monitor the volume of traffic or amount of time the subscriber uses during a session, and send reports to the PCRF. The PCRF can use this information to adjust the policies for a subscriber.

RELATED DOCUMENTATION

Configuring Subscriber-Aware and Application-Aware Traffic Treatment Overview

To configure subscriber-aware and application-aware traffic treatment:

1. Configure service PICs and session PICs.

See "Configuring Service PICs and Session PICs Overview " on page 11.

2. (Optional) Identify Layer 7 applications.

   a. Install application signature packages.

      See "Downloading and Installing Predefined Junos OS Application Signature Packages" on page 23.

   b. Configure custom application signatures.

      See "Configuring Custom Application Signatures" on page 25.

3. (Optional) Configure HTTP header enrichment.

   See "Configuring HTTP Header Enrichment Overview" on page 40.

4. Configure a policy enforcement method.

   • For dynamic policy control, see "Configuring Dynamic Policy Control by PCRF" on page 73.

   • For static policy control, see "Configuring Static Policy Control" on page 74.

   • For RADIUS server policy control, see "Configuring Policy Control by RADIUS Servers" on page 75.

5. Configure the policy enforcement for an IP-based subscriber. An IP-based subscriber session handles traffic for one unique user IP address.

   • If the MX Series router is identified as a RADIUS server for the access gateway, see "Configuring IP-Based TDF Subscriber Setup When MX Series Router Is a RADIUS Server" on page 113

   • If the MX Series router is not identifed as a RADIUS server for the access gateway, see "Configuring IP-Based TDF Subscriber Setup When Accounting Requests Are Snooped" on page 114

6. Configure the policy enforcement for an IFL-based subscriber. An IFL-based subscriber session handles all the traffic received on a specific set of interfaces.

   See "Configuring IFL-Based TDF Subscriber Setup" on page 134.

7. Apply services to a subscriber.

   See "Applying Services to Subscriber-Aware Traffic with a Service Set" on page 144.

8. (Optional) If you configured dynamic policy control, configure Diameter.

   See "Configuring Diameter Overview" on page 149.

RELATED DOCUMENTATION

Subscriber-Aware and Application-Aware Traffic Treatment Overview | 2

2
PART

Applying Subscriber-Aware and Application-Aware Policies and Services

CHAPTER 2

Configuring the Service PIC, Session PIC, and TDF Gateway

TDF Gateway Service PICs and Session PICs for Subscriber-Aware Traffic Treatment

You must configure at least one TDF gateway, one service PIC, and one session PIC to operate subscriber-aware traffic treatment. Each service PIC and session PIC is configured on an MS-MPC, and assigned to a TDF gateway.

TDF Gateway

The traffic detection function (TDF) gateway on the MX Series router establishes a context and framework for configuring subscriber-aware services. You assign service PICs and session PICs to the TDF gateway, and specify the call admission control (CAC) parameters for subscriber sessions.

Service and Session PICs

A service PIC provides subscriber-aware policy enforcement and traffic redirection (*steering*) that is application-aware. Traffic steering refers to the capability to direct or traverse traffic from a specified source to an endpoint or the adjacent network element in a routing path. The service PIC is configured with software plugins to perform the configured or requested services, which include the policy and charging enforcement function (PCEF), application detection and control, HTTP header enrichment, HTTP redirect, and network address translation.

The service PIC also stores the policy and charging control (PCC) rules that it enforces, and holds the subscriber records and rules that are sent from the session PIC.

The subscriber's assigned TDF logical interface (mif) and the service set that is applied to the mif determine the service PIC to which a packet is sent. See .

A session PIC supports access subscriber session setup and management, enabling the steering of subscriber traffic to the correct services PIC. The session PIC also sets up a session with the policy and charging rules function (PCRF) so it can receive subscriber PCC rules from the PCRF and send application-start messages to the PCRF.

shows an overview of a service PIC and a session PIC and their functions.

Figure 2: Service PIC and Session PIC Overview
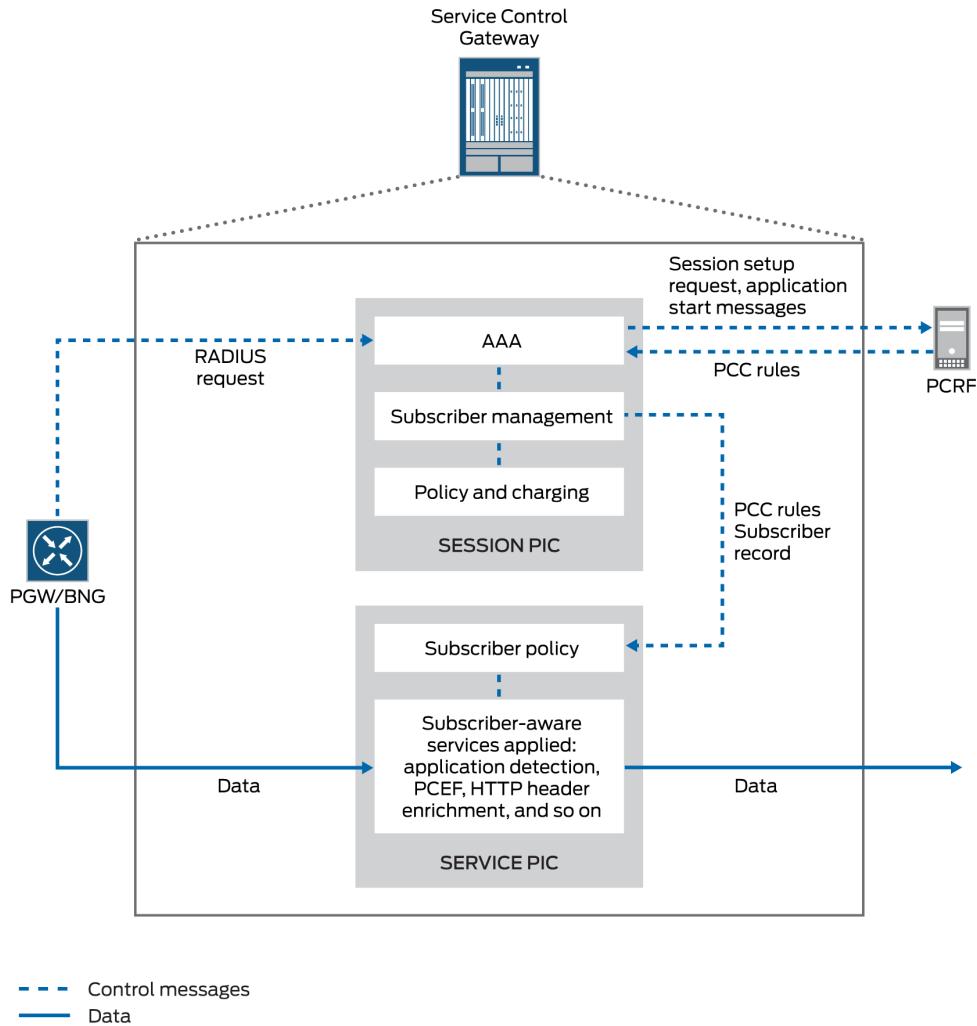
Redundancy for Service PICs and Session PICs

You can configure a service PIC or a session PIC as an individual PIC or with a backup for redundancy. You can configure redundancy by including the interfaces for the primary and the backup PICs in an aggregated multiservices (AMS) interface .

You can configure a session PIC with 1:1 redundancy — a primary session PIC has one backup PIC that does not back up any other session PICs.

You can configure service PICs with N:1 redundancy — multiple service PICs can share the same backup MS-PIC.

In addition to the redundancy configuration, each PIC that is a primary or backup needs to be configured as a session PIC or service PIC at the [edit unified-edge gateways tdf *gateway-name* system] hierarchy level.

Configuring Aggregated Multiservices Interfaces

Configuring Service PICs and Session PICs Overview

You must configure at least one service PIC and one session PIC under a TDF gateway. The service PIC provides subscriber-aware services, such as the policy and charging enforcement function (PCEF), application detection and control, and HTTP header enrichment. The session PIC supports access subscriber sessions, policy and charging rules function (PCRF) sessions, and PCEF library installation from the PCRF.

You can configure service PICs and session PICs on MS-MPCs, and you can configure them either as a member of a redundant group by using an aggregated multiservices (AMS) interface or as a standalone service PIC or session PIC.

To configure service and session PICs:

1. Configure the TDF gateway.

   See "Configuring a TDF Gateway" on page 14.

2. If you want any of the service or session PICs to be members of redundant groups, configure an aggregated multiservices (AMS) interface for each group.

   See *Configuring Aggregated Multiservices Interfaces*.

3. If you want any of the service or session PICs not to be members of redundant groups, configure a services interface.

   See "Configuring a Services Interface for a Session PIC or Service PIC" on page 14.

4. Install predefined groups that are needed for configuration of the service PICs and session PICs.

   See "Making Predefined Groups Available for Session PIC and Service PIC Configuration" on page 16.

5. Configure each service PIC.

   See "Configuring Service PICs" on page 16.

6. Configure each session PIC.

   See "Configuring Session PICs" on page 18.

Preconfigured Groups for Service PICs and for Session PICs Overview

To simplify configuration, Junos Subscriber Aware software includes predefined configuration groups that include the parameters for stable operation of session PICs and service PICs. These groups are included in the `/etc/config/tdf-defaults.conf` file, which you load and then merge with your configuration. Next, you apply the appropriate group to each session PIC and service PIC configuration as follows:

- For each session PIC, apply the `tdf-session-xlp` group.

- For each service PIC that requires application identification but not HTTP header enrichment, apply the `tdf-services-xlp-dpi` group.

- For each service PIC that requires both application identification and HTTP header enrichment, configure the `tdf-services-xlp-dpi-with-hcm` group.

The predefined `tdf-session-xlp` group contains the following statements:

```
[edit groups]
tdf-session-xlp {
    chassis {
        fpc <*> {
            pic <*> {
                adaptive-services {
                    service-package {
                        extension-provider {
                            boot-os embedded-junos64;
                            package jservices-mobile;
                        }
                    }
                }
            }
        }
    }
}
```

The predefined `tdf-services-xlp-dpi` group contains the following statements:

```
[edit groups]
tdf-services-xlp-dpi {
    chassis {
        fpc <*> {
            pic <*> {
                adaptive-services {
                    service-package {
                        extension-provider {
                            boot-os embedded-junos64;
                            package jservices-mss;
                            package jservices-jdpi;
                            package jservices-pcef;
                        }
                    }
                }
            }
        }
    }
}
```

The predefined `tdf-services-xlp-dpi-with-hcm` group contains the following statements:

```
[edit groups]
tdf-services-xlp-dpi-with-hcm {
    chassis {
        fpc <*> {
            pic <*> {
                adaptive-services {
                    service-package {
                        extension-provider {
                            boot-os embedded-junos64;
                            package jservices-mss;
                            package jservices-jdpi;
                            package jservices-pcef;
                            package jservices-hcm;
                            package jservices-crypto-base;
                        }
                    }
                }
```

```
            }
         }
      }
   }
```

Making Predefined Groups Available for Session PIC and Service PIC Configuration **| 16**

Configuring Session PICs **| 18**

Configuring Service PICs **| 16**

Configuring a Services Interface for a Session PIC or Service PIC

If a service PIC or a session PIC is not part of a redundant group (the service interface is not part of an aggregated multiservices interface), you must configure a services interface on the MS-MPC for the service PIC.

- Configure the services interface.

```
[edit]
user@host# set interfaces ms-fpc/pic/0 unit logical-unit-number family family address address
```

Configuring Aggregated Multiservices Interfaces

Configuring Service PICs **| 16**

Configuring Session PICs **| 18**

TDF Gateway Service PICs and Session PICs for Subscriber-Aware Traffic Treatment **| 8**

Configuring a TDF Gateway

To run Junos Subscriber Aware, you must configure a traffic detection function (TDF) gateway on the MX Series router. The TDF gateway establishes a context and framework for configuring subscriber-

aware services for subscriber data that is accessing the network through the MX Series router. You also specify the call admission control (CAC) parameters for the TDF gateway.

To configure the TDF gateway:

1. Configure a name for the TDF gateway.

```
[edit unified-edge gateways]
user@host# set tdf gateway-name
```

2. Configure the threshold for the maximum amount of CPU that the TDF gateway can use as a percentage from 1 through 90.

```
[edit unified-edge gateways tdf gateway-name]
user@host# set cac cpu cpu-pct
```

If the amount of CPU that the TDF gateway uses reaches the threshold, the SNMP trap **jnxScgSMCPUThreshHigh** is generated.

3. Configure the maximum number of TDF subscriber sessions that can be running, expressed in thousands of sessions.

```
[edit unified-edge gateways tdf gateway-name]
user@host# set cac maximum-sessions max-sessions
```

You can configure from 10 through 5000 sessions.

4. Configure the trap threshold for the number of TDF subscriber sessions as a percentage of the maximum number of sessions.

```
[edit unified-edge gateways tdf gateway-name]
user@host# set cac maximum-sessions-trap-percentage max-sessions-pct
```

If the number of subscriber sessions reaches the threshold, the SNMP trap **jnxScgSMSessionThreshHigh** is generated.

5. Configure the threshold for the maximum amount of memory that the TDF gateway can use, as a percentage from 1 through 90.

```
[edit unified-edge gateways tdf gateway-name]
user@host# set cac memory memory-pct
```

If the amount of memory that the TDF gateway uses reaches the threshold, the SNMP trap **jnxScgSMMemoryThreshHigh** is generated.

Making Predefined Groups Available for Session PIC and Service PIC Configuration

You must make the predefined session PIC and service PIC groups available in your configuration. These groups are used when you configure the session PICs and the service PICs.

To make the predefined groups available in your configuration:

- Load and merge the `tdf-defaults.conf` file.

```
[edit]
user@host# load merge /etc/config/tdf-defaults.conf
```

Configuring Service PICs

An MS-MPC must have a service interface configured as a service PIC in order to provide subscriber-aware services, such as the policy and charging enforcement function (PCEF), application detection and control, or HTTP header enrichment. Repeat this procedure for each service interface that you want to serve as a service PIC.

Before you begin to configure a service PIC:

- Make sure that you installed the predefined groups.

- If the service PIC is not part of a redundant group, make sure that you have configured the service interface on the MS-MPC.

- If the service PIC is to function as a member of a redundant group, make sure that you have configured an aggregated multiservices (AMS) interface with the service interface as a member interface.

To configure a service PIC:

1. Add the MS-MPC service interface to the list of service PICs.

```
[edit unified-edge gateways tdf gateway-name system]
user@host# set service-pics interface interface-name
```

where *interface-name* is ams*n* if you have redundancy configured and is ms-*fpc*/*pic*/0 if you do not have redundancy configured.

2. Perform one of the following actions:

- If application identification is required but not HTTP header enrichment, configure the `tdf-services-xlp-dpi` group to run on the PIC.

```
[edit chassis]
user@host# set fpc slot-number pic pic-number apply-groups tdf-services-xlp-dpi
```

- If both application identification and HTTP header enrichment are required, configure the `tdf-services-xlp-dpi-with-hcm` group to run on the PIC.

```
[edit chassis]
user@host# set fpc slot-number pic pic-number apply-groups tdf-services-xlp-dpi-with-hcm
```

3. (Optional) For Next Gen Services, enable subscriber awareness. This steps loads MSS, PCEF, HCM (all subscriber related plugins) on the PIC.

```
[edit chassis]
user@host# set fpc slot-number pic pic-number subscriber-aware-services
```

Configuring Session PICs

An MS-MPC must have a service interface configured as a session PIC in order to support access subscriber sessions, policy and charging rules function (PCRF) sessions, and PCEF library installation from the PCRF. Repeat this procedure for each service interface that you want to serve as a session PIC.

Before you begin to configure a session PIC:

- Make sure that you have installed the predefined groups.

- If the session PIC is not part of a redundant group, make sure that you have configured the service interface on the MS-MPC.

- If the session PIC is to function as a member of a redundant group, make sure that you have configured an aggregated multiservices (AMS) interface with the service interface as a member interface.

To configure a session PIC:

1. Add the MS-MPC service interface to the list of session PICs.

```
[edit unified-edge gateways tdf gateway-name system]
user@host# set session-pics interface interface-name
```

where *interface-name* is ams*n* if you have redundancy configured and is ms-*fpc*/*pic*/0 if you do not have redundancy configured.

2. Configure the `tdf-session-xlp` group to run on the PIC.

```
[edit chassis]
user@host# set fpc slot-number pic pic-number apply-groups tdf-session-xlp
```

Configuring Tracing for TDF Gateway

To configure tracing operations for the TDF gateway:

1. Specify that you want to configure tracing options for the TDF gateway.

```
[edit unified-edge gateways tdf gateway-name]
user@host# edit traceoptions
```

2. Configure the name of the file used for the trace output.

```
[edit unified-edge gateways tdf gateway-name traceoptions]
user@host# set file file-name
```

3. (Optional) Configure the maximum size of each trace file.

```
[edit unified-edge gateways tdf gateway-name traceoptions]
user@host# set file file-name size size
```

4. (Optional) Configure the maximum number of trace files.

```
[edit unified-edge gateways tdf gateway-name traceoptions]
user@host# set file file-name files number
```

5. (Optional) Configure the read permissions for the log file.

```
[edit unified-edge gateways tdf gateway-name traceoptions]
user@host# set file file-name (no-world-readable | world-readable)
```

6. (Optional) Disable remote tracing capabilities.

```
[edit unified-edge gateways tdf gateway-name traceoptions]
user@host# set no-remote-trace
```

7. Configure flags to filter the operations to be logged.

```
[edit unified-edge gateways tdf gateway-name traceoptions]
user@host# set flag flag
```

describes the flags that you can include.

Table 1: Trace Flags

Flag	Description
all	Trace all operations.
bulkjob	Trace events that are handled by bulk jobs in order to prevent system overload.
config	Trace configuration events.
cos-cac	Trace class of service (CoS) and call admission control (CAC) events.
ctxt	Trace user equipment, Packet Data Network (PDN), or bearer context events.
fsm	Trace mobile subscriber finite state machine (FSM) events.
gtpu	Trace GPRS tunneling protocol, user plane (GTP-U) events.
ha	Trace high availability events.
init	Trace initialization events.
pfem	Trace Packet Forwarding Engine Manager events.

Table 1: Trace Flags *(Continued)*

Flag	Description
stats	Trace `stats` events. This flag is used internally by Juniper Networks engineers.
waitq	Trace `waitq` events. This flag is used internally by Juniper Networks engineers.

8. Configure the level of tracing.

```
[edit unified-edge gateways tdf gateway-name traceoptions]
user@host# set level (all | critical | error | info | notice | verbose | warning)
```

RELATED DOCUMENTATION

traceoptions

CHAPTER 3

Configuring Application Identification

Application Identification Overview

Junos Application Aware is an infrastructure plug-in on MS-MPC service PICs and on the MX-SPC3 services card that provides information to clients about application protocol bundles based on deep packet inspection (DPI) of application signatures. These clients can be any of the plug-ins on the MX Series router service chain, such as traffic detection function (TDF), that request application classification data. Starting in Junos OS Release 16.1R4 and Junos OS Release 17.2R1, application identification is available in Junos OS Broadband Subscriber Management. Starting in Junos OS Release 19.3R2, application identification is also supported for Broadband Subscriber Management on the MX-SPC3 services card if you have enabled Next Gen Services on the MX240, MX480 or MX960 router.

In application identification, you can apply application signatures as follows:

- **Predefined signatures**—Junos Application Aware comes with a bundle of predefined, preinstalled application signatures, but we recommend that you download and install the latest version of predefined signatures. As new sets of signatures are supported, they are compiled and made available for you to download.

- **Custom application signatures**—For any application signatures that are not predefined, you can create custom signatures for HTTP, SSL, and stream signature contexts and install them for application identification. After you have configured and committed custom signatures, they are serialized and merged with the predefined application signatures. You can specify the following types of custom application signatures:

  - **Address based**—You can define an application identification based on a specific IP address, or port, or both where a source IP address, destination IP address, or both are used for a known

application in a customer's network. This is useful, for example, when a Session Initiation Protocol (SIP) server initiates a session from its well known port, 5060. The customer can put the SIP server IP address and port 5060 as source IP/port for the SIP application. This method provides efficiency and accuracy of application identification for customer's network.

- **Internet Control Message Protocol (ICMP) based**—Application identification based on types of ICMP messages.

- **IP protocol based**—Application identification based on IP protocol. TCP, UDP, and ICMP are not supported for this method of signature creation.

- **Pattern-matching signatures**—Application based on pattern matching combined with Layer 7 protocol identification.

Change History Table

Feature support is determined by the platform and release you are using. Use Feature Explorer to determine if a feature is supported on your platform.

Release	Description
19.3R1	Starting in Junos OS Release 19.3R2, application identification is also supported for Broadband Subscriber Management on the MX-SPC3 services card if you have enabled Next Gen Services on the MX240, MX480 or MX960 router.
16.1R4	Starting in Junos OS Release 16.1R4 and Junos OS Release 17.2R1, application identification is available in Junos OS Broadband Subscriber Management.

RELATED DOCUMENTATION

Configuring Custom Application Signatures | **25**

Downloading and Installing Predefined Junos OS Application Signature Packages | **23**

Downloading and Installing Predefined Junos OS Application Signature Packages

> (i) **NOTE**: Starting in Junos OS Release 19.3R2 and 19.4R1, application identification is also supported for Broadband Subscriber Management if you have enabled Next Gen Services on the MX240, MX480 or MX960 router with the MX-SPC3 card.

To download, install, and verify the installation of predefined Junos OS application signature packages:

1. Use `download ignore-server-validation` if you want to skip server certification validation during the download. Validation is enabled by default.

```
[edit services application-identification]
user@host# set download ignore-server-validation
```

2. Configure the URL for the application signature packages server.

```
[edit services application-identification]
user@host# set download url https://signatures.juniper.net/cgi-bin/index.cgi
```

3. Download the application signature package.

   - To download the latest signature package, enter the following command:

   ```
   user@host> request services application-identification download
   ```

   - To download a specific, known signature package, include the version number:

   ```
   user@host> request services application-identification download version version-number
   ```

4. Confirm the successful download of the package.

```
user@host> request services application-identification download status
```

```
Downloading application package succeed.
```

5. Install the application signature package.

```
user@host> request services application-identification install
```

6. Confirm the successful installation of the application signature package.

```
user@host> request services application-identification install status
```

```
Compiling application signatures of package version.
```

or

```
Install application package succeed
```

7. View the protocol bundle status:

```
user@host> show services application-identification status
```

Configuring Custom Application Signatures

> ⓘ **NOTE**: Starting in Junos OS Release 19.3R2 and 19.4R1, application identification is also supported for Broadband Subscriber Management if you have enabled Next Gen Services on the MX240, MX480 or MX960 router with the MX-SPC3 card.

You can configure custom application definitions using custom signatures. These definitions enable identification of protocol bundles through deep packet inspection (DPI) for use by interested services in the service chain.

Before you configure custom application signatures, ensure that `jservices-jdpi` is configured on all required interfaces of your MS-MPC, or of your MX-SPC3 services card if you have enabled Next Gen Services on the MX240, MX480, or MX960. To review how to configure the package on your MS-MPC or MX-SPC3 services card:

- For Junos OS Subscriber Aware, see .

- For Junos OS Broadband Subscriber Management, see *Installing Services Packages for Subscriber Management Application-Aware Policy Management*.

To configure one or more custom application signatures:

1. Specify a name for the application.

```
[edit services application-identification]
user@host# edit application application-name
```

For example:

```
[edit services application-identification]
user@host# edit application my:http
```

2. Specify a description for the application.

```
[edit services application-identification application application-name]
user@host# set description description
```

For example:

```
[edit services application-identification application my:http]
user@host# set description "Test application"
```

3. Specify an alternative name for the application.

```
[edit services application-identification application application-name]
user@host# set alt-name alt-name
```

For example:

```
[edit services application-identification application my:http]
user@host# set alt-name my:http-app
```

4. Enable saving of the application system cache (ASC).

```
[edit services application-identification application my:http]
user@host# set cacheable
```

5. Specify the name of the Junos OS release for compatibility.

```
[edit services application-identification application application-name]
user@host# set compatibility junos-compatibility-version
```

For example:

```
[edit services application-identification application my:http]
user@host# set compatibility 17.1
```

6. Specify any desired application tags, consisting of a user-defined name and value.

```
[edit services application-identification application application-name]
user@host# set tags tag-name tag-value
```

For example:

```
[edit services application-identification application my:http]
user@host# set tags traffic-type video-stream
```

7. Specify one or more address-based signatures.

- Specify a destination address and destination port-range.

```
[edit services application-identification application application-name]
user@host# set filter ip 200.0.0.2/24 port-range [80]
```

8. Specify an ICMP-based signature.

a. Specify ICMP type and code.

```
[edit services application-identification application application-name]
user@host# set icmp-mapping type icmp-type code icmp-code
```

For example:

```
[edit services application-identification application my:http]
user@host# set icmp-mapping type 33 code 34
```

9. Specify an IP protocol-based signature.

   a. Specify the IP protocol by protocol number.

```
[edit services application-identification application application-name]
user@host# set ip-protocol-mapping protocol protocol-number
```

   For example:

```
[edit services application-identification application my:http]
user@host# set ip-protocol-mapping protocol 103
```

   All ip-protocol-mappings are allowed except Protocol numbers 1,6,17 are not allowed to be configured under ip-protocol based signatures. If you try to configure protocols 1,6,17 under ip-protocol-mapping you will get commit errors.

10. Specify one or more Layer 4 and Layer 7 signatures using pattern matching in conjunction with a Layer 4 protocol.

    a. Specify a name for the Layer 4 and Layer 7 signature.

```
[edit services application-identification application application-name over protocol-type]
user@host# set signature l4-l7-signature-name
```

   For example:

```
[edit services application-identification application my:http over http]
user@host# set signature myl3l7
```

b. Specify the order to be used if conflicts occur during the application classification. In such a case, the application with lowest order is classified.

```
[edit services application-identification application application-name over protocol-
type signature l4-l7-signature-name member member-name]
user@host# set order order
```

For example:

```
[edit services application-identification application my:http over http signature myl3l7
member m01]
user@host# set order 1
```

c. Specify the priority for using this signature instead of using any matched predefined signatures.

```
[edit services application-identification application application-name over protocol-
type signature l4-l7-signature-name]
user@host# set order-priority (high | low)
```

For example:

```
[edit services application-identification application my:http over http signature myl3l7]
user@host# set order-priority high
```

d. (Optional) Specify the protocol. If you are using Next Gen Services with the MX-SPC3 services card, do not perform this step.

```
[edit services application-identification application application-name over protocol-
type signature l4-l7-signature-name]
user@host# set protocol (http | ssl | tcp | udp)
```

For example:

```
[edit services application-identification application my:http over http signature myl3l7]
user@host# set protocol http
```

e. (Optional) Specify that members are to be matched in order.

```
[edit services application-identification application application-name over protocol-
type signature l4-l7-signature-name]
user@host# set chain-order
```

f. Specify a member. You can repeat this step to define up to four members.

```
[edit services application-identification application application-name over protocol-
type signature l4-l7-signature-name]
user@host# edit member member-name
```

For example:

```
[edit services application-identification application my:http over http signature myl3l7]
user@host# edit member m01
```

g. Specify the member's identifying pattern.

```
[edit services application-identification application application-name over protocol-
type signature l4-l7-signature-name member member-name]
user@host# set pattern pattern
```

For example:

```
[edit services application-identification application my:http over http signature myl3l7
member m01]
user@host# set pattern "www\.facebook\.net"
```

h. Specify the direction of flows to which pattern matching is applied.

```
[edit services application-identification application application-name over protocol-
type signature l4-l7-signature-name member member-name]
user@host# set direction (any | client-to-server | server-to-client)
```

For example:

```
[edit services application-identification application my:http over http signature myl3l7
member m01]
user@host# set direction any
```

i. Specify the number of check-bytes. This option applies to TCP and UDP only.

```
[edit services application-identification application application-name over protocol-
type signature l4-l7-signature-name member member-name]
user@host# set check-bytes max-bytes-to-check
```

For example:

```
[edit services application-identification application my:http over http signature myl3l7
member m01]
user@host# set check-bytes 5000
```

11. (For Next Gen Services with the MX-SPC3 services card only) After you have committed your changes, you can check the status of the custom signature commitment.

```
[edit services application-identification application my:http over http signature myl3l7
member m01]
user@host> show services application-identification commit-status
```

RELATED DOCUMENTATION

Uninstalling a Predefined Junos OS Application Signature Package

(i) **NOTE**: Starting in Junos OS Release 19.3R2 and 19.4R1, application identification is also supported for Broadband Subscriber Management if you have enabled Next Gen Services on the MX240, MX480 or MX960 router with the MX-SPC3 card.

To uninstall the current application signature package:

- Enter the uninstall command.

```
user@host> request service application-identification uninstall
```

RELATED DOCUMENTATION

Downloading and Installing Predefined Junos OS Application Signature Packages | 23

Configuring HTTP Header Enrichment

IN THIS CHAPTER

Junos Web Aware HTTP Header Enrichment Overview

Subscribers accessing Web-based services often need to have content added to the HTTP headers sent back and forth as part of the client-server exchange. You can use Junos Web Aware to configure HTTP header enrichment on the MX Series router. Junos Web Aware allows tag insertions. In addition to the International Mobile Subscriber Identity (IMSI) and mobile station ISDN (MSISDN) tags, you can specify tags for International Mobile Station Equipment Identity (IMEI), TDF gateway IP address, and Subscriber IP address.

Support added in Junos OS Release 20.2R1 for only the insertion of IPv4 or IPv6 tags user addresses in an HTTP headers. No other tags are supported in this release for Next Gen Services.

For example, this feature can add the last line to this sequence of HTTP headers:

```
GET /256k.html HTTP/1.1
Host: 10.45.45.2
Accept */*
Accept-Language: en-us
Accept-Encoding: gzip, deflate
User-Agent: Mozilla/4.0 (compatible; MSIE 6.0; Windows NT 5.1; NET CLR 1.1.4322
name: value
X-MSISDN: <MSISDN #>
```

You can also use HTTP header enrichment to replace a byte of the IPv4 or IPV6 user address in the HTTP header with a value you specify.

Change History Table

Feature support is determined by the platform and release you are using. Use Feature Explorer to determine if a feature is supported on your platform.

Release	Description
20.2R1	Support added in Junos OS Release 20.2R1 for only the insertion of IPv4 or IPv6 tags user addresses in an HTTP headers. No other tags are supported in this release for Next Gen Services.

RELATED DOCUMENTATION

Configuring HTTP Header Enrichment Overview | **40**

hcm

HTTP Content Manager (HCM)

IN THIS SECTION

- Configuring the HTTP-Manager Package on the Router | **35**

HTTP Content Management (HCM) is an application used for inspecting the HTTP traffic transmitted through port 80 (default) or any other port you use to transmit HTTP traffic. HCM can be installed on an MX Series router that is running the corresponding version of the Junos OS release. HCM inspects HTTP traffic even if the default port 80 is not used for HTTP traffic and is interoperable with ms, vms, and ams interface types. It supports fragmented HTTP request packets and GET, PUT, and POST requests.

Support added in Junos OS Release 20.2R1 for only the insertion of IPv4 or IPv6 tags user addresses in an HTTP headers. No other tags are supported in this release for Next Gen Services.

Configuring the HTTP-Manager Package on the Router

1. Before you install the HTTP-Manager package on the router, ensure that you have the appropriate version of the HTTP-Manager package for the Junos OS image you are using on the router. When you have confirmed that you have the right package, use the `request system software add` command to install the HTTP-Manager package. You have to restart the CLI after the package is installed.

```
user@router> request system software add jservices-x86-32-19.4R1.1.tgz
NOTICE: Validating configuration against package-name.
NOTICE: Use the 'no-validate' option to skip this if desired.
Checking compatibility with configuration
```

```
Initializing...
```

```
WARNING: cli has been replaced by an updated version:
CLI release 19.4R1 built by builder on 2020-06-10 02:36:22 UTC
Restart cli using the new version ? [yes,no] (yes)
Restarting cli ...
```

2. When the CLI has restarted, use the **show version** command to see whether the HTTP-Manager packages are installed.

```
user@router> show version
...
HTTP-Manager Management Component [19.4R1-1-A1.2]
HTTP-Manager Dataplane Component [19.4R1-1-A1.2]
user@router>..
```

3. If you want to upgrade the Junos OS image on a router that has the HTTP-Manager package installed, you should first save and then delete the HTTP-Manager configuration from the router.

- To view the HTTP-Manager configuration, use the **user@router>extension juniper-http-manager show <section>** command.

- To delete the HTTP-Manager configuration from the router, use the **user@router>extension juniper-http-manager delete <section>** command.

- Any remnant HTTP-Manager configuration left on the router will be deleted when the Junos OS image is upgraded. So, ensure that you have saved all necessary HTTP Content Management configurations.

- To delete the HTTP-Manager package from the router, use the **user@router> request system software delete <http-manager-package>** command.

- Reinstall the HTTP-Manager package on the router after you upgrade the Junos OS image on the router.

```
user@router> show version
Hostname: router
Model: mx480
JUNOS Base OS boot [19.4R1]
JUNOS Base OS Software Suite [19.4R1]
JUNOS Kernel Software Suite [19.4R1]
JUNOS Crypto Software Suite [19.4R1]
JUNOS Packet Forwarding Engine Support (M/T Common) [19.4R1]
JUNOS Packet Forwarding Engine Support (MX Common) [19.4R1]
JUNOS Online Documentation [19.4R1]
JUNOS Voice Services Container package [19.4R1]
JUNOS Border Gateway Function package [19.4R1]
JUNOS Services AACL Container package [19.4R1]
JUNOS Services LL-PDF Container package [19.4R1]
JUNOS Services PTSP Container package [19.4R1]
JUNOS Services Stateful Firewall [19.4R1]
JUNOS Services NAT [19.4R1]
JUNOS Services Application Level Gateways [19.4R1]
JUNOS Services Captive Portal and Content Delivery Container package [19.4R1]
JUNOS Services RPM [19.4R1]
JUNOS Services HTTP Content Management package [19.4R1]
JUNOS AppId Services [19.4R1]
JUNOS IDP Services [19.4R1]
JUNOS Services Crypto [19.4R1]
JUNOS Services SSL [19.4R1]
JUNOS Services IPSec [19.4R1]
JUNOS Runtime Software Suite [19.4R1]
 JUNOS Routing Software Suite [19.4R1]
```

```
HTTP-Manager Management Component [19.4R1-1-A1.2]
HTTP-Manager Dataplane Component [19.4R1-1-A1.2]
```

```
user@router> configure
Entering configuration mode
```

```
[edit]
user@router# extension juniper-http-manager show
## Last changed: 2020-06-07 13:21:36 PDT
services {
    http-manager {
        traceoptions {
            level all;
            flag all;
        }
```

```
    }
}
```

```
[edit]
user@router# extension juniper-http-manager delete
```

```
[edit]
user@router# extension juniper-http-manager show
```

```
[edit]
 user@router# commit
commit complete
```

```
[edit]
user@router# exit
Exiting configuration mode
```

```
user@router> request system software delete http-manager-services
Removing package 'http-manager-services' ...
Removing /opt/sdk/service-packages/http-manager-services ...
Removing http-manager-services-xlr-19.4R1-1-A1.2.tgz from /var/sw/pkg ...
Notifying mspd ...
```

```
user@router> request system software delete http-manager-mgmt
Removing package 'http-manager-mgmt' ...
Reloading /config/juniper.conf.gz ...
Activating /config/juniper.conf.gz ...
mgd: commit complete
```

```
Restarting mgd ...
Restarting http-manager ...
```

```
WARNING: cli has been replaced by an updated version:
CLI release 11.4R3.7 built by builder on 2020-05-14 19:51:45 UTC
Restart cli using the new version ? [yes,no] (yes)
```

```
Restarting cli ...
user@router>
```

```
user@router> show version
Hostname: router
Model: mx480
JUNOS Base OS boot [19.4R1]
JUNOS Base OS Software Suite [19.4R1]
JUNOS Kernel Software Suite [19.4R1]
JUNOS Crypto Software Suite [19.4R1]
JUNOS Packet Forwarding Engine Support (M/T Common) [19.4R1]
JUNOS Packet Forwarding Engine Support (MX Common) [19.4R1]
JUNOS Online Documentation [19.4R1]
JUNOS Voice Services Container package [19.4R1]
JUNOS Border Gateway Function package [19.4R1]
JUNOS Services AACL Container package [19.4R1]
JUNOS Services LL-PDF Container package [19.4R1]
JUNOS Services PTSP Container package [19.4R1]
JUNOS Services Stateful Firewall [19.4R1]
JUNOS Services NAT [19.4R1]
JUNOS Services Application Level Gateways [19.4R1]
JUNOS Services Captive Portal and Content Delivery Container package [19.4R1]
JUNOS Services RPM [19.4R1]
JUNOS Services HTTP Content Management package [19.4R1]
JUNOS AppId Services [19.4R1]
JUNOS IDP Services [19.4R1]
JUNOS Services Crypto [19.4R1]
JUNOS Services SSL [19.4R1]
JUNOS Services IPSec [19.4R1]
JUNOS Runtime Software Suite [19.4R1]
JUNOS Routing Software Suite [19.4R1]
```

Change History Table

Feature support is determined by the platform and release you are using. Use Feature Explorer to determine if a feature is supported on your platform.

Release	Description
20.2R1	Support added in Junos OS Release 20.2R1 for only the insertion of IPv4 or IPv6 tags user addresses in an HTTP headers. No other tags are supported in this release for Next Gen Services.

RELATED DOCUMENTATION

show services hcm statistics

Configuring HTTP Header Enrichment Overview

You configure HTTP header enrichment by configuring tag rules and an HCM profile that points to specific tag rules. Tag rules identify the HTTP enrichment actions to take when the conditions in the tag rule are matched. For subscriber traffic under static policy control, a tag rule is used if it is included in the HCM profile specified in a PCC rule that the traffic matches. For subscribers under dynamic policy control, a message from the PCRF identifies the configured HCM profile to use for HTTP header enrichment.

Support added in Junos OS Release 20.2R1 for only the insertion of IPv4 or IPv6 tags user addresses in an HTTP headers. No other tags are supported in this release for Next Gen Services.

If you change the configuration of tag rules during an existing subscriber data session, the changes do not impact the existing session. The configuration changes are used by any new subscriber data sessions.

To configure HTTP header enrichment for a subscriber:

1. Configure one or more tag rules to specify the HTTP header enrichment actions.

   See "Configuring Tag Rules" on page 41.
2. Configure an HCM profile and assign tag rules to it.

   See "Configuring HCM Profiles and Assigning Tag Rules" on page 48.
3. (For subscribers under static policy control) Assign the HCM profile to a PCC action profile.

   See "Configuring Policy and Charging Control Action Profiles For Junos OS Subscriber Aware" on page 80.
4. (For subscribers under static policy control) Configure a PCC rule that includes the PCC action profile.

   See "Configuring Policy and Charging Control Rules" on page 83.

5. Enable HTTP enrichment for a subscriber's service set.

See .

Change History Table

Feature support is determined by the platform and release you are using. Use Feature Explorer to determine if a feature is supported on your platform.

Release	Description
20.2R1	Support added in Junos OS Release 20.2R1 for only the insertion of IPv4 or IPv6 tags user addresses in an HTTP headers. No other tags are supported in this release for Next Gen Services.

RELATED DOCUMENTATION

Junos Web Aware HTTP Header Enrichment Overview | **33**

Configuring Tag Rules

Tag rules include one or more `term` statements that identify the HTTP enrichment actions to take when the conditions in the `term` are matched. You must configure at least one `tag` in the `then` clause of a `term`, and you can configure multiple tags.

Terms are evaluated in the order they are configured. If a data packet matches all the criteria in the `from` statement in a `term`, then the actions specified in the `then` statement of the `term` are applied. If the `from` statement does not identify any criteria, then all traffic matches. After a data packet matches a term, further terms are not evaluated. If no terms match, then the HTTP header is not enriched.

To configure a tag rule:

1. Configure the list of tag attributes that may be used in tag rules.

```
[edit services hcm]
user@host# set tag-attribute tag-attr-name
```

The tag attributes currently supported for Adaptive Services are `apn`, `ggsnipv4`, `ggsnipv6`, `imei`, `imsi`, `ipv4addr`, `ipv6addr`, and `msisdn`. To configure multiple tag attributes, include them in square brackets ([ ]). Starting in Junos 20.2R1 IPv4 and IPv6 tags for HTTP Header Enrichment are supported for Next Gen Services on MX240, MX480 and MX960. No other tags are supported for Next Gen Services in this release.

For example:

```
[edit services hcm]
user@host# set tag-attribute [msisdn apn]
```

2. Configure a name for the tag rule.

```
[edit services hcm]
user@host# set tag-rule rule-name
```

For example:

```
[edit services hcm]
user@host# set tag-rule rule1
```

3. Configure a term for the tag rule.

```
[edit services hcm set tag-rule rule-name]
user@host# set term term-number
```

> **ⓘ** **NOTE**: The term argument must have a numeric value.

For example:

```
[edit services hcm set tag-rule rule1]
user@host# set term 1
```

4. (Optional) Specify the prefix that the HTTP request destination IP address must match.

```
[edit services hcm tag-rule rule-name term term-number from]
user@host# set destination-address prefix
```

For example:

```
[edit services hcm tag-rule rule1 term 1 from]
user@host# set destination-address 192.0.2.0/24
```

You can also specify the type of address to match:

```
[edit services hcm tag-rule rule-name term term-number from]
user@host# set destination-address (any-ipv4 | any-ipv6 | any-unicast)
```

You can specify multiple prefixes or address types by including the destination-address statement multiple times.

5. (Optional) Specify an IP address range that the HTTP request destination IP address must match.

```
[edit services hcm tag-rule rule-name term term-number from]
user@host# set destination-address-range low address high address
```

For example:

```
[edit services hcm tag-rule rule1 term 1 from]
user@host# set destination-address-range  low 10.10.10.1 high 10.10.10.255
```

You can specify multiple address ranges by including the destination-address-range statement multiple times.

6. (Optional) Specify the destination prefix list that the HTTP request destination IP address must match. The prefix list must already be defined at the [edit policy-options prefix-list] hierarchy level.

```
[edit services hcm tag-rule rule-name term term-number from]
user@host# set destination-prefix-list prefix-name
```

For example:

```
[edit services hcm tag-rule rule1 term 1 from]
user@host# set destination-prefix-list customer1
```

You can specify multiple prefix lists by including the destination-prefix-list statement multiple times.

7. (Optional) Specify any addresses that you want to exclude from matching the HTTP request destination IP address with the except statement. To exclude addresses, you must also configure addresses that do match in a destination-address, destination-address-range, or destination-prefix-list statement at the [edit services hcm tag-rule rule-name term term-number from] hierarchy level.

For example:

```
[edit services hcm tag-rule rule1 term 1 from]
user@host# set destination-address-range  low 10.10.10.1 high 10.10.10.255
user@host# set destination-address 10.10.10.9/32 except
```

This matches all the addresses in the destination range except 10.10.10.9.

You can use except in the following statements at the [edit services hcm tag-rule *rule-name* term *term-number* from] hierarchy level:

```
destination-address {
    any-ipv4 except;
    any-ipv6 except;
    any-unicast except;
    prefix except;
}
destination-address-range {
    high address low address except;
}
destination-prefix-list {
    prefix-name except;
}
```

8. (Optional) Specify a port range that the HTTP request destination port number must match.

```
[edit services hcm tag-rule rule-name term term-number from]
user@host# set destination-port-range high port-number low port-number
```

You can specify multiple port ranges by including the destination-port-range statement multiple times.

> *(i)* **NOTE**: If you do not specify any ports or port ranges to match, then all ports are matched.

9. (Optional) Specify the HTTP request destination port number that must be matched.

```
[edit services hcm tag-rule rule-name term term-number from]
user@host# set destination-ports value
```

You can specify multiple ports by including the destination-ports statement multiple times.

10. (Optional) Specify that you want to apply all HTTP header enrichment actions specified in the `then` statement of the tag rule to all HTTP requests by not including any matching conditions in the `from` statement. You must include a `from` statement in each `term` of a tag rule.

```
[edit services hcm tag-rule rule-name term term-number ]
user@host# set from
```

For example:

```
[edit services hcm tag-rule rule2 term 1]
user@host# set from
[edit services hcm tag-rule rule2 term 1]
user@host# set then count
```

11. Configure a name for a tag.

```
[edit services hcm tag-rule rule-name term term-number then]
user@host# set tag tag-name
```

For example:

```
[edit services hcm tag-rule rule1 term 1 then]
user@host# set tag msisdn-tag
```

12. Configure the tag header that the tag applies to the HTTP header.

```
[edit services hcm tag-rule rule-name term term-number then tag tag-name]
user@host# set tag-header header
```

For example:

```
[edit services hcm tag-rule rule1 term 1 then tag msisdn-tag]
user@host# set tag-header X_MSISDN
```

You can configure a maximum of 16 unique tag headers.

The *header* values cannot be `accept`, `accept-charset`, `accept-encoding`, `accept-language`, `authorization`, `expect`, `host`, `if-match`, `if-modified-since`, `if-none-match`, `if-range`, `if-unmodified-since`, `max-forwards`, `proxy-`

`authorization`, `referer`, `user-agent`, or `x-moz`. These header values are reserved; you cannot configure them.

13. Specify the tag attribute that the tag applies to the HTTP header. To specify multiple attributes at one time, include the attributes in square brackets ([]).

```
[edit services hcm tag-rule rule-name term term-number then tag tag-name]
user@host# set tag-attribute [tag-attr-name]
```

> ⓘ **NOTE**: The tag attribute must be listed in the tag attributes configured in Step 1.

For example:

```
[edit services hcm tag-rule rule1 term 1 then tag msisdn-tag]
user@host# set tag-attribute msisdn
```

14. Specify the separator that the tag uses in the HTTP header.

```
[edit services hcm tag-rule rule-name term term-number then tag tag-name]
user@host# set tag-separator separator
```

For example:

```
[edit services hcm tag-rule rule1 term 1 then tag msisdn-tag]
user@host# set tag-separator /
```

15. (Optional) Specify a hash method and prefix key for the insertion of the tag in the HTTP header.

```
[edit services hcm tag-rule rule-name term term-number then tag tag-name encrypt]
user@host# set hash algorithm prefix hash-prefix
```

Currently, only the `md5` hash method is supported.

For example:

```
[edit services hcm tag-rule rule1 term 1 then tag msisdn-tag encrypt]
user@host# set hash md5 prefix gatewaykey1
```

16. (Optional) Enable the collection of statistics for HTTP header enrichment for the tag rule.

```
[edit services hcm tag-rule rule-name term term-number then
user@host# set count
```

17. (Optional) Configure how the tag replaces a byte of the IPv4 or IPv6 user address with a different value in the HTTP header.

```
[edit services hcm tag-rule rule-name term term-number then tag tag-name]
user@host# set (ipv4-mask ipv4-mask | ipv6-mask ipv6-mask) (ipv4-or-value ipv4-or-value |
ipv6-or-value ipv6-or-value)
```

To identify the byte you want to replace, enter **255** for IPv4 or **ff** for IPv6 in the corresponding byte of the `ipv4-mask` or `ipv6-mask` and enter zero in the other bytes.

To specify the new value for that byte, enter the value in the corresponding byte of the `ipv4-or-value` or the `ipv6-or-value` and enter zero in the other bytes.

For example, the following replaces the first byte of the IPv4 user address with the value 168:

```
[edit services hcm tag-rule tag1 term term1 then tag subscip4]
user@host# set ipv4-mask 255.0.0.0 ipv4-or-value 168.0.0.0
```

18. If you want to configure more tags for the `then` statement in the term, repeat Step 11 through Step 17.

19. If you want to configure another `term` statement for the tag rule, repeat Step 3 through Step 18.

Change History Table

Feature support is determined by the platform and release you are using. Use Feature Explorer to determine if a feature is supported on your platform.

Release	Description
20.2R1	Starting in Junos 20.2R1 IPv4 and IPv6 tags for HTTP Header Enrichment are supported for Next Gen Services on MX240, MX480 and MX960. No other tags are supported for Next Gen Services in this release.

RELATED DOCUMENTATION

Configuring HTTP Header Enrichment Overview | 40

Configuring HCM Profiles and Assigning Tag Rules

The HCM profile for a subscriber specifies the tag rules to apply to a subscriber's traffic. Tag rules identify the HTTP enrichment actions to take when the conditions in the tag rule are matched. You can have a maximum of 100 HCM profiles.

Support added in Junos OS Release 20.2R1 for only the insertion of IPv4 or IPv6 tags user addresses in an HTTP headers. No other tags are supported in this release for Next Gen Services.

For subscriber-aware traffic under static policy control, a tag rule is used if it is included in the HCM profile specified in a PCC rule that the traffic matches. For subscriber-aware traffic under dynamic policy control, a message from the PCRF identifies the configured HCM profile and tag rules to use for HTTP header enrichment.

To configure an HCM profile:

1. Configure the HCM profile name.

```
[edit services hcm]
user@host# set profile profile-name
```

For example:

```
[edit services hcm]
user@host# set profile hcm1
```

2. Assign a tag rule to the HCM profile.

```
[edit services hcm profile profile-name]
user@host# set tag-rule rule-name
```

For example:

```
[edit services hcm profile hcm1]
user@host# set tag-rule rule1
```

Change History Table

Feature support is determined by the platform and release you are using. Use Feature Explorer to determine if a feature is supported on your platform.

Release	Description
20.2R1	Support added in Junos OS Release 20.2R1 for only the insertion of IPv4 or IPv6 tags user addresses in an HTTP headers. No other tags are supported in this release for Next Gen Services.

RELATED DOCUMENTATION

Configuring HTTP Header Enrichment Overview | **40**

Configuring Tag Rules | **41**

Junos Web Aware HTTP Header Enrichment Overview | **33**

Configuring Policy and Charging Control Action Profiles For Junos OS Subscriber Aware | **80**

Configuring Policy and Charging Enforcement

IN THIS CHAPTER

Understanding Junos Subscriber Aware Policy and Charging Enforcement Function (PCEF)

IN THIS SECTION

The policy and charging enforcement function (PCEF) of Junos Subscriber Aware enforces policy and charging control (PCC) rules for the treatment of a subscriber's packets. A PCC rule is installed on, and enforced by, the PCEF. The PCC rules can be under static control, under dynamic control of the policy and charging rules function (PCRF), or under activation/deactivation control of a RADIUS server, depending on the PCEF profile that is assigned to a subscriber.

Static Policy Control

For static policies, the PCEF enforces PCC rules you predefined on the MX Series router with no interaction from the PCRF or a RADIUS server, as shown in .

Figure 3: Static Policy Control
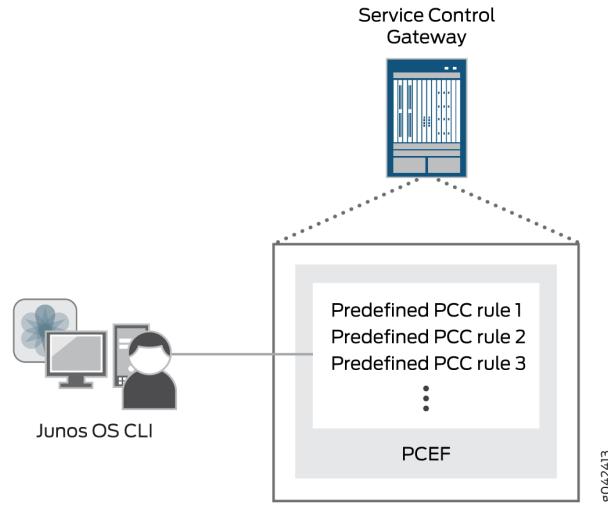
Dynamic Policy Control

For dynamic policies, the PCEF acts upon messages received from the PCRF. The PCRF is the central entity that makes policy and charging decisions based on input from different sources, such as mobile operator configuration, user subscription information, and services information. The PCC rules are either provisioned by the PCRF and sent to the PCEF over the Gx interface using Diameter AVPs, or predefined on the MX Series router and activated by a Diameter message from the PCRF. The PCEF also provides the PCRF with subscriber and access information. See Figure 4 on page 53.

When PCC rules are under dynamic control, the PCEF gives precedence to rules sent by the PCRF over rules that are predefined on the PCEF.

Figure 4: Dynamic Policy Control
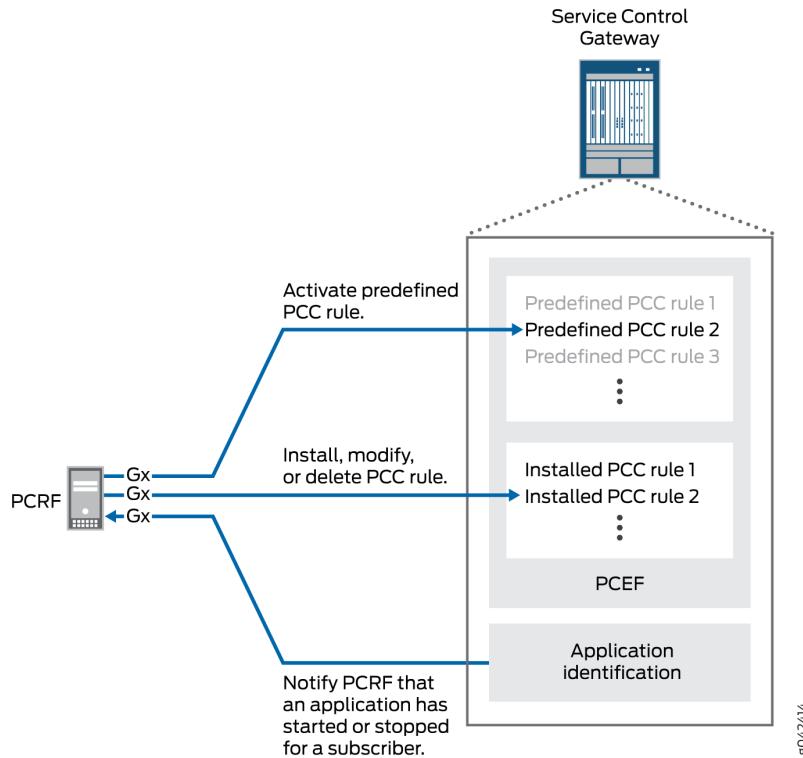
RADIUS Server Policy Control

For policies under control of a RADIUS server, a RADIUS server activates and deactivates policy and PCC rules that you have predefined on the MX Series router, as shown in .

A PCEF profile with RADIUS server control requires an AAA profile, which provides the policy control attributes for RADIUS servers. Subscription-Id-Data in CCR is sourced from the RADIUS server and Calling-Station-Id is received with RADIUS requests.

Usage monitoring is supported through the Third Generation Partnership Project (3GPP) with Gx profile for subscriber services using the `dynamic-profile` configuration.

Figure 5: RADIUS Server Policy Control
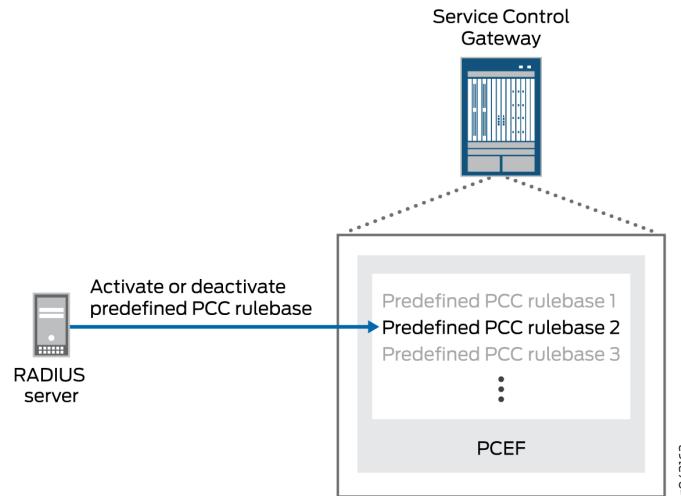
Understanding Predefined Policy and Charging Control Rules for Subscriber-Aware Traffic Treatment

IN THIS SECTION

- Understanding Service Data Flow Filters | **55**
- Understanding Application Filters | **55**
- Understanding PCC Action Profiles | **56**

You can configure policy and charging control (PCC) rules to define the treatment to apply to specific service data flows or to packets associated with specific applications. A PCC rule is applicable to a subscriber's traffic if the rule is in the subscriber's PCEF profile.

These predefined PCC rules contain a `from` clause that identifies the service data flows or applications, and a `then` clause that specifies the PCC action profile that identifies the treatment to apply.

A predefined PCC rule can be used in three ways:

- When PCC rules are under static control, predefined rules are the only rules used. The provisioning of PCC rules involves no interaction from the policy and charging rules function (PCRF) or a RADIUS server.

- When PCC rules are under dynamic control, a predefined PCC rule must be activated by the PCRF. (With dynamic control, PCC rules can also be sent from the PCRF.)

- When PCC rules are under RADIUS server control, a predefined PCC rule must be activated by the RADIUS server.

This topic includes the following sections:

Understanding Service Data Flow Filters

Service data flow (SDF) filters (flow identifiers) are specified in the `from` clause of a PCC rule to identify IP packets belonging to a particular Layer 3 or Layer 4 service data flow. If the IP packet matches the SDF filter in a PCC rule, the treatment specified in the PCC action profile in the `then` clause of the rule is applied.

To configure Layer 3 or Layer 4 SDF filters, you specify one or more of the following parameters:

- Source IP address

- Destination IP address

- Source port

- Destination port

- Layer 4 protocol (UDP or TCP)

Understanding Application Filters

Applications or application groups are specified in the `from` clause of a PCC rule to identify IP packets belonging to a specific application. If the IP packet is for an application identified in a PCC rule, the treatment specified in the PCC action profile in the `then` clause of the rule is applied.

To configure application-aware PCC rules, you can specify one or more of the following parameters:

- application—Specifies the name of an application. This can be a Layer 7 protocol (for example, HTTP) or a particular application running on a Layer 7 protocol, such as Facebook and Yahoo Messenger.

- application-group—Specifies the name of an application group, which can be used to process a number of applications or subgroups at the same time.

> **NOTE**: Application-aware PCC rules that reference specified applications can include wildcard or specific Layer-3 SDF filters, Layer-4 SDF filters, or both.

Understanding PCC Action Profiles

A PCC rule configuration includes an action profile in the `then` clause that defines the treatment to apply to a service data flow or to a packet belonging to an application identified in the `from` clause of the rule. You can configure a PCC action profile that is used in one or more PCC rules to provide the following functionality:

- HTTP redirection—Specifies HTTP redirection to a URL. You can use this action only for PCC rules that match only HTTP-based applications and all flows.

- HTTP Steering path—Specifies an IPv4 or IPv6 address for steering HTTP packets. You can use this action only for PCC rules that match only HTTP-based applications and all flows.

> **NOTE**: A single PCC rule can support either HTTP redirection or HTTP steering path, but not both.

- Steering with a routing instance—Specifies a routing instance for steering of packets to a third-party server to apply services or to a local or external service chain. You can configure different routing instances for traffic from the subscriber (uplink) and traffic to the subscriber (downlink).

- Keep existing steering—Specifies that steering attributes configured in a PCC action profile that a PCC rule applies to a data flow session when it begins will continue to be applied to the data flow when the PCC rule match conditions are modified, deleted, or added to.

- Forwarding class—Specifies the forwarding class that you want assigned to the packet.

- Maximum bit rate—Specifies the maximum bit rate for uplink and for downlink traffic.

- HCM profile—Specifies the profile that identifies the HTTP header enrichment rules to apply. You can use this action only for PCC rules that match only HTTP-based applications and all flows.

- Gating status—Specifies whether to block or to forward IP packets.

Understanding How Subscriber-Aware Policy and Charging Control Rules Are Provisioned Dynamically by a PCRF

IN THIS SECTION

- Policy Decisions | **57**
- Supported Operations | **58**
- Methods for Provisioning PCC Rules | **59**

With dynamic policy control, the policy and charging rules function (PCRF) controls the provisioning of policy and charging control (PCC) rules on the Junos Subscriber Aware PCEF for a subscriber. Dynamic policy control is enabled when a dynamic-policy-control policy and charging enforcement function (PCEF) profile is assigned to a subscriber. Dynamic policy control requires Junos Policy Control.

This topic includes the following sections:

Policy Decisions

The PCRF is central in making policy and charging control decisions and can install, activate, modify, or deactivate a PCC rule on the PCEF at any time. The PCRF can make its policy and charging control decisions based on different sources, including:

- Subscription information for the user equipment that is received from the subscription profile repository (SPR)

- Operator configuration in the PCRF

- Information from the access network about the access technology

- Information from the PCEF, such as the name of the application that the subscriber is using

The Gx interface is used to send PCC rule provisioning information from the PCRF to the PCEF, and to provide notification of traffic-plane events from the PCEF to the PCRF.

Supported Operations

Junos Subscriber Aware and Junos Policy Control support the following operations with the PCRF:

- Install or modify rules—The PCRF sends the `Charging-Rule-Install` AVP to install a PCC rule that is not already installed or modify an existing rule on the PCEF.

- Remove rules—The PCRF sends the `Charging-Rule-Remove` AVP to remove a PCC rule that is already installed.

- Activate rules—The PCRF sends the `Rule-Activation-Time` AVP to indicate the time at which to activate the rule, and it is contained within the `Charging-Rule-Install` AVP. This operation results in a single activation of the rule, not a recurring activation schedule.

- Deactivate rules—The PCRF sends the `Rule-Deactivation-Time` AVP to indicate the time at which to deactivate the rule, and it is contained within the `Charging-Rule-Install` AVP. This operation results in a single deactivation of the rule, not a recurring deactivation schedule.

- PCEF session revalidation—The PCRF sends the `Revalidation-Time` AVP along with the `Event-Trigger` AVP with the value `REVALIDATION_TIMEOUT` to indicate the time at which the PCEF must request PCEF session revalidation from the PCRF. When the specified time is reached, the PCEF sends an event trigger with the value `REVALIDATION_TIMEOUT` to request PCEF session revalidation.

- Report application start or stop—The PCEF sends an event trigger when it detects the start or stop of an application.

The containers for the PCC rules are named `Charging-Rule-Definition`. Multiple `Charging-Rule-Definition` containers can be sent within a `Charging-Rule-Install` or `Charging-Rule-Remove`, each of which is applied per subscriber.

If a time zone is configured on the router, the activation and deactivation settings apply to the configured time zone and are adjusted for transitions to and from daylight saving time.

Methods for Provisioning PCC Rules

The PCRF uses one of the following procedures to specify the PCC rules that the PCEF applies:

- Pull mode during TDF subscriber creation—Applies when the MX Series gateway receives a request for a new TDF subscriber. The PCEF sends a credit control request initial (CCR-I) message to the PCRF with information about the subscriber. The PCRF downloads PCC rules to the PCEF in a credit control answer initial (CCA-I) message, which may also include any activation and deactivation times that apply to the rules and the time at which the PCEF must re-request PCC rules from the PCRF. Figure 6 on page 59 shows the message flow for a pull procedure during TDF subscriber creation.

Figure 6: Message Flow for Pull Mode During TDF Subscriber Creation
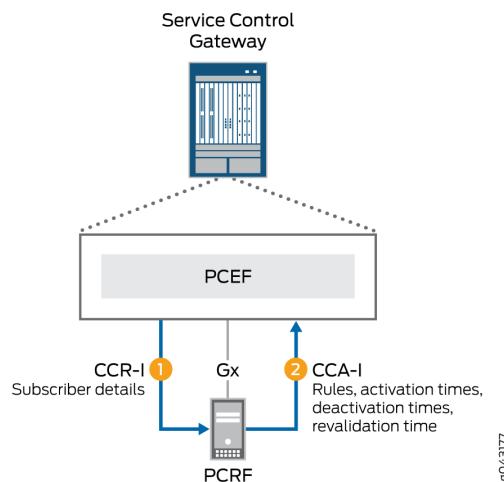
- Pull mode after PCEF event trigger—Applies when the PCEF sends an event trigger to the PCRF. This can occur when the MX Series router detects a new application start or stop or when the revalidation time has occurred. The PCEF sends a credit control request update (CCR-U) message along with the appropriate event trigger to the PCRF. The PCRF might download new rules to the PCEF in a credit control answer update (CCA-U) message, which may also include any activation and deactivation times that apply to the rules and the time at which the PCEF must re-request PCC rules from the PCRF. Figure 7 on page 60 shows the message flow for a pull procedure after a PCEF event trigger.

Figure 7: Message Flow for Pull Mode After PCEF Event Trigger
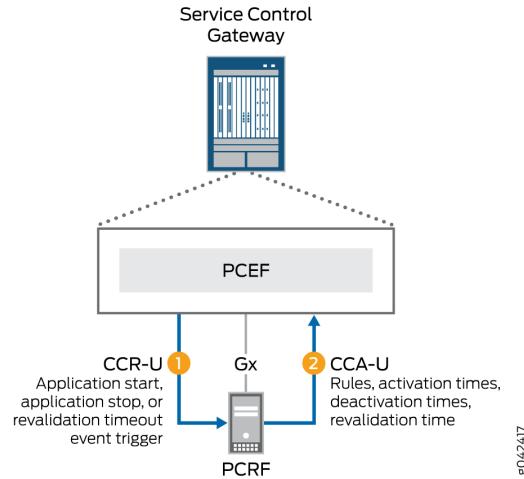
- Push mode—Applies when the PCRF provisions PCC rules without obtaining a request from the PCEF. The PCRF sends the PCC rules in a re-authorization request (RAR) to the PCEF based on information sent to the PCRF through the Rx interface or in response to a trigger within the PCRF. The RAR may also include any activation and deactivation times that apply to the rules and the time at which the PCEF must re-request PCC rules from the PCRF. The PCRF includes these PCC rules in an RAR message because the PCC rules were not requested by the PCEF, and no credit control request (CCR) or credit control answer (CCA) messages are triggered by the RAR. The PCEF responds with a re-authorization answer (RAA) message. Figure 8 on page 61 shows the message flow for a push procedure.

Figure 8: Message Flow for Push Mode
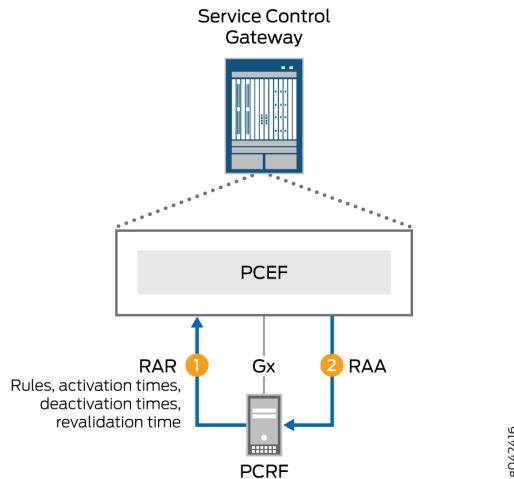
RELATED DOCUMENTATION

Understanding Junos Subscriber Aware Policy and Charging Enforcement Function (PCEF) | **51**

Understanding Predefined Policy and Charging Control Rules for Subscriber-Aware Traffic Treatment | **54**

Configuring Dynamic Policy Control by PCRF | **73**

Understanding How Subscriber-Aware Policy and Charging Control Rules Are Provisioned Statically

Static policy control is enabled when a static-policy-control policy and charging enforcement function (PCEF) profile is assigned to a subscriber. The policy and charging control (PCC) rules that you configure on the MX Series router and assign to the PCEF profile are active, and are *not* controlled by the policy and charging rules function (PCRF) or RADIUS server.

RELATED DOCUMENTATION

Understanding Junos Subscriber Aware Policy and Charging Enforcement Function (PCEF) | **51**

Understanding Predefined Policy and Charging Control Rules for Subscriber-Aware Traffic Treatment | **54**

Configuring Static Policy Control | **74**

Understanding How a RADIUS Server Controls Policy and Charging Control Rules

Policy control by a RADIUS server takes place when an aaa-policy-control policy and charging enforcement function (PCEF) profile is assigned to a subscriber. A RADIUS server activates and deactivates policy and charging control (PCC) rules that you have configured on the MX Series router and assigned to the PCEF profile. A network element, which is a load-balanced group of RADIUS servers, is assigned to the subscriber.

This topic includes the following sections:

Rule Activation When TDF Session Begins

When the traffic detection function (TDF) subscriber session begins, the Junos Subscriber Aware PCEF sends an access request to the RADIUS server. This is shown in Figure 9 on page 63. This access request includes the subscriber username, IP address, and other relevant AVP information that Subscriber Aware received from the broadband network gateway or Packet Data Network Gateway during the subscriber session setup.

The RADIUS server responds to the PCEF with an access-accept message, which contains the names of the rulebases to activate. You can configure the AVP that carries the name of a rulebase to be activated; by default the PCEF looks for a rulebase name in the ERX-Service-Activate Juniper vendor-specific attributes (VSA).

Figure 9: RADIUS Server Message Flow When TDF Session Begins
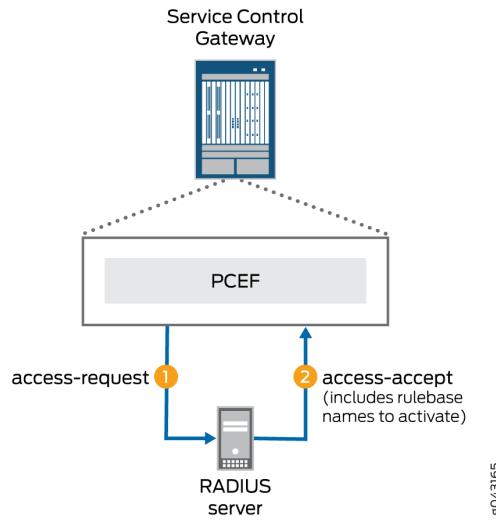
Rule Activation and Deactivation When RADIUS Server Sends Request

The RADIUS server can initiate the activation or deactivation of rulebases by sending a change of authorization (CoA) request to the PCEF, as shown in . You can configure the AVP that carries the name of a rulebase to be activated; by default the PCEF looks for a rulebase name in the ERX-Service-Activate Juniper VSA. You can also configure the AVP that carries the name of a rulebase to be deactivated; by default the PCEF looks for a rulebase name in the ERX-Service-Deactivate Juniper VSA.

The PCEF responds to the CoA request by sending a CoA Ack to the RADIUS server.

Figure 10: Message Flow When RADIUS Server Sends Request
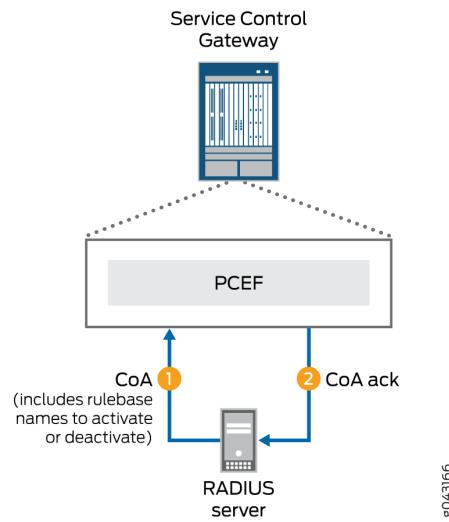
Supported Attributes in RADIUS Messages

The following tables list the RADIUS attributes, 3GPP VSAs, and Juniper Networks VSAs that are supported in the RADIUS messages between the MX Series router and a RADIUS server.

Table 2 on page 64 lists the RADIUS attributes and 3GPP VSAs that are supported in the access-request messages sent to the RADIUS server.

Table 2: Attributes Supported in Access-Request Messages

Attribute Number	Attribute Name	Description	Content
1	User-Name	Username for the TDF subscriber if it is provided in the RADIUS accounting request received from the Packet Data Network Gateway (PGW) or broadband network gateway (BNG). This is a RADIUS IETF attribute.	String
2	User-Password	User password configured in the subscriber's PCEF profile. This is a RADIUS IETF attribute.	String

Table 2: Attributes Supported in Access-Request Messages *(Continued)*

Attribute Number	Attribute Name	Description	Content
4	NAS-IP-Address	IPv4 address of the MX Series router for communication with the RADIUS server. This is a RADIUS IETF attribute.	IPv4 address
8	Framed-IP-Address	IPv4 address for the TDF subscriber if it is provided in the RADIUS accounting request received from the PGW or BNG. This is a RADIUS IETF attribute.	IPv4 address
31	Calling-Station-ID	Identifier for the mobile station of the TDF subscriber if it is provided in the RADIUS accounting request received from the PGW or BNG. This is a RADIUS IETF attribute.	MSISDN in international format, UTF-8 encoded decimal characters
32	NAS-Identifier	Identifier of the NAS originating the request. This is a RADIUS IETF attribute.	String
44	Acct-Session-ID	User Session identifier generated by Subscriber Aware for the TDF subscriber. This is a RADIUS IETF attribute.	UTF-8 encoded string
97	Framed-IPv6-Prefix	IPv6 prefix for the TDF subscriber if it is provided in the RADIUS accounting request received from the PGW or BNG. This is a RADIUS IETF attribute.	Value indicating the prefix, as specified in RFC 3162

Table 2: Attributes Supported in Access-Request Messages *(Continued)*

Attribute Number	Attribute Name	Description	Content
26/10415/1 (3GPP type 1)	3GPP-IMSI	IMSI for the TDF subscriber if it is provided in the RADIUS accounting request received from the PGW or BNG. This is a 3GPP VSA.	UTF-8 encoded string

lists the VSAs that are supported in the Access-Accept messages sent from the RADIUS server to the PCEF.

Table 3: Attributes Supported in Access-Accept Messages

Attribute Number	Attribute Name	Description	Content
26-65	ERX-Service-Activate	Specifies a PCC rulebase to activate for the subscriber. Tagged VSA, which supports 8 tags (1-8). This is a Juniper Networks VSA and is the default VSA for carrying rulebase activations; you can also specify a different AVP code and vendor ID.	string: *rulebase-name*

lists the VSAs that are supported in the CoA messages sent from the RADIUS server to the PCEF.

Table 4: Attributes Supported in CoA Messages

Attribute Number	Attribute Name	Description	Content
26-65	ERX-Service-Activate	Specifies a PCC rulebase to activate for the subscriber. Tagged VSA, which supports 8 tags (1-8). This is a Juniper Networks VSA and is the default VSA for carrying rulebase activations; you can also specify a different AVP code and vendor ID.	string: *rulebase-name*
26-66	ERX-Service-Deactivate	Specifies a PCC rulebase to deactivate for the subscriber. This is a Juniper Networks VSA and is the default VSA for carrying rulebase deactivations; you can also specify a different AVP code and vendor ID.	string: *rulebase-name*

RELATED DOCUMENTATION

Understanding PCEF Profiles

A policy and charging enforcement function (PCEF) profile defines whether policy and charging control (PCC) rules for a subscriber are under static control, under dynamic control of the policy and charging rules function, or under activation/deactivation control of a RADIUS server by using the `static-policy-control`, `dynamic-policy-control`, or `aaa-policy-control` statement, respectively, in the PCEF profile configuration. The PCEF profile also identifies the predefined PCC rules and rulebases that the subscriber can use, and assigns a precedence value to each predefined rule. A subscriber is assigned a

PCEF profile during the TDF subscriber session setup. See "Understanding Selection of Policy-Control Properties for an IP-based TDF Subscriber" on page 108.

A PCEF profile with dynamic policy control requires a Diameter Gx profile, which provides network access information for the Diameter application.

A PCEF profile with RADIUS server control requires an AAA profile, which provides the policy control attributes for RADIUS servers

RELATED DOCUMENTATION

Understanding Network Elements

IN THIS SECTION

- Load Balancing Within Network Elements | 69
- Server Priority | 69
- Dead Server Detection | 69
- Maximum Pending Requests for a Network Element | 69

A network element is a load-balanced group of RADIUS servers providing policy management for TDF subscribers.

Network elements are specified in the AAA profile that is applied to a policy and charging enforcement function (PCEF) profile. A subscriber is assigned to a PCEF profile.

Load Balancing Within Network Elements

The Junos Subscriber Aware PCEF distributes requests to RADIUS servers across the servers in the network element.

Server Priority

Within a network element, a RADIUS server can be assigned a priority of 1 through 16, with 1 being the highest priority. You can have multiple servers with the same priority in a network element. All access requests are load balanced among the highest priority servers. If all the servers with the highest priority in the network element fail, then requests are load balanced among servers with the next highest priority level.

Dead Server Detection

To determine whether a RADIUS server in a network element has failed, the PCEF keeps track of how often requests sent to a server time out and must be retransmitted. If the number of times that requests need to be retransmitted reaches a configured limit within a configured time interval, PCEF marks the server as `dead` and starts sending requests to the next available server in the network element with the same priority.

At the same time, the PCEF starts a timer for the RADIUS server. After this timer expires, the PCEF marks the dead server as alive again, and includes it in the rotation for sending RADIUS messages.

Maximum Pending Requests for a Network Element

You can configure the maximum number of requests that can be queued to the network element. When the pending request queue is full, any additional requests are dropped.

You can also configure a high and a low watermark that are percentages of the maximum number of requests that can be queued. If the number of pending requests reaches this high watermark, a **flow control on** message is generated. When the number of pending requests then falls below the low watermark, a **flow control off** message is generated.

RELATED DOCUMENTATION

Understanding AAA Profiles

An AAA profile is a collection of attributes to specify how the Junos Subscriber Aware PCEF interacts with RADIUS servers that control the activation and deactivation of policy and charging control (PCC) rules. An AAA profile is assigned to a subscriber's policy and charging enforcement function (PCEF) profile, which specifies the PCC rulebases for the subscriber.

Network Elements

In the AAA profile, you specify a network element (load-balanced RADIUS server group) to be used for authorization of policy control. If the RADIUS servers in a Network Element cannot initiate a change of authorization (CoA) request without an accounting record, then the AAA profile must specify the network element for accounting as well as for authorization, and the AAA profile must enable CoA accounting.

RADIUS Attributes That Carry Rulebase Names for Activation and Deactivation

You can specify the RADIUS AVPs that carry the PCC rulebase names for activation or deactivation. By default, the PCC rulebase name for activations is carried in the ERX-Service-Activate Juniper vendor-specific attributes (VSA). By default, the PCC rulebase name for deactivations is carried in the ERX-Service-Deactivate Juniper VSA.

RELATED DOCUMENTATION

Understanding Static Time-of-Day PCC Rule Activation and Deactivation

With static time-of-day policy and charging control (PCC) rules activation and deactivation, you can specify a schedule for activating and deactivating PCC rules or rulebases within a static PCEF profile. The rule or rulebase activation and deactivation settings take effect for subscribers assigned to that static PCEF profile.

The activation and deactivation settings can consist of the time of day, the day, and the month of the year. The day can be expressed as a day of the week, as a numbered day of the month, or as the last day of the current month. If a day is not specified, then the rule activation and deactivation occurs daily at the specified times. If you configure a day of the month, you can also configure a month of the year.

If a day is not specified and the deactivation time of day setting is earlier than the activation time of day setting, then a rule is deactivated the day after it is activated.

If a time zone is configured on the router, the time-of-day settings apply to the configured time zone and are adjusted for transitions to and from daylight saving time.

You cannot use static time-of-day settings for dynamic PCC rules.

RELATED DOCUMENTATION

Understanding Usage Monitoring for TDF Subscribers

IN THIS SECTION

For TDF subscribers that are assigned to a dynamic policy and charging enforcement function (PCEF) profile, you can monitor the subscriber use during a session as a volume of traffic, an amount of time, or both, and send reports to the policy and charging rules function (PCRF) when a threshold is exceeded or when the PCRF requests a report. Data volume and the amount of time used can be tracked for individual or multiple data flows or applications that appear in specific policy and charging control (PCC) rules, or for the entire subscriber session.

This topic includes the following sections:

Tracked Resource Identification

Data usage for a subscriber session is tracked through an object called a monitoring key, which the PCRF configures. Traffic for a particular data flow, application, or combination of data flows and applications can be tracked as a data set by assigning a monitoring key to the PCC rules that identify those flows or applications. For predefined PCC rules, you specify the monitoring key with the PCC rule's action profile. For dynamic PCC rules, the PCRF specifies the monitoring key for a rule.

Data usage can also be tracked for the entire TDF subscriber session by configuring the monitoring key level as SESSION.

Threshold Configuration

The PCRF specifies a threshold for reporting data usage when it configures a monitoring key. The threshold can be a combination of uplink volume, downlink volume, total volume, and time used. The MX Series router reports the usage information to the PCRF when this limit is exceeded, and resets the volume to zero.

Messages and AVPs That Are Used

The PCRF must first request usage monitoring by sending the Event-Trigger AVP with the value USAGE_REPORT. This request can be sent to the MX Series router in a CCA-I, CCA-U, or RAR message.

The PCRF configures a monitoring key by sending a Usage Monitoring Information (UMI) AVP that includes the following in a CCA-I, CCA-U, or RAR message to the MX Series router:

- Monitoring-key AVP, which is the identifier.

- Granted-Service-Unit AVP, which specifies the volume threshold, time threshold, or both.

- Usage-Monitoring-Level AVP, which indicates whether the monitoring key applies to the entire subscriber session or to particular PCC/ePCC rules.

The PCRF requests usage monitoring for traffic that matches a PCC rule's data flows or applications by sending the following in a CCA-I, CCA-U, or RAR message to the MX Series router:

- Charging-Rule-Definition AVP, which identifies the rule.

- UMI AVP that includes the Monitoring-key AVP, which identifies the monitoring key to which the rule is associated.

The MX Series router reports usage to the PCRF by sending a UMI AVP that includes the following in a CCR-U message:

- Monitoring-key AVP, which is the identifier.

- Used-Service-Unit AVP, which gives a combination of uplink volume, downlink volume, total volume, and time used.

The PCRF can request a usage report, regardless of whether the threshold is reached, by sending a UMI AVP that includes the following in a CCA-U or RAR message:

- Monitoring-key AVP, which is the identifier.

- Usage-Monitoring-Report AVP, which is set to the value USAGE_MONITORING_REPORT_REQUIRED (0).

The PCRF requests that usage monitoring be disabled for a monitoring key by sending a UMI AVP that includes the following in a CCA-U or RAR message:

- Monitoring-key AVP, which is the identifier.

- Usage-Monitoring-Support, which is set to the value USAGE_MONITORING_DISABLED (0).

RELATED DOCUMENTATION

Configuring Dynamic Policy Control by PCRF

You can configure policy management that is dynamically controlled by the policy and charging rules function (PCRF), which can both provision policy and charging control (PCC) rules on the MX Series router and activate PCC rules that are predefined on the MX Series router.

To configure policy management that is dynamically controlled by a PCRF:

1. (Optional) Configure any flow identifiers to be used in PCC rules.

See "Configuring Service Data Flow Filters" on page 76.

2. (Optional) Configure any custom applications to be used in PCC rules.

   See "Configuring Custom Application Signatures" on page 25.

3. (Optional) Configure the PCC action profiles to be used in PCC rules.

   See "Configuring Policy and Charging Control Action Profiles For Junos OS Subscriber Aware" on page 80

4. (Optional) Configure PCC rules.

   See "Configuring Policy and Charging Control Rules" on page 83.

5. (Optional) Configure PCC rulebases.

   See "Configuring a Policy and Charging Control Rulebase" on page 86.

6. Configure a Diameter Gx profile.

   See "Configuring Diameter Profiles" on page 150.

7. Configure a dynamic PCEF profile.

   See "Configuring a Policy and Charging Enforcement Function Profile for Junos OS Subscriber Aware Dynamic Policies" on page 94

8. (Optional) Configure an NTP server if you want the PCRF to send activation, deactivation, or revalidation times.

   See "Configuring the NTP Server" on page 99.

RELATED DOCUMENTATION

Understanding Junos Subscriber Aware Policy and Charging Enforcement Function (PCEF) | **51**

Understanding How Subscriber-Aware Policy and Charging Control Rules Are Provisioned Dynamically by a PCRF | **57**

Configuring Static Policy Control

You can configure static policy management that is controlled entirely by predefined policy and charging control (PCC) rules that you have configured on the MX Series router.

To configure static policy control:

1. Configure any flow identifiers to be used in PCC rules.

   See "Configuring Service Data Flow Filters" on page 76.

2. Configure any custom applications to be used in PCC rules.

   See "Configuring Custom Application Signatures" on page 25.

3. Configure the PCC action profiles to be used in PCC rules.

See "Configuring Policy and Charging Control Action Profiles For Junos OS Subscriber Aware" on page 80

4. Configure PCC rules.

See "Configuring Policy and Charging Control Rules" on page 83.

5. (Optional) Configure PCC rulebases.

See "Configuring a Policy and Charging Control Rulebase" on page 86.

6. Configure a policy and charging enforcement function (PCEF) profile for static policy control.

See "Configuring a Policy and Charging Enforcement Function Profile for Junos OS Subscriber Aware Static Policies" on page 96

RELATED DOCUMENTATION

Understanding Junos Subscriber Aware Policy and Charging Enforcement Function (PCEF) | **51**

Understanding How Subscriber-Aware Policy and Charging Control Rules Are Provisioned Statically | **61**

Configuring Policy Control by RADIUS Servers

You can configure policy management that is controlled by RADIUS servers. A RADIUS server activates and deactivates policy and charging control (PCC) rules that have been configured on the MX Series router.

To configure policy management that is controlled by RADIUS servers:

1. Configure any flow identifiers to be used in PCC rules.

   See "Configuring Service Data Flow Filters" on page 76.

2. Configure any custom applications to be used in PCC rules.

   See "Configuring Custom Application Signatures" on page 25.

3. Configure the PCC action profiles to be used in PCC rules.

   See "Configuring Policy and Charging Control Action Profiles For Junos OS Subscriber Aware" on page 80

4. Configure PCC rules.

   See "Configuring Policy and Charging Control Rules" on page 83.

5. Configure PCC rulebases.

   See "Configuring a Policy and Charging Control Rulebase" on page 86.

6. Configure RADIUS servers.

   See "Configuring RADIUS Servers" on page 88.

7. Configure RADIUS network elements.

   See "Configuring RADIUS Network Elements" on page 91.

8. Configure an AAA profile.

   See "Configuring an AAA Profile" on page 92.

9. Configure a policy and charging enforcement function (PCEF) profile for policy control by a RADIUS server.

   See "Configuring a Policy and Charging Enforcement Function Profile for Junos OS Subscriber Aware Policies That a RADIUS Server Controls" on page 97

RELATED DOCUMENTATION

Configuring Service Data Flow Filters

> **(i)** **NOTE**: Starting in Junos OS Release 19.3R2, PCC rules are also supported if you have enabled Next Gen Services on the MX240, MX480 or MX960 router with the MX-SPC3 card.

A service data flow (SDF) filter is specified as a matching condition in the `from` clause of a policy and charging control (PCC) rule. Each SDF filter can have one or more flows associated with it; each flow is a five-tuple match.

> **(i)** **NOTE**: If you configure an SDF filter without specifying a remote address, port, port range, or protocol , then the SDF filter matches IP packets that have any value configured for the corresponding attribute. If you configure an SDF filter, you must configure at least one of the following attributes: direction, local port or local port range, protocol, remote address, or remote port or remote port range.

You can configure SDF filters for Junos OS Subscriber Aware or for Junos OS Broadband Subscriber Management, but you use a different CLI hierarchy level for each product.

- If you are using Junos OS Subscriber Aware, configure SDF filters at the `[edit unified-edge pcef]` hierarchy level.

- If you are using Junos OS Broadband Subscriber Management, configure SDF filters at the `[edit services pcef]` hierarchy level.

To configure Layer 3 and Layer 4 SDF filters:

1. Specify a name for the SDF filter.

   For Junos OS Subscriber Aware:

   ```
   [edit unified-edge pcef]
   user@host# set flow-descriptions flow-identifier
   ```

   For Junos OS Broadband Subscriber Management:

   ```
   [edit services pcef]
   user@host# set flow-descriptions flow-identifier
   ```

2. Specify the flow direction for the SDF filter.

   > **NOTE**: If you do not specify a flow direction, then the SDF filter is applied in both the uplink and downlink directions.

   For Junos OS Subscriber Aware:

   ```
   [edit unified-edge pcef flow-descriptions flow-identifier]
   user@host# set direction (uplink | downlink | both)
   ```

   For Junos OS Broadband Subscriber Management:

   ```
   [edit services pcef flow-descriptions flow-identifier]
   user@host# set direction (uplink | downlink | both)
   ```

3. Specify a remote address (IPv4 or IPv6) for the SDF filter:

   > **NOTE**: You can specify an IPv4 subnet or an IPv6 subnet but not both.

   - Specify an IPv4 address for the SDF filter.

For Junos OS Subscriber Aware:

```
[edit unified-edge pcef flow-descriptions flow-identifier]
user@host# set remote-address ipv4-address ipv4-address
```

For Junos OS Broadband Subscriber Management:

```
[edit services pcef flow-descriptions flow-identifier]
user@host# set remote-address ipv4-address ipv4-address
```

- Specify an IPv6 address for the SDF filter.

  For Junos OS Subscriber Aware:

```
[edit unified-edge pcef flow-descriptions flow-identifier]
user@host# set remote-address ipv6-address ipv6-address
```

  For Junos OS Broadband Subscriber Management:

```
[edit services pcef flow-descriptions flow-identifier]
user@host# set remote-address ipv6-address ipv6-address
```

4. Specify a protocol (using the standard protocol number) for the SDF filter.

   For Junos OS Subscriber Aware:

```
[edit unified-edge pcef flow-descriptions flow-identifier]
user@host# set protocol number
```

   For Junos OS Broadband Subscriber Management:

```
[edit services pcef flow-descriptions flow-identifier]
user@host# set protocol number
```

5. Specify a local port or a list of port numbers for the SDF filter. To specify a list of port numbers (up to a maximum of three), enclose the port numbers in square brackets ([]).

> **ⓘ NOTE**: You can configure a local port or local port range but not both in the same SDF filter.

For Junos OS Subscriber Aware:

```
edit unified-edge pcef flow-descriptions flow-identifier]
user@host# set local-ports number
```

For Junos OS Broadband Subscriber Management:

```
edit services pcef flow-descriptions flow-identifier]
user@host# set local-ports number
```

6. Specify a local port range for the SDF filter.

   For Junos OS Subscriber Aware:

```
[edit unified-edge pcef flow-descriptions flow-identifier]
user@host# set local-port-range low low-value high high-value
```

   For Junos OS Broadband Subscriber Management:

```
[edit services pcef flow-descriptions flow-identifier]
user@host# set local-port-range low low-value high high-value
```

7. Specify a remote port or list of remote ports for the SDF filter. To specify a list of port numbers (up to a maximum of three), enclose the port numbers in square brackets ([]).

> **ⓘ NOTE**: You can configure a remote port or remote port range but not both in the same SDF filter.

   For Junos OS Subscriber Aware:

```
[edit unified-edge pcef flow-descriptions flow-identifier]
user@host# set remote-ports number
```

For Junos OS Broadband Subscriber Management:

```
[edit services pcef flow-descriptions flow-identifier]
user@host# set remote-ports number
```

8. Specify a remote port range for the SDF filter.

   For Junos OS Subscriber Aware:

```
[edit unified-edge pcef flow-descriptions flow-identifier]
user@host# set remote-port-range low low-value high high-value
```

   For Junos OS Broadband Subscriber Management:

```
[edit services pcef flow-descriptions flow-identifier]
user@host# set remote-port-range low low-value high high-value
```

RELATED DOCUMENTATION

Configuring Application-Aware Policy Control for Subscriber Management

Understanding Application-Aware Policy Control for Subscriber Management

Understanding Predefined Policy and Charging Control Rules for Subscriber-Aware Traffic Treatment | **54**

Configuring Policy and Charging Control Rules | **83**

Configuring Policy and Charging Control Action Profiles For Junos OS Subscriber Aware

A PCC action profile defines the treatment to be applied to specific service data flows or to packets associated with specific applications. A PCC action profile is specified in the then clause of a PCC rule.

> (i) **NOTE**: To make a change to a PCC action profile, you must be in maintenance mode. (See "Changing PCEF Profiles, PCC Rules, PCC Rulebases, Diameter Profiles, Flow Descriptions, and PCC Action Profiles" on page 217).

To configure PCC action profiles:

1. Specify a name for the PCC action profile.

```
[edit unified-edge pcef]
user@host# edit pcc-action-profiles profile-name
```

2. Configure the maximum bit rate for uplink and downlink subscriber traffic.

```
[edit unified-edge pcef pcc-action-profiles profile-name]
user@host# set maximum-bit-rate uplink mbr-uplink-value downlink mbr-downlink-value
```

The range is 0 through 6144000 Kbps.

3. Configure HTTP redirection to a URL.

```
[edit unified-edge pcef pcc-action-profiles profile-name redirect]
user@host# set url url-name
```

> **NOTE**: A PCC action profile that includes HTTP redirection can only be used in PCC rules that match only HTTP-based applications and all flows.

4. Configure the steering of traffic to a third-party server for applying services or to a service chain with one of the following methods:

- Specify the IP address of the third-party server for HTTP traffic.

```
[edit unified-edge pcef pcc-action-profiles profile-name steering path]
user@host# set (ipv4-address ipv4-address | set ipv6-address ipv6-address)
```

> **NOTE**: A PCC action profile that includes a steering path can only be used in PCC rules that match only HTTP-based applications and all flows.

- Specify the routing instance to use to reach the third-party server or service chain.

```
[edit unified-edge pcef pcc-action-profiles profile-name steering]
user@host# set routing-instance downlink downlink-vrf-name uplink uplink-vrf-name
```

The downlink routing instance is applied to traffic going to the access side, and the uplink routing instance is applied to traffic being sent from the access side.

5. Specify that steering attributes configured in a PCC action profile that a PCC rule applies to a data flow session when it begins will continue to be applied to the data flow when the PCC rule match conditions are modified, deleted, or added to.

```
[edit unified-edge pcef pcc-action-profiles profile-name steering]
user@host# set keep-existing-steering
```

6. Specify the HCM profile that you want to use for determining which HTTP header enrichment rules are applied.

```
[edit unified-edge pcef pcc-action-profiles profile-name]
user@host# set hcm-profile hcm-profile-name
```

> (i) **NOTE**: A PCC action profile that includes an HCM profile can only be used in PCC rules that match only HTTP-based applications and all flows.

7. Specify the forwarding class that you want packets to be assigned.

```
[edit unified-edge pcef pcc-action-profiles profile-name]
user@host# set forwarding-class class-name
```

8. Configure the gating status by enabling or disabling the forwarding of packets.

```
[edit unified-edge pcef pcc-action-profiles profile-name]
user@host# set gate-status (disable-both | downlink | uplink | uplink-downlink)
```

RELATED DOCUMENTATION

Understanding Predefined Policy and Charging Control Rules for Subscriber-Aware Traffic Treatment | **54**

Configuring Policy and Charging Control Rules | **83**

Configuring TDF Subscriber Usage Monitoring for Traffic That Matches Predefined PCC Rules | **101**

Configuring Policy and Charging Control Rules

A policy and charging control (PCC) rule defines the treatment to be applied to packets associated with specific applications or to specific service data flows.

You can configure PCC rules for Junos OS Subscriber Aware or for Junos OS Broadband Subscriber Management, but you use a different CLI hierarchy level for each product.

- If you are using Junos OS Subscriber Aware, configure PCC rules at the `[edit unified-edge pcef]` hierarchy level.

- If you are using Junos OS Broadband Subscriber Management, configure PCC rules at the `[edit services pcef]` hierarchy level.

> **NOTE**: If you are using Junos OS Subscriber Aware, you must be in maintenance mode to make a change to a PCC rule. (See Changing PCEF Profiles, PCC Rules, PCC Rulebases, Diameter Profiles, Flow Descriptions, and PCC Action Profiles).

> **NOTE**: If you are using Junos OS Broadband Subscriber Management, you cannot change a PCC rule while it is being used by a subscriber. To modify the rule, you must log off the subscribers that are using the rule.

Before you configure PCC rules, you must do the following:

- Configure the service data flow (SDF) filters that the PCC rules reference.

- Configure the application groups and any custom applications that you want to reference in application-aware PCC rules.

- Configure the PCC action profiles that the PCC rules reference.

> **NOTE**: When specifying application-aware PCC rules in a PCEF profile, you must also configure a default Layer 3 or Layer 4 wildcard PCC rule to ensure that the default charging characteristics are applied to unmatched subscriber traffic without dropping that traffic. For example, the default Layer 3 or Layer 4 wildcard PCC rule prevents traffic based on DNS queries from being dropped. In addition, the policy (PCEF profile) that includes application-aware PCC rules must also include a wildcard Layer 3 or Layer 4 PCC rule at a lower precedence.

To configure PCC rules:

1. Specify a name for the PCC rule.

For Junos OS Subscriber Aware:

```
[edit unified-edge pcef]
user@host# edit pcc-rules rule-name
```

For Junos OS Broadband Subscriber Management:

```
[edit services pcef]
user@host# edit pcc-rules rule-name
```

2. In a `from` statement, specify an SDF filter to use Layer 3 or Layer 4 match conditions for filtering subscriber traffic.

For Junos OS Subscriber Aware:

```
[edit unified-edge pcef pcc-rules  rule-name]
user@host# set from flows flow-identifier
```

For Junos OS Broadband Subscriber Management:

```
[edit services pcef pcc-rules  rule-name]
user@host# set from flows flow-identifier
```

If you do not want to filter subscriber traffic based on SDF filters, use the `any` option.

For Junos OS Subscriber Aware:

```
[edit unified-edge pcef pcc-rules  rule-name]
user@host# set from flows any
```

For Junos OS Broadband Subscriber Management:

```
[edit services pcef pcc-rules  rule-name]
user@host# set from flows any
```

3. (Optional) Specify an application as a match condition for filtering subscriber traffic.

For Junos OS Subscriber Aware:

```
[edit unified-edge pcef pcc-rules  rule-name]
user@host# set from applications application-name
```

For Junos OS Broadband Subscriber Management:

```
[edit services pcef pcc-rules  rule-name]
user@host# set from applications application-name
```

4. (Optional) Specify multiple applications instead of specifying each application separately by specifying an application group as a match condition for filtering subscriber traffic.

   For Junos OS Subscriber Aware:

```
[edit unified-edge pcef pcc-rules rule-name]
user@host# set from application-groups application-group-name
```

   For Junos OS Broadband Subscriber Management:

```
[edit services pcef pcc-rules rule-name]
user@host# set from application-groups application-group-name
```

5. Specify the PCC rules action profile that defines the treatment to be applied to specific service data flows or to packets associated with specific applications.

   > **NOTE**: You can use PCC action profiles with HTTP redirection or HCM profiles only in PCC rules that match only HTTP-based applications and any flows.

   For Junos OS Subscriber Aware:

```
[edit unified-edge pcef pcc-rules rule-name]
user@host# set then pcc-action-profile profile-name
```

   For Junos OS Broadband Subscriber Management:

```
[edit services pcef pcc-rules rule-name]
user@host# set then pcc-action-profile profile-name
```

Configuring a Policy and Charging Control Rulebase

A policy and charging control (PCC) rulebase contains a set of PCC rules. Each rule specified in the PCC rulebase is assigned a precedence to designate the priority in which PCC rules are evaluated for selection in a policy and charging enforcement function (PCEF) profile.

> ⓘ **NOTE**: Starting in Junos OS Release 19.3R1, application-aware policy control is also supported for Broadband Subscriber Management if you have enabled Next Gen Services on the MX240, MX480 or MX960 router with the MX-SPC3 card.

You can configure PCC rulebases for Junos OS Subscriber Aware or for Junos OS Broadband Subscriber Management, but you use a different CLI hierarchy level for each product.

- If you are using Junos OS Subscriber Aware, configure PCC rulebases at the `[edit unified-edge pcef]` hierarchy level.

- If you are using Junos OS Broadband Subscriber Management, configure PCC rulebases at the `[edit services pcef]` hierarchy level.

> ⓘ **NOTE**: If you are using Junos OS Subscriber Aware, you must be in maintenance mode to make a change to a PCC rulebase. (See Changing PCEF Profiles, PCC Rules, PCC Rulebases, Diameter Profiles, Flow Descriptions, and PCC Action Profiles).

> ⓘ **NOTE**: If you are using Junos OS Broadband Subscriber Management, you cannot change a PCC rulebase while it is being used by a subscriber. To modify the rulebase, you must log off the subscribers that are using the rule.

Before you configure a PCC rulebase, you must do the following:

- Configure service data flow filters.

- Configure PCC action profiles.

- Configure PCC rules.

To configure a PCC rulebase:

1. Specify a name for the rulebase.

   For Junos OS Subscriber Aware:

   ```
   [edit unified-edge pcef ]
   user@host#  edit pcc-rulebases rulebase-name
   ```

   For Junos OS Broadband Subscriber Management:

   ```
   [edit services pcef ]
   user@host#  edit pcc-rulebases rulebase-name
   ```

2. Specify the PCC rules that the rulebase references and a precedence value (1 through 65,535) for each rule.

   > **(i) NOTE**:
   >
   > - The same rule can be configured in different rulebases and can have a different precedence.
   >
   > - The precedence assigned must be unique among the configured PCC rules.
   >
   > - A lower precedence value indicates a higher precedence. For example, if a PCC rulebase has two PCC rules with precedence 5 and 10 respectively, the PCC rule with precedence 5 is evaluated first.

   For Junos OS Subscriber Aware:

   ```
   [edit unified-edge pcef pcc-rulebases rulebase-name]
   user@host#  set pcc-rule rule-name precedence number
   user@host#  set pcc-rule rule-name precedence number
   user@host#  set pcc-rule rule-name precedence number
   ```

For Junos OS Broadband Subscriber Management:

```
[edit services pcef pcc-rulebases rulebase-name]
user@host#  set pcc-rule rule-name precedence number
user@host#  set pcc-rule rule-name precedence number
user@host#  set pcc-rule rule-name precedence number
```

Configuring RADIUS Servers

You must configure RADIUS servers before you can configure a RADIUS network element. A network
element is a load-balanced group of RADIUS servers providing policy management for TDF subscribers.

To configure a RADIUS server:

1. Configure a name for the RADIUS server.

   ```
   [edit access radius]
   user@host# set servers name
   ```

2. Specify the IP address of the RADIUS server.

   ```
   [edit access radius servers name]
   user@host# set address server-address
   ```

3. Configure an interface and IPv4 address to specify the source for RADIUS requests. The MX Series router sends RADIUS requests to the RADIUS server using this source address.

```
[edit access radius servers name]
user@host# set source-interface interface [ipv4-address address]
```

4. Configure a shared secret (password) to be used by the MX Series router and the RADIUS server.

```
[edit access radius servers name]
user@host# set secret password
```

5. Configure the port number to which the RADIUS requests are sent.

```
[edit access radius servers name]
user@host# set port port-number
```

6. Specify the RADIUS server port number to which the MX Series router sends RADIUS accounting-start and accounting-stop requests. RADIUS accounting-start and accounting-stop requests are used when the RADIUS server is not able to initiate a change of authorization (CoA) request without an accounting record.

```
[edit access radius servers name]
user@host# set accounting-port port-number
```

7. Configure the secret password to be used when sending accounting-start requests to the RADIUS server if the accounting secret password is different from the authentication secret password. RADIUS accounting-start requests are used when the RADIUS server is not able to initiate a CoA request without an accounting record.

```
[edit access radius servers name]
user@host# set accounting-secret password
```

8. Configure the number of attempts to contact the RADIUS server that the MX Series router is allowed to make when it does not receive a response to its initial request. You can specify from 1 through 10 retries. The default is 3.

```
[edit access radius servers name]
user@host# set retry attempts
```

9. Configure the amount of time that the MX Series router waits to receive a response from a RADIUS server before retrying a request. By default, the MX Series router waits 3 seconds. You can configure the timeout to be from 1 through 90 seconds.

```
[edit access radius servers name]
user@host# set timeout seconds
```

10. Allow dynamic requests from the RADIUS server so that CoA requests can be received.

```
[edit access radius servers name]
user@host# set allow-dynamic-requests
```

11. Configure the secret password to be used for CoA requests from the RADIUS server.

```
[edit access radius servers name]
user@host# set dynamic-requests-secret password
```

12. Configure a limit to the number of request retries within a specified time interval that the MX Series router can send to the RADIUS server. If the number of retries reaches this limit, the RADIUS server is marked as dead, and the MX Series router begins to send requests to other RADIUS servers in the network element.

```
[edit access radius servers name]
user@host# set dead-criteria-retries retry-number interval seconds
```

13. Configure the amount of time that must pass after a RADIUS server is first marked dead until it is marked as alive by the MX Series router. When the MX Series router marks the RADIUS server as alive, it can again send requests to the RADIUS server.

```
[edit access radius servers name]
user@host# set revert-interval seconds
```

RELATED DOCUMENTATION

Configuring RADIUS Network Elements

A network element is a load-balanced group of RADIUS servers providing policy management for TDF subscribers.

Before you configure a network element, you must do the following:

- Configure the RADIUS servers that are to be part of the network element.

To configure a network element:

1. Specify a name for the network element.

```
[edit access radius]
user@host# set network-elements name
```

2. Specify the RADIUS servers that make up the network element.

```
[edit access radius network-elements name]
user@host# set server name
```

3. Assign each server in the network element a priority from 1 through 16 (1 is the highest priority). You can have multiple servers with the same priority in a network element. All access requests are load balanced among the highest priority servers. If all the servers with the highest priority in the network element fail, then requests are load balanced among servers with the next highest priority level.

```
[edit access radius network-elements name server name]
user@host# set priority priority
```

4. Configure the maximum number of requests that can be queued to the network element. When the pending-request queue is full, any additional requests are dropped.

```
[edit access radius network-elements name]
user@host# set maximum-pending-reqs-limit number
```

5. Configure the pending-request queue high watermark for the network element. This is a percentage of the maximum number of requests that can be queued to the network element, which is configured

in the `maximum-pending-reqs-limit` *number* statement. When the queue size reaches the high watermark, a `flow control on` message is generated.

```
[edit access radius network-elements name]
user@host# set pending-queue-watermark watermark
```

6. Configure the pending-request queue low watermark for the network element. This is a percentage of the maximum size of the pending-request queue, which is configured in the `maximum-pending-reqs-limit` *watermark* statement. When the number of pending requests drops below this low watermark value after having exceeded the high watermark, a `flow control off` message is generated.

```
[edit access radius network-elements name]
user@host# set pending-queue-watermark-abate abate-watermark
```

RELATED DOCUMENTATION

Understanding How a RADIUS Server Controls Policy and Charging Control Rules | **62**

Understanding Network Elements | **68**

Configuring RADIUS Servers | **88**

Configuring an AAA Profile

An AAA profile is a collection of attributes to specify how the MX Series router interacts with RADIUS servers that control the activation and deactivation of policy and charging control (PCC) rules.

Before you configure an AAA profile, you must do the following:

- Configure the network elements that are to be included in the AAA profile.

To configure an AAA profile:

1. Configure a name for the AAA profile.

```
[edit unified-edge aaa]
user@host# set profiles aaa-profile-name
```

2. Specify the network element providing policy management for TDF subscribers.

```
[edit unified-edge aaa profiles aaa-profile-name radius authentication]
user@host# set network-element network-element-name
```

3. If the RADIUS servers in the network element providing policy management for TDF subscribers cannot initiate a change of authorization (CoA) request without an accounting record, specify that the network element is used for accounting.

```
[edit unified-edge aaa profiles aaa-profile-name radius accounting]
user@host# set network-element network-element-name
```

4. If the RADIUS servers in the network element providing policy management for TDF subscribers cannot initiate a CoA request without an accounting record, enable the initiation of a RADIUS accounting start from the MX Series router to the RADIUS servers.

```
[edit unified-edge aaa profiles aaa-profile-name radius policy]
user@host# set coa-accounting enable
```

5. Configure the RADIUS attribute that you want to carry the PCC rulebase name for rulebase activations from the RADIUS policy server to the MX Series router. By default, the rulebase name is carried in the ERX-Service-Activate Juniper vendor-specific attribute (VSA).

   a. Specify the numeric value for the RADIUS AVP.

   ```
   [edit unified-edge aaa profiles aaa-profile-name radius policy activation-attribute]
   user@host# set code numeric-code
   ```

   b. If the RADIUS AVP is vendor-specific, specify the vendor identification.

   ```
   [edit unified-edge aaa profiles aaa-profile-name radius policy activation-attribute]
   user@host# set vendor-id vendor-id
   ```

6. Configure the RADIUS attribute that you want to carry the PCC rulebase name for rulebase deactivations from the RADIUS policy server to the MX Series router. By default, the rulebase name is carried in the ERX-Service-Deactivate Juniper VSA.

a. Specify the numeric value for the RADIUS AVP.

```
[edit unified-edge aaa profiles aaa-profile-name radius policy deactivation-attribute]
user@host# set code numeric-code
```

b. If the RADIUS AVP is vendor-specific, specify the vendor identification.

```
[edit unified-edge aaa profiles aaa-profile-name radius policy deactivation-attribute]
user@host# set vendor-id vendor-id
```

RELATED DOCUMENTATION

Configuring a Policy and Charging Enforcement Function Profile for Junos OS Subscriber Aware Dynamic Policies

When a policy and charging enforcement function (PCEF) profile is configured with dynamic policy control, the policy and charging rules function (PCRF) can both provision policy and charging control (PCC) rules and activate PCC rules that are predefined on the Junos Subscriber Aware PCEF.

Before you configure a PCEF profile for dynamic policies, you must do the following:

- Configure a Diameter Gx profile.

- (Optional) Configure service data flow (SDF) filters.

- (Optional) Configure a PCC action profile.

- (Optional) Configure PCC rules, PCC rulebases, or both.

> (i) **NOTE**: You can add PCC rules or PCC rulebases to a dynamic PCEF profile without being in maintenance mode. To make other changes to a dynamic PCEF profile, you must be in maintenance mode.

> **NOTE**: When a PCEF profile includes application-aware PCC rules, you must also include a default Layer 3 or Layer 4 wildcard PCC rule to ensure that the default charging characteristics are applied to unmatched subscriber traffic without dropping that traffic. For example, the default Layer 3 or Layer 4 wildcard PCC rule prevents traffic based on DNS queries from being dropped. In addition, the PCEF profile that includes application-aware PCC rules must also include a wildcard Layer 3 or Layer 4 PCC rule at a lower precedence.

To configure a PCEF profile for dynamic policies:

1. Specify a name for the PCEF profile.

```
[edit unified-edge pcef]
user@host#  edit profiles profile-name
```

2. Specify one or more PCC rules and a precedence for each rule for dynamic policy control. A lower precedence value indicates a higher precedence.

```
[edit unified-edge  pcef profiles profile-name]
user@host#  set dynamic-policy-control pcc-rules rule-name precedence number
```

> **NOTE**: If the profile includes application-aware PCC rules, you must also include a wildcard Layer 3 or Layer 4 PCC rule at a lower precedence.

3. Specify one or more PCC rulebases for dynamic policy control.

```
[edit unified-edge  pcef profiles profile-name]
user@host#  set dynamic-policy-control pcc-rulebases rulebase-name
```

> **NOTE**: Make sure that the PCC rules and PCC rulebases configured in a PCEF profile do not overlap.

4. Specify a Diameter Gx profile.

```
[edit unified-edge  pcef profiles profile-name dynamic-policy-control]
user@host#  set diameter-profile gx-profile-name
```

Configuring a Policy and Charging Enforcement Function Profile for Junos OS Subscriber Aware Static Policies

A policy and charging enforcement function (PCEF) profile configured for static policy control specifies that policy and charging control (PCC) rules are provisioned by the Junos Subscriber Aware PCEF with no interaction from the policy and charging rules function (PCRF).

> (i) **NOTE**: To make a change to a static PCEF profile, you must be in maintenance mode. (See "Changing PCEF Profiles, PCC Rules, PCC Rulebases, Diameter Profiles, Flow Descriptions, and PCC Action Profiles" on page 217).

Before you configure a PCEF profile for static policies, you must do the following:

- Configure service data flow filters for PCC rules.

- Configure PCC action profiles for PCC rules.

- Configure PCC rules.

- (Optional) Configure PCC rulebases.

To configure a PCEF profile for static policies:

1. Specify a name for the PCEF profile.

```
[edit unified-edge  pcef]
user@host#  edit profiles profile-name
```

2. Specify one or more PCC rules and a precedence for each rule for static policy control. A lower precedence value indicates a higher precedence.

```
[edit unified-edge  pcef profiles profile-name]
user@host#  set static-policy-control pcc-rules rule-name precedence number
```

3. Specify one or more PCC rule bases for static policy control.

```
[edit unified-edge  pcef profiles profile-name]
user@host#  set static-policy-control pcc-rulebases rulebase-name
```

RELATED DOCUMENTATION

Configuring a Policy and Charging Enforcement Function Profile for Junos OS Subscriber Aware Policies That a RADIUS Server Controls

A policy and charging enforcement function (PCEF) profile configured for policy control by a RADIUS server specifies that the RADIUS server activates and deactivates policy and charging control (PCC) rulebases that you have predefined on the MX Series router.

Before you configure a PCEF profile for policies controlled by a RADIUS server, you must do the following:

- Configure PCC rulebases.

- Configure an AAA profile.

To configure a PCEF profile for policies controlled by a RADIUS server:

1. Specify a name for the PCEF profile.

```
[edit unified-edge  pcef]
user@host#  edit profiles profile-name
```

2. Specify one or more PCC rule bases for policy control by a RADIUS server.

```
[edit unified-edge  pcef profiles profile-name]
user@host#  set aaa-policy-control pcc-rulebases rulebase-name
```

3. Specify the AAA profile that identifies the RADIUS server policy control parameters.

```
[edit unified-edge  pcef profiles profile-name]
user@host#  set aaa-policy-control aaa-profile aaa-profile-name
```

4. Configure the user password for subscribers assigned to this PCEF profile.

```
[edit unified-edge  pcef profiles profile-name]
user@host#  set aaa-policy-control user-password password
```

RELATED DOCUMENTATION

Configuration of Static Time-of-Day PCC Rule Activation and Deactivation Overview

You configure static time-of-day PCC rule activation and deactivation to specify when a rule or rulebase within a static PCEF profile is active.

To configure static time-of-day PCC rules activation and deactivation:

1. Configure an NTP server.

    See "Configuring the NTP Server" on page 99.

2. Configure the activation and deactivation settings and apply them to a rule or rulebase.

    See "Configuring Static Time-of-Day PCC Rule Activation and Deactivation in a Junos OS Subscriber Aware PCEF Profile" on page 100

RELATED DOCUMENTATION

Understanding Static Time-of-Day PCC Rule Activation and Deactivation | **71**

Configuring the NTP Server

Before you use the static or dynamic time-of-day functionality for PCC rules, you must configure an NTP server.

To configure the NTP server:

1. Specify the IP address of the NTP server.

    ```
    [edit system]
    user@host# set ntp server ip-address
    ```

2. Enable the NTP process on the router.

    ```
    [edit system]
    user@host# set processes ntp enable
    ```

RELATED DOCUMENTATION

Configuring Static Time-of-Day PCC Rule Activation and Deactivation in a Junos OS Subscriber Aware PCEF Profile | **100**

Understanding Static Time-of-Day PCC Rule Activation and Deactivation | **71**

Configuring Static Time-of-Day PCC Rule Activation and Deactivation in a Junos OS Subscriber Aware PCEF Profile

You configure static time-of-day PCC rule activation and deactivation to specify when to activate or deactivate a rule or rulebase within a static PCEF profile.

Before you configure static time-of-day PCC rule activation and deactivation, configure the NTP server.

To configure static time-of-day PCC rule or rulebase activation and deactivation within a PCEF profile:

1. Specify a name for a time-of-day profile.

   ```
   [edit unified-edge pcef]
   user@host# set pcc-time-of-day-profiles profile-name
   ```

2. Specify the activation time in the time-of-day profile.

   ```
   [edit unified-edge pcef pcc-time-of-day-profiles profile-name]
   user@host# set rule-activation-time <day-of-week | day-of-month month> <hour:min>
   ```

   You can specify the time of day, the day, and the month of the year. The day can be expressed as the day of the month (**DAY1** through **DAY31** or **Last-day-of-month**) or the day of the week (for example, **MONDAY**). If you specify the day of the month, you can also specify the month of the year. If a time zone is configured on the router, the time-of-day settings apply to the configured time zone.

3. Specify the deactivation time in the time-of-day profile. Use the same combination of options that you used in Step 2.

   ```
   [edit unified-edge pcef pcc-time-of-day-profiles profile-name]
   user@host# set rule-deactivation-time <day-of-week | day-of-month month> <hour:min>
   ```

   If a day is not specified and the deactivation time of day setting is earlier than the activation time of day setting, then a rule is deactivated the day after it is activated.

4. Within a static PCEF profile, apply the time-of-day profile to individual rules or rulebases.

   ```
   [edit unified-edge pcef profiles profile-name static-policy-control]
   user@host# set pcc-rules rule-name precedence number time-of-day-profile profile-name
   user@host# set pcc-rulebases rulebase-name time-of-day-profile profile-name
   ```

   Those rules or rulebases use the activation and deactivation settings for subscribers assigned to the PCEF profile.

Configuring TDF Subscriber Usage Monitoring for Traffic That Matches Predefined PCC Rules

You can configure usage monitoring of TDF subscriber traffic that matches particular data flows or applications that are identified in a predefined PCC rule by identifying the appropriate monitoring key in the `pcc-action-profile` of the PCC rule. This monitoring key controls usage reporting for all the predefined PCC rules that use this `pcc-action-profile`.

To configure usage monitoring for a predefined PCC rule:

- For the `pcc-action-profile` that is used in the predefined PCC rule, specify the monitoring key that controls reporting:

```
[edit unified-edge pcef pcc-action-profiles profile-name]
user@host# set monitoring-key key_string
```

CHAPTER 6

Configuring TDF Subscribers

IN THIS CHAPTER

IP-Based and IFL-Based TDF Subscribers Overview

IN THIS SECTION

Junos Subscriber Aware implements the Third-Generation Partnership Project (3GPP) traffic detection function (TDF), enabling subscriber-aware policy enforcement and traffic steering that is application-aware. Before a user's data traffic can undergo TDF processing, a TDF subscriber session must be set up.

You can configure two types of TDF subscribers:

IP-Based Subscribers

IP-based subscriber sessions are initiated when Junos Subscriber Aware processes a RADIUS accounting start request for a potential subscriber from a gateway GPRS support node (GGSN), Packet Data Network Gateway (PGW), or broadband network gateway (BNG). An IP-based subscriber session is for one unique user IP address.

IFL-Based Subscribers

IFL-based subscriber sessions are initiated when you configure the TDF subscriber and assign it a set of interfaces. All traffic that the MX Series router receives on those interfaces shares the same IFL-based subscriber session.

RELATED DOCUMENTATION

IP-Based Subscriber Setup Overview

Junos Subscriber Aware initiates an IP-based subscriber session when it receives a RADIUS accounting request from a gateway GPRS support node (GGSN), Packet Data Network Gateway (PGW), or

broadband network gateway (BNG). An individual subscriber session is created for each unique source IP address.

The MX Series router can receive a RADIUS accounting request in two ways:

- When the MX Series router is identified as a RADIUS server for the GGSN, PGW, or BNG, you configure the GGSN, PGW, or BNG as a RADIUS client of the MX Series router. The RADIUS client sends the accounting request to a designated interface and IP address on the MX Series router, which sends it to the subscriber processing module.

- When the GGSN, PGW, or BNG does not treat the MX Series router as a RADIUS server, you configure a filter called a *snoop segment*. Junos OS examines RADIUS accounting requests that pass through the MX Series router to determine whether they match the filter, which is known as *snooping*. When an accounting request matches the filter, Junos OS copies the request and sends it to the subscriber processing module.

You specify how an IP-based subscriber session is created and how a subscriber's traffic is processed by configuring TDF domains and PCEF profiles, and configuring a selection process for applying them to subscribers. The selection process identifies the attribute-value pair (AVP) values in the RADIUS accounting start request that must be matched to select a particular TDF domain or PCEF profile.

Understanding the Definition of a Set of IP-Based Subscriber Properties with a TDF Domain

A traffic detection function (TDF) domain identifies a set of properties for creating a TDF IP-based subscriber session and specifying how TDF subscriber traffic is processed. You can create several TDF domains if you have multiple categories of subscribers. You configure a selection process to assign IP-based subscribers to a TDF domain. Multiple subscribers can be assigned to the same TDF domain.

IP-based TDF domains include the following information:

- An IP-based type of subscriber.

- The TDF logical interface (mif) that handles the subscriber traffic. A TDF interface is distinct from other types of interfaces and is used to associate a TDF domain's subscribers with an access interface in a virtual routing and forwarding table (VRF). The TDF logical interface also identifies the TDF service set that is applied to the traffic.

- (Optional) The PCEF profile that must be applied to the TDF subscriber. The PCEF profile specifies how to apply policy and charging rules to the TDF subscriber traffic. If the TDF domain does not specify a PCEF profile, you must configure a PCEF profile selection process in addtion to the TDF domain selection process.

- Source IP addresses for uplink traffic and destination IP addresses for downlink traffic that you do *not* want to undergo TDF processing.

- Idle timeout and maximum number of subscribers for the TDF domain.

- Source IP addresses for users who can become TDF subscribers, using address pools.

- (Not applicable to snooped messages) The enabling or disabling of an immediate RADIUS response message from the MX Series router to the accounting start message received from a gateway GPRS support node (GGSN), Packet Data Network Gateway (PGW), or broadband network gateway (BNG) RADIUS client.

- The method for constructing the Subscription-Id for the Diameter credit control request (CCR) message that is sent from the TDF to the PCRF for a TDF subscriber.

- The local policy (drop/forward packets, maximum bit rate, burst size) to apply to the subscriber packets entering the access interface of the TDF domain if a TDF subscriber session does not exist.

- One or more interfaces that face the access network and can carry traffic for the TDF subscriber.

RELATED DOCUMENTATION

Understanding Source IP Filtering with Address Pools in TDF Domains for IP-Based Subscribers

The TDF domain that is assigned to an IP-based subscriber can identify a set of source IP addresses of packets that need to undergo TDF processing. These sets of IP addresses are configured using address pools. Address pools can then be added to a TDF domain.

Address pools contain a set of IP addresses specified by network prefixes. You can configure more than one set of addresses in an address pool. You can configure address pools to contain IPv4 addresses or IPv6 addresses, but not both.

You can configure an address pool as a default pool, and a TDF domain uses the default address pool when an address pool is not explicitly specified for the TDF domain.

RELATED DOCUMENTATION

Understanding Selection of Properties for an IP-Based TDF Subscriber

When the MX Series router receives a RADIUS accounting start request from the access network's gateway GPRS support node (GGSN), Packet Data Network Gateway (PGW), or broadband network gateway (BNG) for an IP-based subscriber, it needs to select the properties to apply to a subscriber by selecting a traffic detection function (TDF) domain before setting up a TDF subscriber session. The domain-selection configuration identifies the values that various AVPs (such as the 3GPP IMSI or the IPv4 address) in the RADIUS request must match to select a particular TDF domain. For RADIUS requests that were snooped, the domain-selection configuration can identify the snoop segment that matched the request.

The domain-selection configuration includes one or more `term` statements, each of which includes `from` statements that must all be matched, and a `then` statement that identifies the name of the TDF domain. When a term matches, further terms are not evaluated if a PCEF profile is specified in either the selected TDF domain or in the `then` statement. If a PCEF profile is not specified in either the selected TDF domain or in the `then` statement, further terms are evaluated to find a PCEF profile for the subscriber.

If no TDF domain is selected, then the TDF subscriber session is not set up.

Before you can configure the TDF domain selection, you must configure a TDF gateway, the TDF domains, and the RADIUS client.

The match conditions for TDF domain selection include:

- (Not applicable to snooped messages) The RADIUS client (GGSN, PGW, or BNG) that is sending the accounting start request

- Values for called-station-id, calling-station-id, class, framed-ip-address, framed-ipv6-prefix, 3gpp-imsi, nas-ip-address, or user-name AVPs

- Values for other AVPs you identify

Figure 11 on page 107 shows an overview of the IP-based subscriber setup process.

Figure 11: IP-Based Subscriber Setup Process
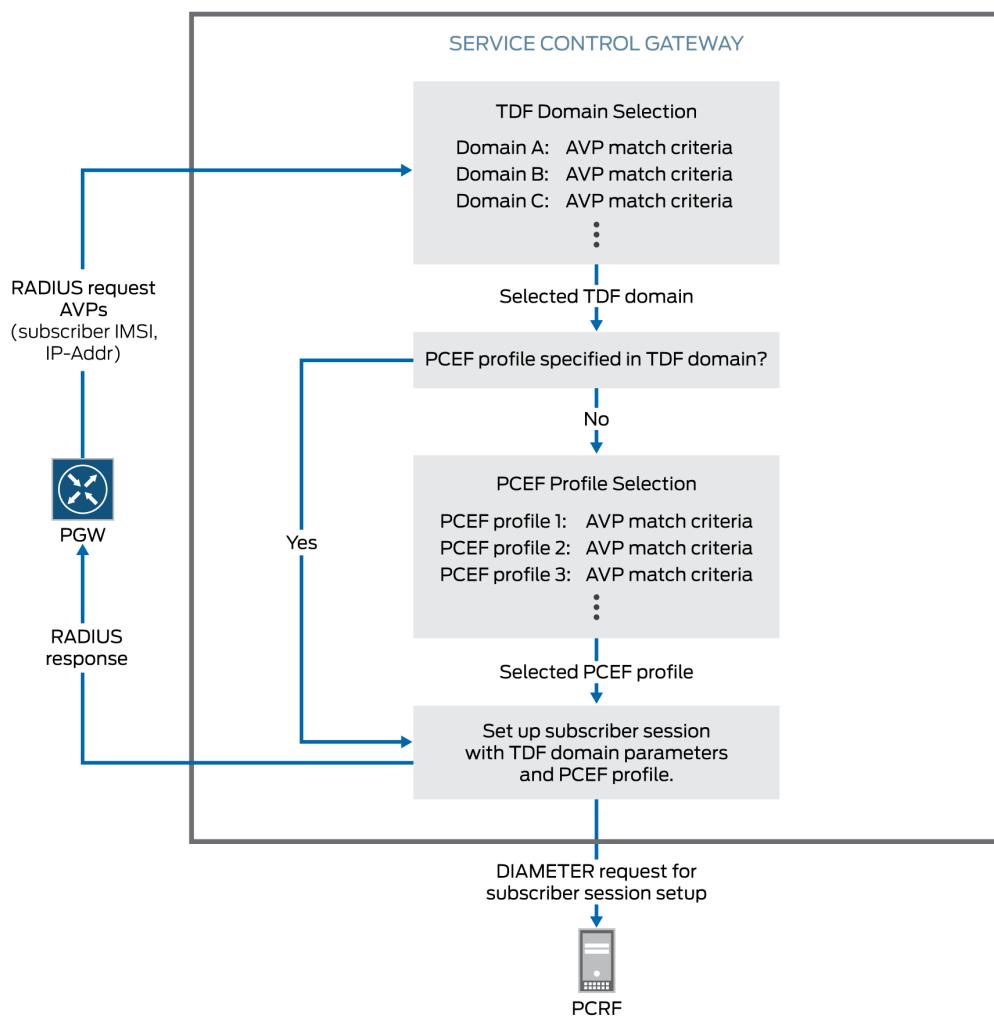
Understanding Selection of Policy-Control Properties for an IP-based TDF Subscriber

As part of the traffic detection function (TDF) subscriber session creation, the subscriber is assigned a policy and charging enforcement function (PCEF) profile, which specifies how policy and charging control (PCC) rules are defined on the TDF.

If every IP-based subscriber assigned to a TDF domain can share the same PCEF profile, then the PCEF profile can be specified within the TDF domain, under the `[edit unified-edge gateways tdf` *gateway-name* `domains]` hierarchy level. (For IFL-based subscribers, the PCEF profile *must* be specified within the TDF domain.)

If all of the IP-based subscribers assigned to the same TDF domain cannot share the same PCEF profile, the TDF domain does not specify a PCEF profile, and the PCEF profile selection must be configured under the `[edit unified-edge gateways tdf` *gateway-name* `domain-selection term]` hierarchy level. The `domain-selection term` consists of a `from` and a `then` statement.

The `from` statement identifies the match conditions for the subscriber. This includes the RADIUS client (GGSN, PGW, or BNG) that is sending the accounting start request for the subscriber and the values for particular AVPs in the message.

The `then` statement identifies the PCEF profile to assign to the subscriber. The `then` statement can also include the name of the TDF domain to assign to the subscriber. If the `then` statement only includes the PCEF profile, then another `domain-selection term` must assign a TDF domain to the subscriber.

When both a PCEF profile and a TDF domain are assigned to a subscriber in a `domain-selection term` statement, that PCEF profile is used even if the TDF domain specifies another PCEF profile.

Example: The TDF domain `domain1` specifies a PCEF profile. The `domain-selection term` does not need to specify a PCEF profile.

```
[edit unified-edge gateways tdf tdf1]
domain-selection {
    term 1 {
        from {
            client {
                client1;
            }
            user-name matches carrierA
        }
        then {
            domain domain1;
        }
    }
}
```

Example: The TDF domain `domain2` does not specify a PCEF profile. A `domain-selection term` must specify a PCEF profile. In this example, the PCEF profile is specified in the same term as the TDF domain.

```
[edit unified-edge gateways tdf tdf1]
domain-selection {
    term 1 {
        from {
            framed-ip-address equals 192.0.2.1/32
        }
        then {
            domain domain2;
            pcef-profile pcef3;
        }
    }
}
```

Example: The TDF domain `domain2` does not specify a PCEF profile. A `domain-selection term` must specify a PCEF profile. In this example, only the first term selects the TDF domain, so other terms must be added to select the PCEF profile.

```
[edit unified-edge gateways tdf tdf1]
domain-selection {
    term 1 {
```

```
        from {
            client {
                client2;
            }
            user-name matches carrierB
        }
        then {
            domain domain2;
        }
    }
    term 2 {
        from {
            framed-ip-address equals 192.0.2.1/32
        }
        then {
            pcef-profile pcef3;
        }
    }
    term 3{
        from {
            framed-ip-address equals 198.51.100.2/32
        }
        then {
            pcef-profile pcef4;
        }
    }
}
```

RELATED DOCUMENTATION

Snooping RADIUS Accounting Requests for IP-Based Subscribers Overview

When the gateway GPRS support node (GGSN), Packet Data Network Gateway (PGW), or broadband network gateway (BNG) does not identify the MX Series router as a RADIUS server, RADIUS accounting requests are not sent to a particular MX Series router IP address and interface configured for RADIUS messages. In this situation, you can configure the MX Series router to actively examine RADIUS accounting requests that pass through the MX Series router. This process is known as *snooping*. Junos OS identifies accounting requests that match a filter you configure, copies those requests, and sends them to the subscriber processing module.

To configure snooping, you configure filters called *snoop segments*. You can include the following conditions in a snoop segment:

- Destination IP address of the accounting request

- Shared secret between the accounting request sender and the MX Series router

- (Optional) Destination port of the accounting request

- (Optional) MX Series router interface that receives the accounting request

- (Optional) Source IP address of accounting requests from a GGSN, PGW, or BNG

You can also configure the length of time to cache the accounting request that was snooped. Any duplicate request that is received by the MX Series router within this time is dropped.

You can configure multiple snoop segments.

RELATED DOCUMENTATION

IP-Based Subscriber Setup Overview | 103

Configuring IP-Based TDF Subscriber Setup When Accounting Requests Are Snooped | 114

Configuring Snooping of RADIUS Accounting Requests for IP-Based Subscribers | 131

Understanding IFL-Based Subscriber Setup

You use the CLI to configure an IFL-based subscriber for a particular interface or set of access interfaces. All user traffic that uses these interfaces belongs to the same subscriber session. The IFL-based subscriber session becomes active when at least one of its access interfaces is up.

You can specify the following types of interfaces:

- Physical Layer 3 Ethernet interface

- Layer 3 Aggregated Ethernet interface

- Integrated routing and bridging (IRB) interface

- IRB that contains Ether-channel and physical interface members

- Logical Tunnel interface

You specify how an IFL-based subscriber's traffic is processed by configuring the properties of the TDF domain in which the IFL-based subscriber is configured, which includes a pointer to the PCEF profile to assign to the subscriber.

When an IFL-based subscriber session is created, it is anchored on a session PIC based on a round-robin selection process. If a stand-alone session PIC goes down and any IFL-based subscribers are anchored on that PIC, Junos OS re-anchors a subscriber onto another session PIC.

An IFL-based subscriber session is deleted in the following situations:

- All of the subscriber's access interfaces are down. When at least one interface comes back up, the subscriber session is restored.

- Subscriber is removed from the configuration with the CLI.

- Subscriber is set to deactivate with the CLI.

- Subscriber is cleared with the CLI. You can later restore the subscriber by using the revert option with the clear command. (See *clear unified-edge tdf subscribers*.)

RELATED DOCUMENTATION

Understanding the Definition of a Set of IFL-Based Subscriber Properties with a TDF Domain

A traffic detection function (TDF) domain identifies a set of properties for the IFL-based subscribers configured in the TDF domain. You can create several TDF domains if you have multiple categories of subscribers. Multiple subscribers can be assigned to the same TDF domain.

TDF domains include the following information:

- Logical interface-based type of subscriber.

- Name of each subscriber.

- Interfaces that belong to a subscriber. An interface can belong to only one subscriber.

- The TDF logical interface (mif) that handles the subscriber traffic. A TDF interface is distinct from other types of interfaces and is used to associate a TDF domain's subscribers with an access interface in a virtual routing and forwarding table (VRF). The TDF logical interface also identifies the TDF service set that is applied to the traffic.

- The PCEF profile that must be applied to the TDF subscriber. The PCEF profile specifies how to apply policy and charging rules to the TDF subscriber traffic.

- Source IP addresses for uplink traffic and destination IP addresses for downlink traffic you do not want to undergo TDF processing.

RELATED DOCUMENTATION

Configuring IP-Based TDF Subscriber Setup When MX Series Router Is a RADIUS Server

This task describes how to configure IP-based TDF subscriber setup when the gateway GPRS support node (GGSN), Packet Data Network Gateway (PGW), or broadband network gateway (BNG) identifies the MX Series router as a RADIUS server. An IP-based TDF subscriber is defined by the AVP values in the RADIUS accounting request received.

Before you configure the subscriber setup, you must do the following:

- Configure the access interfaces on the MX Series router chassis.

- Configure the PCEF profile.

- Configure the interface and IP address that you want to receive RADIUS requests on the MX Series router.

- Configure a TDF gateway.

To configure IP-based subscriber setup when the MX Series router acts as a RADIUS server:

1. Configure the TDF interfaces that can be used by TDF subscribers.

   See "Configuring a TDF Logical Interface" on page 139.

2. Associate the TDF interface to an access interface in a VRF routing instance.

   See "Configuring TDF Interface to Access Interface Associations in VRFs" on page 139.

3. Configure sets of source IP addresses that TDF domains can use to accept traffic.

   See "Configuring Address Pools for Source-IP Filtering of IP-Based Subscribers" on page 115.

4. Configure TDF domains that can be assigned to subscribers.

   See "Configuring a Set of IP-Based TDF Subscriber Properties with a TDF Domain" on page 117.

5. Configure RADIUS clients that can send the subscriber accounting requests.

   See "Configuring RADIUS Clients That Send Accounting Requests for IP-Based Subscribers" on page 123.

6. Configure how Junos OS selects TDF domains and PCEF profiles for subscribers.

   See "Configuring Assignment of TDF Subscriber Properties and Policy-Control Properties to IP-Based Subscribers" on page 125.

RELATED DOCUMENTATION

Configuring IP-Based TDF Subscriber Setup When Accounting Requests Are Snooped

This task describes how to configure IP-based TDF subscriber setup when the gateway GPRS support node (GGSN), Packet Data Network Gateway (PGW), or broadband network gateway (BNG) does not identify the MX Series router as a RADIUS server.

Before you configure the subscriber setup, you must do the following:

- Configure the PCEF profile.

- Configure a TDF gateway.

To configure IP-based subscriber setup when the MX Series router *does not* act as a RADIUS server:

1. Configure the TDF interfaces that can be used by TDF subscribers.

   See "Configuring a TDF Logical Interface" on page 139.

2. Associate the TDF interface to an access interface.

See "Configuring TDF Interface to Access Interface Associations in VRFs" on page 139.

3. Configure sets of source IP addresses that TDF domains can use to accept traffic.

   See "Configuring Address Pools for Source-IP Filtering of IP-Based Subscribers" on page 115.

4. Configure TDF domains that can be assigned to subscribers.

   See "Configuring a Set of IP-Based TDF Subscriber Properties with a TDF Domain" on page 117.

5. Configure the snooping filters that examine RADIUS accounting requests.

   See "Configuring Snooping of RADIUS Accounting Requests for IP-Based Subscribers" on page 131.

6. Configure how Junos OS selects TDF domains and PCEF profiles for subscribers.

   See "Configuring Assignment of TDF Subscriber Properties and Policy-Control Properties to IP-Based Subscribers" on page 125.

RELATED DOCUMENTATION

IP-Based Subscriber Setup Overview | 103

Snooping RADIUS Accounting Requests for IP-Based Subscribers Overview | 111

Configuring Address Pools for Source-IP Filtering of IP-Based Subscribers

Address pools identify a set of IP addresses that a TDF domain for IP-based subscribers uses to determine which packets undergo TDF processing.

To configure address pools:

1. Specify a name for the address pool.

```
[edit access address-assignment]
user@host# set address-pools name
```

The pool name can contain letters, numbers, and hyphens (-) and can be up to 63 characters long.

2. Specify the protocol family (inet for IPv4 addresses and inet6 for IPv6 addresses) for the address pool.

```
[edit access address-assignment]
user@host# set address-pools name family (inet | inet6)
```

For example, to configure an address pool named *mbg-pool1* with IPv4 addresses:

```
[edit access address-assignment]
user@host# set address-pools mbg-pool1 family inet
```

3. Specify the network prefix for the address pool for the configured protocol family.

```
[edit access address-assignment]
user@host# set address-pools name family (inet | inet6) network [network-prefix] external-
assigned
```

> **(i) NOTE**: A address pool must have at least one network prefix configured. You can configure more than one network prefix by including the `network` statement multiple times.
>
> The `external-assigned` statement is required.

For example, to configure an address pool with network prefixes 10.100.0.0/16 and 192.168.0.0/16:

```
[edit access address-assignment]
user@host# set address-pools mbg-pool1 family inet network 10.100.0.0/16 external-assigned
user@host# set address-pools mbg-pool1 family inet network 192.168.0.0/16 external-assigned
```

4. (Optional) Specify that the address pool is the default pool.

A TDF domain uses the default address pool to specify the source addresses of packets that undergo TDF processing when an address pool is not specified for the TDF domain.

```
[edit access address-assignment]
user@host# set address-pools name default-pool
```

RELATED DOCUMENTATION

Configuring a Set of IP-Based TDF Subscriber Properties with a TDF Domain

You define a set of properties for processing IP-based subscriber traffic and for setting up the subscriber session by configuring a TDF domain. You can create multiple TDF domains.

A potential IP-based subscriber is assigned to a TDF domain through a TDF domain-selection process that you configure in another topic.

Before you begin to create a TDF domain for IP-based subscribers, make sure that you have done the following:

- Configured the TDF interface (mif-) that the TDF domain uses.

- Configured the access-facing interfaces that the TDF domain uses.

- Configured a VRF routing instance that includes the TDF interface and the access-facing interfaces.

- Configured the PCEF profile if the TDF domain specifies one.

- Configured the address pool that contains source IP addresses of packets that are excluded from TDF processing for the TDF domain.

To configure a TDF domain for IP-based subscribers:

Configuring the TDF Domain Name and AAA Parameters

To configure the TDF domain name and the AAA parameters that are used by the TDF domain to create TDF IP-based subscriber sessions:

1. Specify a name for the TDF domain. The name can be from 1 through 50 characters long.

```
[edit unified-edge gateways tdf gateway-name]
user@host# set domains domain-name
```

2. (Optional) Configure the TDF domain for IP-based subscribers.

```
[edit unified-edge gateways tdf gateway-name domains domain-name]
user@host# set subscriber-type ip
```

You may omit this step because the default `subscriber-type` for TDF domains is `ip`.

3. Specify one or more methods for constructing the Subscription-Id for the Diameter credit control request (CCR) message that is sent from the TDF to the PCRF for subscribers belonging to the TDF domain.

   a. Specify the type of information to use for the Subscription-Id.

   You can specify multiple types, and the order of preference matches the order in which you enter the types. describes the types.

   ```
   [edit unified-edge gateways tdf gateway-name domains domain-name ip-subscriber]
   user@host# set subscription-id subscription-id-options entry-name id-components [use-class
   | use-imsi | use-msisdn | use-nai | use-nas-port | use-nas-port-id | use-realm | use-
   username]
   ```

   You can specify multiple methods by including the *entry-name* variable multiple times.

   b. If you selected `use-class` in Step a, you can also configure a regular expression to parse the Class attribute contents, specify characters to insert between the resulting regular expression groups, and specify the subscription ID type.

   ```
   [edit unified-edge gateways tdf gateway-name domains domain-name ip-subscriber
   subscription-id]
   user@host# set use-class regex "value"
   user@host# set use-class pattern "pattern"
   user@host# set use-class subscription-id-type (imsi | msisdn | nai | private | sip-uri)
   ```

   where *value* is a regular expression and *pattern* indicates the characters to insert between regular expression groups, which are identified with **\n** for a group number.

For example, the following configuration generates " `000118191129|ALICE:DRAV3:`" out of " `000118191129#000118191129#ALICE:DRAV3:#7168#nflat#ADSL##`" and sets the type to IMSI:

```
[edit unified-edge gateways tdf TDF1 domains domain1 ip-subscriber subscription-id ]
user@host# set use-class regex "[^#]*#\([^#]*\)\#\([^#]*\)"
user@host# set use-class pattern "\1|\2"
user@host# set use-class subscription-id-type imsi
```

c. Specify a constant string for the Subscription-Id-Data value.

This constant value is used if none of the `subscription-id-options` methods can be used. In such a case, the Subscription-Id-Type is END_USER_PRIVATE.

```
[edit unified-edge gateways tdf gateway-name domains domain-name ip-subscriber]
user@host# set subscription-id constant value
```

Table 5: Options for id-components of Subscription-Id

Option	Subscription-Id Type	Subscription-Id Data
use-class	Configurable	Entire Class attribute by default. Class attribute value can be parsed with `regex` option under the [edit unified-edge gateways tdf gateway-name domains domain-name subscription-id use-class] hierarchy.
use-imsi	END_USER_IMSI	3GPP-IMSI
use-msisdn	END_USER_E164	Calling-Station-Id
use-nai	END_USER_NAI	User-Name
use-nas-port	END_USER_PRIVATE	NAS-Port
use-nas-port-id	END_USER_PRIVATE	NAS-Port-Id
use-realm	END_USER_PRIVATE	Realm portion of the User-Name in NAI format

Table 5: Options for id-components of Subscription-Id *(Continued)*

Option	Subscription-Id Type	Subscription-Id Data
use-username	END_USER_PRIVATE	Username portion of the User-Name in NAI format

4. (Not applicable to snooped messages) Enable or disable the sending of an immediate RADIUS response message to the accounting start message received from a gateway GPRS support node (GGSN), Packet Data Network Gateway (PGW), or broadband network gateway (BNG) RADIUS client (disabled is the default).

    If the option is disabled, the response is sent after the TDF subscriber session creation is complete.

    ```
    [edit unified-edge gateways tdf gateway-name domains domain-name ip-subscriber]
    user@host# set immediate-accounting-response (enabled | disabled)
    ```

Configuring Address Filtering

To restrict the traffic that undergoes TDF processing for the TDF domain by identifying source IP addresses for uplink traffic and destination IP addresses for downlink traffic:

1. Identify the network prefix of source and destination IP addresses for packets that *do not* undergo TDF processing. Specify **inet** for IPv4 prefixes and **inet6** for IPv6 prefixes.

    ```
    [edit unified-edge gateways tdf gateway-name domains domain-name]
    user@host# set subscriber-exclude-prefix family (inet | inet6) network address net-mask
    ```

2. Identify the address pool that contains source and destination IP addresses of packets that undergo TDF processing. Specify **inet** for IPv4 prefixes and **inet6** for IPv6 prefixes.

    ```
    [edit unified-edge gateways tdf gateway-name domains domain-name ip-subscriber ]
    user@host# set subscriber-address (inet | inet6) pool pool-name
    ```

    > **(i)** **NOTE**: The address pool must be configured at the `[edit access address-assignment]` hierarchy level.

Configuring Subscriber Services and Policies

To configure the services and policies for IP-based subscribers that belong to the TDF domain:

1. Identify the TDF interface for the TDF domain.

   The TDF domain uses the service set that is applied to this TDF interface.

   ```
   [edit unified-edge gateways tdf gateway-name domains domain-name]
   user@host# set tdf-interface mif.number
   ```

   > **NOTE**: The TDF interface (mif) must have been previously configured at the `[edit interfaces]` hierarchy level.

2. (Optional) Identify the PCEF profile that the TDF domain uses to apply policies.

   If you do not identify a PCEF profile, then the PCEF profile must be assigned under the `[edit unified-edge gateways tdf gateway-name domain-selection term]` hierarchy.

   ```
   [edit unified-edge gateways tdf gateway-name domains domain-name]
   user@host# set pcef-profile name
   ```

   > **NOTE**: The PCEF profile must have been previously configured at the `[unified-edge pcef]` hierarchy level.

Configuring Access Interfaces

To configure the interfaces that face the access network and carry traffic to and from the IP-based subscribers that belong to the TDF domain:

Specify at least one interface. You can specify multiple interfaces.

```
[edit unified-edge gateways tdf gateway-name domains domain-name ip-subscriber]
user@host# set access-interfaces interface-name
```

Configuring Session Controls

To configure the TDF session controls for subscribers that belong to the TDF domain:

1. Configure the idle timeout (in minutes) for the TDF subscriber session. The range is 0 through 300.

   ```
   [edit unified-edge gateways tdf gateway-name domains domain-name ip-subscriber]
   user@host# set idle-timeout idle-timeout
   ```

2. Configure the default TDF subscriber maximum bit rate (MBR) for uplink and downlink traffic.

   Uplink traffic originates from the subscriber towards the public data network (PDN), and downlink traffic comes from the PDN and is destined for the subscriber. The range is 0 through 6,144,000 Kbps.

   ```
   [edit unified-edge gateways tdf gateway-name domains domain-name]
   user@host# set maximum-bit-rate uplink mbr-uplink-value downlink mbr-downlink-value
   ```

3. Configure the default TDF subscriber allowed burst size for uplink and downlink traffic.

   Uplink traffic originates from the subscriber towards the public data network (PDN), and downlink traffic comes from the PDN and is destined for the subscriber. The range is 1500 through 1,500,000,000 bytes.

   ```
   [edit unified-edge gateways tdf gateway-name domains domain-name]
   user@host# set burst-size uplink uplink-burst-size downlink downlink-burst-size
   ```

4. Configure the maximum number of subscriber sessions allowed (in thousands) for the TDF domain. The range is 100 thousands through 5000 thousands.

   ```
   [edit unified-edge gateways tdf gateway-name domains domain-name ip-subscriber]
   user@host# set maximum-subscribers number
   ```

Configuring Default Policy

To configure the default local policy for handling subscriber traffic entering the access interface of the TDF domain if a TDF subscriber session does not exist:

1. Configure the flow action to take on the subscriber's traffic.

   ```
   [edit unified-edge gateways tdf gateway-name domains domain-name ip-subscriber]
   user@host# set default-local-policy flow-action (drop | forward)
   ```

2. Configure the maximum bit rate for the subscriber's traffic.

   ```
   [edit unified-edge gateways tdf gateway-name domains domain-name ip-subscriber]
   user@host# set default-local-policy maximum-bit-rate uplink mbr-uplink-value downlink mbr-downlink-value
   ```

   Uplink traffic originates from the subscriber towards the public data network (PDN); downlink traffic comes from the PDN and is destined for the subscriber. The range is 0 through 6144000 Kbps.

3. Configure the allowed burst size for the subscriber's traffic.

```
[edit unified-edge gateways tdf gateway-name domains domain-name ip-subscriber]
user@host# set default-local-policy burst-size uplink uplink-burst-size downlink downlink-
burst-size
```

Uplink traffic originates from the subscriber towards the public data network (PDN); downlink traffic comes from the PDN and is destined for the subscriber. The range is 1500 through 1,500,000,000 bytes.

RELATED DOCUMENTATION

Configuring RADIUS Clients That Send Accounting Requests for IP-Based Subscribers

You specify an MX Series router RADIUS client for each gateway GPRS support node (GGSN), Packet Data Network Gateway (PGW), or broadband network gateway (BNG) that sends IP-based subscriber session requests and identifies the MX Series router as a RADIUS server. This task is not used for snooped accounting requests.

Before you begin to configure a RADIUS client, make sure that you have configured the interface and IP address that you want to receive RADIUS requests on the MX Series router.

To configure the RADIUS clients:

1. Configure the name of the RADIUS client.

```
[edit access radius]
user@host# set clients client-name
```

2. Specify the IP address from which the RADIUS client sends the RADIUS requests.

```
[edit access radius]
user@host# set clients client-name address client-address
```

3. Specify the MX Series router interface and IPv4 address that receive RADIUS requests from the GGSN, PGW, or BNG.

```
[edit access radius]
user@host# set clients client-name source-interface interface ipv4-address address
```

4. Configure a shared secret to be used by the MX Series router and the RADIUS client for accounting.

```
[edit access radius]
user@host# set clients client-name accounting secret password
```

5. (Optional) Specify that the framed-ip-address is used for subscriber creation when both the framed-route and framed-ip-address attributes are in the RADIUS accounting request from the RADIUS client. The framed-ip-netmask is also used for subscriber creation if it is in the request.

```
[edit access radius]
user@host# set clients client-name prefer-framed-ip-address
```

By default, the framed-route attribute is used for subscriber creation when both the framed-route and framed-ip-address attributes are in the RADIUS accounting request.

6. (Optional) Specify that the framed-ipv6-prefix is used for subscriber creation when both the delegated-ipv6-prefix and framed-ipv6-prefix attributes are in the RADIUS accounting request from the RADIUS client.

```
[edit access radius]
user@host# set clients client-name prefer-framed-ipv6-prefix
```

By default, the delegated-ipv6-prefix attribute is used for subscriber creation when both the delegated-ipv6-prefix and framed-ipv6-prefix attributes are in the RADIUS accounting request.

7. Configure the duration, in seconds, that the RADIUS response messages (sent for request messages) are stored in the MX Series router response cache before they time out.

```
[edit access radius]
user@host# set clients client-name accounting response-cache-timeout seconds
```

8. Enable the RADIUS client for a specific TDF gateway.

```
[edit unified-edge gateways tdf gateway-name]
user@host# set aaa clients client-name
```

Use the *client-name* that you configured in Step 1.

RELATED DOCUMENTATION

IP-Based Subscriber Setup Overview | **103**

Configuring Assignment of TDF Subscriber Properties and Policy-Control Properties to IP-Based Subscribers

IN THIS SECTION

- Configuring the Term Name | **126**
- Configuring Match Conditions for the RADIUS Client | **126**
- Configuring Match Conditions for Snoop Segments | **127**
- Configuring Match Conditions for Predefined AVPs | **127**
- Configuring Match Conditions for Custom AVP Attributes | **129**
- Configuring the TDF Domain to Select | **130**
- Configuring the PCEF Profile to Select | **130**

You must configure the criteria that Junos OS uses to select a TDF domain for an IP-based subscriber, which determines how the subscriber session is set up and how the subscriber traffic is treated. (The domain-selection process does not apply to IFL-based subscribers, who are automatically assigned to

the TDF domain in which they are configured.) You configure a `term` to identify conditions that must be matched in the incoming RADIUS request in order to select a particular TDF domain.

You configure the selection of the policy-control properties by selecting a PCEF profile. The PCEF profile can be identified in the selected TDF domain, or you can independently configure the criteria for the selection of a PCEF profile.

Before you begin to configure TDF domain or PCEF profile selection, make sure that you have done the following:

- Configured a TDF gateway.

- Configured the TDF domains.

- Configured the PCEF profiles.

- Configured the RADIUS client.

To configure a term for TDF domain or PCEF profile selection, perform the following tasks and repeat this process for each term you want to configure:

Configuring the Term Name

To configure the name for the `term` that contains the `from` statements and the `then` statement:

- Configure a term name that is 1 through 50 characters in length.

```
[edit unified-edge gateways tdf gateway-name domain-selection]
user@host# set term term-name
```

Configuring Match Conditions for the RADIUS Client

Before you begin to configure a match condition for a RADIUS client, you must ensure that you have configured the RADIUS client at the `[edit access radius clients]` hierarchy level, and specified it as the `aaa-client` at the `[edit unified-edge gateways tdf gateway-name]` hierarchy level.

To configure a match condition for the RADIUS client that sent the incoming RADIUS request:

- Specify the client.

```
[edit unified-edge gateways tdf gateway-name domain-selection term term-name]
user@host# set from client client-name
```

Configuring Match Conditions for Snoop Segments

For RADIUS requests that were snooped, the domain-selection configuration can identify the snoop segment that matched the request.

To configure a match condition for the snoop segment:

- Specify the snoop segment.

```
[edit unified-edge gateways tdf gateway-name domain-selection term term-name]
user@host# set from snoop-segment snoop-segment-name
```

Configuring Match Conditions for Predefined AVPs

To configure match conditions for the called-station-id, calling-station-id, class, framed-ip-address, framed-ipv6-prefix, 3gpp-imsi, nas-ip-address, or user-name AVP in the incoming RADIUS request from the subscriber:

1. Configure any called-station-id match condition.

```
[edit unified-edge gateways tdf gateway-name domain-selection term term-name]
user@host# set from called-station-id (equals | matches) value
```

Use **equals** to specify a value the called-station-id must equal or use **matches** to specify a regular expression the called-station-id must match.

2. Configure any calling-station-id match condition.

```
[edit unified-edge gateways tdf gateway-name domain-selection term term-name]
user@host# set from calling-station-id equals value
```

or

```
[edit unified-edge gateways tdf gateway-name domain-selection term term-name]
user@host# set from calling-station-id matches value
```

Use **equals** to specify a value the calling-station-id must equal or use **matches** to specify a regular expression the calling-station-id must match.

3. Configure any class match condition.

```
[edit unified-edge gateways tdf gateway-name domain-selection term term-name]
user@host# set from class (equals | has-prefix | has-suffix | matches) value
```

Use **equals** to specify a value the class must equal, use **has-prefix** to specify the prefix that the class must have, use **has-suffix** to specify the suffix that the class must have, or use **matches** to specify a regular expression the class must match.

4. Configure any framed-ip-address match condition.

```
[edit unified-edge gateways tdf gateway-name domain-selection term term-name]
user@host# set from framed-ip-address equals value
```

5. Configure any framed-ipv6-prefix match condition.

```
[edit unified-edge gateways tdf gateway-name domain-selection term term-name]
user@host# set from framed-ipv6-prefix equals value
```

6. Configure any 3gpp-imsi match condition.

```
[edit unified-edge gateways tdf gateway-name domain-selection term term-name]
user@host# set from 3gpp-imsi (equals | has-prefix | has-suffix | matches) value
```

Use **equals** to specify a value the 3gpp-imsi must equal, use **has-prefix** to specify the prefix that the 3gpp-imsi must have, use **has-suffix** to specify the suffix that the 3gpp-imsi must have, or use **matches** to specify a regular expression the 3gpp-imsi must match.

7. Configure any nas-ip-address match condition.

```
[edit unified-edge gateways tdf gateway-name domain-selection term term-name]
user@host# set from nas-ip-address equals value
```

8. Configure any user-name match condition.

```
[edit unified-edge gateways tdf gateway-name domain-selection term term-name]
user@host# set from user-name (equals | has-prefix | has-suffix | matches) value
```

Use **equals** to specify a value the user-name must equal, use **has-prefix** to specify the prefix that the user-name must have, use **has-suffix** to specify the suffix that the user-name must have, or use **matches** to specify a regular expression the user-name must match.

Configuring Match Conditions for Custom AVP Attributes

To configure match conditions for up to five custom AVP attributes (other than the called-station-id, calling-station-id, class, framed-ip-address, framed-ipv6-prefix, 3gpp-imsi, nas-ip-address, or user-name) in the incoming RADIUS request from the subscriber:

1. Configure an attribute name that is 1 through 50 characters in length.

```
[edit unified-edge gateways tdf gateway-name domain-selection term term-name]
user@host# set from attribute name
```

2. Configure any match condition for the custom attribute's AVP code.

```
[edit unified-edge gateways tdf gateway-name domain-selection term term-name from attribute
name]
user@host# set code numeric-code
```

3. Configure any match condition for the custom attribute's vendor-id.

```
[edit unified-edge gateways tdf gateway-name domain-selection term term-name from attribute
name]
user@host# set vendor-id vendor-id
```

4. Configure any match condition for custom attribute data in integer format.

```
[edit unified-edge gateways tdf gateway-name domain-selection term term-name from attribute
name]
user@host# set format integer (equals | greater-than | less-than) value
```

5. Configure any match condition for custom attribute data in string format.

```
[edit unified-edge gateways tdf gateway-name domain-selection term term-name from attribute
name]
user@host# set format string (equals | has-prefix | has-suffix | matches) value
```

Use **equals** to specify a value the string must equal, use **has-prefix** to specify the prefix that the string must have, use **has-suffix** to specify the suffix that the string must have, or use **matches** to specify a regular expression the string must match.

6. Configure any match condition for custom attribute data in time format.

```
[edit unified-edge gateways tdf gateway-name domain-selection term term-name from attribute
name]
user@host# set format time (equals | greater-than | less-than) value
```

7. Configure any match condition for custom attribute data in IPv4 address format.

```
[edit unified-edge gateways tdf gateway-name domain-selection term term-name from attribute
name]
user@host# set format v4address equals value
```

8. Configure any match condition for custom attribute data in IPv6 address format.

```
[edit unified-edge gateways tdf gateway-name domain-selection term term-name from attribute
name]
user@host# set format v6address equals value
```

9. Configure any match condition for custom attribute data in IPv6 address prefix format.

```
[edit unified-edge gateways tdf gateway-name domain-selection term term-name from attribute
name]
user@host# set format v6prefix equals value
```

Configuring the TDF Domain to Select

To specify the TDF domain to select when the `from` conditions in the `term` have been matched:

- Specify the TDF domain name.

```
[edit unified-edge gateways tdf gateway-name domain-selection term term-name]
user@host# set then domain tdf-domain-name
```

Configuring the PCEF Profile to Select

If a particular TDF domain does not specify a PCEF profile or you want different members of the same TDF domain to have different PCEF profiles, you must specify the PCEF profile under the `[edit unified-edge gateways tdf gateway-name domain-selection]` hierarchy level.

To specify the PCEF profile to select when the `from` conditions in the `term` have been matched, use one of the following methods:

- Specify the PCEF profile name in the same `term` statement that specifies the TDF domain.

```
[edit unified-edge gateways tdf gateway-name domain-selection term term-name]
user@host# set from {...}
user@host# set then domain tdf-domain-name
user@host# set then pcef-profile pcef-profile-name
```

- Specify the PCEF profile name in a different `term` statement.

```
[edit unified-edge gateways tdf gateway-name domain-selection term term-name]
user@host# set from {...}
user@host# set then pcef-profile pcef-profile-name
```

RELATED DOCUMENTATION

Configuring Snooping of RADIUS Accounting Requests for IP-Based Subscribers

If a gateway GPRS support node (GGSN), Packet Data Network Gateway (PGW), or broadband network gateway (BNG) does not treat the MX Series router as a RADIUS server, Junos OS must actively snoop RADIUS accounting requests from that gateway to set up TDF subscriber sessions. Snooping uses a filter called a *snoop segment* to identify the requests to send to the subscriber management module.

To configure snooping of RADIUS accounting requests:

1. Configure a name for the snoop segment.

```
[edit access radius]
user@host# set snoop-segments snoop-segment-name
```

For example:

```
[edit access radius]
user@host# set snoop-segments 123
```

2. Specify the destination IP address of accounting requests to snoop.

```
[edit access radius snoop-segments snoop-segment-name]
user@host# set destination-ip-address destination-address
```

For example:

```
[edit access radius snoop-segments 123]
user@host# set destination-ip-address 10.102.30.102
```

3. (Optional) Specify the destination port of accounting requests to snoop.

```
[edit access radius snoop-segments snoop-segment-name]
user@host# set destination-port destination-port
```

For example:

```
[edit access radius snoop-segments 123]
user@host# set destination-port 52000
```

If this statement is not included, the destination port is set to 1813.

4. (Optional) Specify the source IP address of accounting requests from a GGSN, PGW, or BNG to snoop.

```
[edit access radius snoop-segments snoop-segment-name]
user@host# set source-ip-address source-address
```

For example:

```
[edit access radius snoop-segments 123]
user@host# set source-ip-address 10.11.11.11
```

If the source IP address is not included, snooping of accounting requests is not restricted by their source.

5. Specify the MX Series router interface on which the accounting requests to be snooped are received.

```
[edit access radius snoop-segments snoop-segment-name]
user@host# set source-interface source-interface
```

For example:

```
[edit access radius snoop-segments 123]
user@host# set source-interface ge-0/0/0.0
```

If the source interface is not included, snooping of accounting requests is not restricted by the interface that receives the request.

6. Specify the shared secret for the MX Series router and the accounting request sender.

```
[edit access radius snoop-segments snoop-segment-name]
user@host# set shared-secret secret
```

For example:

```
[edit access radius snoop-segments 123]
user@host# set shared-secret juniper
```

If the shared secrets do not match, the subscriber session is not set up.

7. (Optional) Configure the number of seconds to cache the accounting request that was snooped. If the same request is received by the MX Series router within this time, it is considered a duplicate request and is dropped.

```
[edit access radius snoop-segments snoop-segment-name]
user@host# set request-cache-timeout timeout
```

For example:

```
[edit access radius snoop-segments 123]
user@host# set request-cache-timeout 4
```

8. Repeat Steps "1" on page 132 through "7" on page 133 to configure additional snoop segments.

9. Assign one or more snoop segments to the TDF gateway.

```
[edit unified-edge gateways tdf gateway-name aaa]
user@host# set snoop-segments [snoop-segment-name]
```

For example, the following configures gateway1 to snoop accounting requests destined for the RADIUS server 10.102.30.102 on port 52000 that originate from IP address 10.11.11.11 and are received on interface ge-0/0/0.0:

```
[edit unified-edge gateways tdf gateway1 aaa]
user@host# set snoop-segments 123
```

RELATED DOCUMENTATION

Snooping RADIUS Accounting Requests for IP-Based Subscribers Overview | 111

Configuring IP-Based TDF Subscriber Setup When Accounting Requests Are Snooped | 114

IP-Based Subscriber Setup Overview | 103

Configuring IFL-Based TDF Subscriber Setup

This task describes how to configure IFL-based TDF subscriber setup.

Before you configure the subscriber setup, you must do the following:

- Configure the interfaces on the MX Series router chassis.

- Configure the PCEF profile.

- Configure a TDF gateway.

To configure IFL-based subscriber setup:

1. Configure the TDF interfaces that TDF subscribers can use.

   See "Configuring a TDF Logical Interface" on page 139.

2. Associate the TDF interface to an access interface in a VRF routing instance.

   See "Configuring TDF Interface to Access Interface Associations in VRFs" on page 139.

3. Configure the IFL-based subscribers.

   See "Configuring IFL-Based TDF Subscribers and Properties with a TDF Domain" on page 135.

RELATED DOCUMENTATION

Understanding IFL-Based Subscriber Setup | 111

Configuring IFL-Based TDF Subscribers and Properties with a TDF Domain

IN THIS SECTION

-
-
-
-
-

You configure one or more IFL-based TDF subscribers and a set of properties for processing the traffic for those subscribers by configuring a TDF domain. You can create multiple TDF domains.

Before you begin to create a TDF domain for IFL-based subscribers, make sure that you have done the following tasks:

- Configured the TDF interface (mif-) that the TDF domain uses.

- Configured the interfaces that the TDF domain uses.

- Configured a VRF routing instance that includes the TDF interface and the interfaces that the TDF domain uses.

- Configured the PCEF profile that the TDF domain uses.

To configure a TDF domain for IFL-based subscribers, perform the following:

Configuring the TDF Domain Name and Type

To configure the TDF domain name and type:

1. Specify a name for the TDF domain. The name can be from 1 through 50 characters long.

```
[edit unified-edge gateways tdf gateway-name]
user@host# set domains domain-name
```

For example:

```
[edit unified-edge gateways tdf TDF1]
user@host# set domains ifl-1
```

2. Configure the subscriber type for IFL-based subscribers.

```
[edit unified-edge gateways tdf gateway-name domains domain-name]
user@host# set subscriber-type ifl
```

Configuring IFL-Based Subscribers

To configure IFL-based subscribers:

1. Configure the name for a subscriber.

```
[edit unified-edge gateways tdf gateway-name domains domain-name]
user@host# set ifl-subscriber subscriber-name
```

For example:

```
[edit unified-edge gateways tdf TDF1 domains ifl-1]
user@host# set ifl-subscriber ifl-sub1
```

2. Configure one or more interfaces for the subscriber.

```
[edit unified-edge gateways tdf gateway-name domains domain-name  ifl-subscriber subscriber-
name]
user@host# set access-interfaces [interface-name]
```

For example:

```
[edit unified-edge gateways tdf TDF1 domains ifl-1 ifl-subscriber ifl-sub1]
user@host# set access-interfaces ae0.736
```

You can assign only one IFL-based subscriber to an interface.

3. Repeat Step 1 and Step 2 for each IFL-based subscriber you want to configure in the TDF domain.

Configuring Address Filtering

To restrict the traffic that undergoes TDF processing for the TDF domain by identifying source IP addresses for uplink traffic and destination IP addresses for downlink traffic:

- Identify the network prefix of source and destination IP addresses for packets that *do not* undergo TDF processing. Specify **inet** for IPv4 prefixes and **inet6** for IPv6 prefixes.

```
[edit unified-edge gateways tdf gateway-name domains domain-name]
user@host# set subscriber-exclude-prefix family (inet | inet6) network address net-mask
```

Configuring Subscriber Services and Policies

To configure the services and policies for IFL-based subscribers that belong to the TDF domain:

1. Identify the TDF interface for the TDF domain.

   The TDF domain uses the service set that is applied to this TDF interface.

```
[edit unified-edge gateways tdf gateway-name domains domain-name]
user@host# set tdf-interface mif.number
```

> **NOTE**: The TDF interface (mif) must have been previously configured at the [edit interfaces] hierarchy level.

2. Identify the PCEF profile that the TDF domain uses to apply policies.

```
[edit unified-edge gateways tdf gateway-name domains domain-name]
user@host# set pcef-profile name
```

> **NOTE**: The PCEF profile must have been previously configured at the `[unified-edge pcef]` hierarchy level.

Configuring Session Controls

To configure the TDF session controls for subscribers that belong to the TDF domain:

1. Configure the default TDF subscriber maximum bit rate (MBR) for uplink and downlink traffic.

   Uplink traffic originates from the subscriber towards the public data network (PDN), and downlink traffic comes from the PDN and is destined for the subscriber. The range is 0 through 6,144,000 Kbps.

```
[edit unified-edge gateways tdf gateway-name domains domain-name]
user@host# set maximum-bit-rate uplink mbr-uplink-value downlink mbr-downlink-value
```

2. Configure the default TDF subscriber allowed burst size for uplink and downlink traffic.

   Uplink traffic originates from the subscriber towards the public data network (PDN), and downlink traffic comes from the PDN and is destined for the subscriber. The range is 1500 through 1,500,000,000 bytes.

```
[edit unified-edge gateways tdf gateway-name domains domain-name]
user@host# set burst-size uplink uplink-burst-size downlink downlink-burst-size
```

RELATED DOCUMENTATION

Understanding IFL-Based Subscriber Setup | **111**

Understanding the Definition of a Set of IFL-Based Subscriber Properties with a TDF Domain | **112**

Configuring a Policy and Charging Enforcement Function Profile for Junos OS Subscriber Aware Dynamic Policies | **94**

Configuring a Services Interface for a Session PIC or Service PIC | **14**

Configuring a TDF Logical Interface

A TDF logical interface is distinct from other types of interfaces and is used to associate a TDF domain's subscribers with an access interface in a virtual routing and forwarding (VRF) table and with a TDF service set. You need to configure one TDF interface logical interface (unit) for every TDF domain.

To configure a TDF interface, you configure one or more logical interfaces (units) for the interface:

1. Configure a TDF logical interface. Repeat this step for each TDF domain.

```
[edit interfaces]
user@host# set mif unit interface-unit-number family family-name
```

2. (Optional) Configure the maximum transmission unit (MTU) size for the TDF logical interface.

```
[edit interfaces]
user@host# set mtu mtu-size
```

RELATED DOCUMENTATION

Configuring TDF Interface to Access Interface Associations in VRFs

Junos associates TDF interfaces (mif) with access interfaces. You must configure a virtual routing and forwarding (VRF) table for each TDF domain. The VRF must include the TDF interface and one or more access interfaces for the TDF domain.

Before you begin, make sure that you have done the following:

- Configured the access interfaces on the MX Series router chassis.

- Configured the TDF interfaces.

To configure a TDF interface-to-access port mapping in a VRF, specify the VRF and place both the TDF interface (unit) and the physical access interface unit in the same VRF.

- Configure the VRF routing instance.

```
[edit routing-instances]
user@host# set routing-instance interface mif.n
user@host# set routing-instance interface interface-name
```

RELATED DOCUMENTATION

Enhanced CoS Handling and Dynamic Subscriber Management (ACX7100-48L, ACX7332, ACX7348)

SUMMARY

Enhanced CoS Handling and Dynamic Subscriber Management for Junos OS Evolved release 25.1R1

IN THIS SECTION

Overview

Starting Junos OS Evolved Release 25.1R1, subscriber management and CoS handling capabilities support dynamic interface configurations with CoS IFLSET and ACI VLAN, enabling enhanced traffic management and prioritization, ensuring that critical services receive appropriate bandwidth and low-latency handling. The dynamic nature of these configurations supports real-time adjustments based on subscriber needs, enhancing service quality and network performance.

Low-rate mode for queue shaping supports granular traffic shaping, starting from 64 Kbps for low-rate applications that require precise bandwidth allocation. Even minimal traffic flows are managed efficiently, preventing congestion and maintaining optimal performance across the network. This feature is useful in environments with diverse traffic patterns and stringent service level agreements.

Support for dynamic updates to classifier, rewrite, and drop profiles without disrupting existing subscriber sessions. You can adjust traffic management policies dynamically, responding to changing network conditions or service requirements without causing downtime. Additionally, CoS support for LAC subscribers includes Layer-2 classification and Layer-3 rewrite in the uplink direction, improves traffic handling and service quality for subscribers. Configuring these CoS features enhances the network's ability to deliver high-quality services to LAC subscribers efficiently.

Benefits of Subscriber Management and CoS Handling

- Enhances traffic management and prioritization for dynamic subscriber interfaces, resulting in improved service quality and user experience.

- Provides precise traffic shaping capabilities for low-rate applications, optimizing network resource utilization.

- Facilitates dynamic adjustments to traffic management policies without disconnecting existing subscribers, ensuring seamless service.

- Improves traffic classification and management for LAC subscribers, enhancing overall network performance and service quality.

- Simplifies the configuration and management of egress filters for IPv4, reducing administrative overhead and configuration errors.

Configuration Commands

To effectively utilize the enhanced CoS and subscriber management features, you can use various CLI commands. These commands provide detailed insights into the configurations to maintain network performance, optimization and troubleshooting:

- `show user-plane class-of-service classifier`: Displays the details of CoS classifiers, allowing you to verify and troubleshoot classification rules.

- `show user-plane class-of-service rewrite-rule`: Views CoS rewrite rules, ensuring that traffic is being marked and handled correctly.

- `show user-plane class-of-service scheduler-map`: Shows the scheduler map details, which are crucial for understanding how traffic is prioritized and scheduled.

- `show user-plane class-of-service traffic-control-profile`: Displays the traffic control profile details, helping you manage and optimize traffic control settings.

- `show user-plane class-of-service interface <interface-name>`: Provides CoS information for a specific interface, useful for interface-specific troubleshooting and configuration validation.

RELATED DOCUMENTATION

User-Defined Variables in Dynamic Profiles

show user-plane class-of-service classifier

show user-plane class-of-service interface

show user-plane class-of-service rewrite-rule

show user-plane class-of-service traffic-control-profile

show user-plane class-of-service scheduler-map

CHAPTER 7

Configuring Services

Overview of Applying Services to Subscribers

Subscriber-aware services are enabled for the subscribers belonging to a specific TDF domain by creating a subscriber-aware service set. This service set is applied to the TDF domain's TDF interface (mif). These services are carried out on the service PIC that is identified by the service interface in the service set.

Subscriber-aware services are applied to a subscriber's traffic based on policy and control (PCC) rules. The PCC rules are either under local control, under PCRF dynamic control, or under activation and deactivation control by a RADIUS server, depending on the PCEF profile for the TDF domain.

You may also apply network address translation (NAT) services independently of the PCC rules by specifying NAT rules in the service set.

shows the relationships among subscriber-aware service sets and other configured objects.

Figure 12: Subscriber-Aware Service Set Relationships
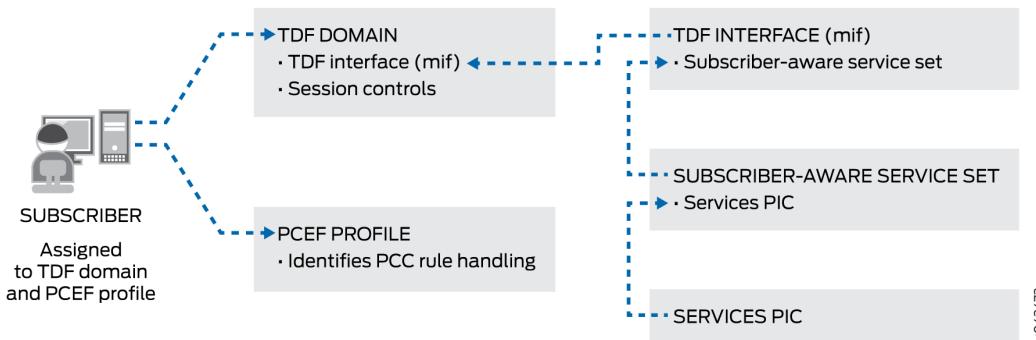
Applying Services to Subscriber-Aware Traffic with a Service Set

Junos OS supports subscriber-aware services for the subscribers belonging to a particular TDF domain through the configuration of a subscriber-aware service set. The service set is assigned to the TDF domain's TDF interface (mif).

Before you configure the service set, complete the following tasks:

- Configure the service PIC for the service set.

- Configure the TDF interface (mif).

- Configure the PCEF profile at the `[edit unified-edge pcef]` hierarchy level.

- Configure any NAT rules or rule sets that you want to apply.

To configure the subscriber-aware services for a TDF domain's subscribers:

1. Configure a PCEF profile at the [services] hierarchy level by specifying a name for the PCEF profile. This profile is a placeholder profile with no configuration options, but it must be created.

```
[edit services]
user@host# set pcef profile pcef-profile-name
```

2. Configure an application identification profile by specifying a name for the profile. This profile is a placeholder profile with no configuration options, but it must be created.

```
[edit services application-identification]
user@host# set profile app-id-profile-name
```

3. Configure an HTTP header enrichment profile by specifying a name for the profile. This profile is a placeholder profile with no configuration options, but it must be created.

```
[edit services hcm]
user@host# set profile hcm-profile-name
```

4. Define a subscriber-aware service set.

```
[edit services]
user@host# set service-set service-set-name service-set-options subscriber-awareness
```

5. Enable PCEF services for the service set. Use the profile name that you configured in Step 1.

```
[edit services service-set service-set-name]
user@host# set pcef-profile pcef-profile-name
```

6. Enable application identification for the service set. Use the profile name that you configured in Step 2.

```
[edit services service-set service-set-name]
user@host# set application-identification-profile app-id-profile-name
```

7. Enable HTTP header enrichment for the service set. Use the profile name that you configured in Step 3.

```
[edit services service-set service-set-name]
user@host# set hcm-profile hcm-profile-name
```

8. Specify NAT rules or rule-sets for the service set.

```
[edit services service-set service-set-name]
user@host# set ([nat-rules rule-name] | nat-rule-sets rule-set-name)
```

9. Specify the services PIC interface on which the services are performed.

```
[edit services service-set service-set-name]
user@host# set interface-service service-interface interface-name
```

The *interface-name* is ams*n* if you have redundancy configured and is ms-*fpc*/*pci*/0 if you do not have redundancy configured.

10. Apply the service set to the TDF interface (mif) that is part of the TDF domain.

```
[edit interfaces mif unit number family family service]
user@host# set input service-set service-set-name
user@host# set output service-set service-set-name
```

> **NOTE**: The output service set for the mif is not used by the MX Series router, but it must be configured so that the configuration commit does not fail.

RELATED DOCUMENTATION

Configuring Diameter

IN THIS CHAPTER

Diameter Profiles Overview

The Diameter profile provides network access information for the Diameter application. The Diameter profile specifies prioritized targets, or endpoints, for particular applications. The target specifies the destination realm, network element, and priority associated with the target.

Target selection is based on priority. A lower number has a higher priority. For load balancing, targets have the same priority.

From the prioritized list of targets for a Diameter profile, the target is selected as follows:

- The target with the highest priority (lowest number) is selected.

- In the event of a tie, where the priority is the same, target selection alternates among the peers with the same priority.

> ⓘ **NOTE**: Failover handling depends on what enables the policy for the application. Switching between targets based on priority, such as failing over between primary and secondary online charging servers, only occurs if the failover handling policy enables it.

After you configure the Diameter profiles, the Diameter applications can reference them. For example, when configuring transport profiles for online charging, you can associate the configured Diameter profile with the transport profile to interact with the online charging server. Similarly, when configuring profiles for provisioning Policy Charging and Control application rules, you can associate the configured Diameter profile with the policy and charging enforcement function (PCEF) profile to interact with the policy and charging rules function (PCRF).

RELATED DOCUMENTATION

Configuring Diameter Profiles | **150**

Juniper Networks Diameter AVPs for Subscriber Aware Policy Control

Diameter conveys information by including various attribute-value pairs (AVPs) in Diameter messages.

Table 6 on page 148 lists the AVPs for subscriber policy control.

Table 6: Juniper Networks Diameter AVPs for Subscriber Policy Control

Attribute Number	Diameter AVP	Description	Type
1100	TDF-Application-Instance-Identifier-Base	Identifies the application-group.	UTF8String
1101	Service-Chaining-Information	Provides service chaining information for dynamic steering of packets.	UTF8String
1102	LRF-Profile-Name	Provides the name of the logging and reporting framework (LRF) profile.	UTF8String
1103	HCM-Profile-Name	Provides the name of the HTTP content module.	UTF8String

Table 6: Juniper Networks Diameter AVPs for Subscriber Policy Control *(Continued)*

Attribute Number	Diameter AVP	Description	Type
1104	Forwarding-Class-Name	Provides the forwarding class name on the router.	UTF8String
1105	Redirect-VRF	Specifies whether redirection is supported. If the application flows support redirection, Redirect-VRF specifies the redirect address and address type.	UTF8String
1106	Requested-Burstsize-UL	Provides the uplink burst size specified in a QoS policy.	Integer32
1107	Requested-Burstsize-DL	Provides the downlink burst size specified in a QoS policy.	Integer32
1108	Steering-Information	Specifies an optional grouped AVP that contains Steering-Uplink-VRF, Steering-Downlink-VRF, and Steering-IP-Address.	Grouped
1109	Steering-Uplink-VRF	Provides the address of uplink destination for packets if dynamic steering is supported.	UTF8String
1110	Steering-Downlink-VRF	Provides the address of downlink destination for packets if dynamic steering is supported.	UTF8String
1111	Steering-IP-Address	Identifies the IP address for HTTP redirect.	Address

Configuring Diameter Overview

If you are using a PCRF to dynamically control subscriber-aware policies, you must configure Diameter.

To configure Diameter for PCRF-controlled subscriber-aware policies:

1. Configure the remote peer to which the MX Series router sends Diameter messages.

   See .

2. Identify the session PIC and PIC interfaces for a Diameter network element.

   See "Configuring Diameter Bindings" on page 152.

3. Configure the peers in a Diameter network element.

   See "Configuring Diameter Network Elements" on page 153.

4. Configure network access information in a Diameter profile.

   See "Configuring Diameter Profiles" on page 150.

5. (Optional) Specify the Diameter attribute-value pairs (AVPs) to include and exclude in the credit control request (CCR) messages.

   See "Configuring Diameter AVPs for Gx Applications" on page 154.

6. Configure the Diameter transport.

   See "Configuring the Diameter Transport" on page 159.

7. Configure the information to be advertised in Diameter messages.

   See "Configuring Advertisements in Diameter Messages" on page 160.

8. Configure the maximum number of pending requests for a Diameter application.

   See "Configuring Parameters for Diameter Applications" on page 161.

9. Configure the endpoint node that originates Diameter messages.

   See "Configuring the Origin Attributes of the Diameter Instance" on page 162.

RELATED DOCUMENTATION

Diameter Profiles Overview | 147

Configuring Diameter Profiles

The Diameter profile provides network access information for the Diameter application.

> ⓘ **NOTE**: To make a change to a Diameter profile, you must be in maintenance mode. (See "Changing PCEF Profiles, PCC Rules, PCC Rulebases, Diameter Profiles, Flow Descriptions, and PCC Action Profiles" on page 217).

To configure the Diameter profile:

1. Create the Diameter profile for the Gx application (`gx-profile`).

```
[edit]
user@host# set unified-edge diameter-profiles gx-profile profile-name
```

2. Set up the target for the profile.

```
[edit unified-edge diameter-profiles gx-profile profile-name]
user@host# set targets target-name
```

3. Specify the destination realm associated with the target.

```
[edit unified-edge diameter-profiles gx-profile profile-name targets target-name]
user@host# set destination-realm realm-name
```

4. Specify the priority associated with the target.

   The prioritization determines failover or load-balancing behavior. For load balancing, configure the targets with the same priority.

```
[edit unified-edge diameter-profiles gx-profile profile-name targets target-name]
user@host# set priority priority-value
```

5. Specify the network element associated with the target.

```
[edit unified-edge diameter-profiles gx-profile profile-name targets target-name]
user@host# set network-element element-name
```

6. (Optional) Specify the destination host associated with the target.

```
[edit unified-edge diameter-profiles gx-profile profile-name targets target-name]
user@host# set destination-host hostname
```

RELATED DOCUMENTATION

Configuring Diameter Bindings

You can configure a Diameter network element to run on a specific session PIC. You can organize other session PICs in a group around the selected session PIC on which the configured network element runs. When organized in a group, the selected session PIC can send and receive messages for other session PICs in the group. By default, every Diameter network element runs on every session PIC.

> (i) **NOTE**: If you want to set up Diameter bindings for session PICs on the broadband gateway, contact Juniper Networks Professional Services for assistance.

To configure the Diameter binding for network elements:

1. Configure the network element used for the Diameter binding on the broadband gateway.

   ```
   [edit]
   user@host# set unified-edge tdf gateway gateway-name diameter network-element element-name
   ```

2. Specify the session PICs group that serves the network element.

   ```
   [edit unified-edge tdf gateway gateway-name diameter network-element element-name]
   user@host# set session-pics group group-name
   ```

3. Specify the session PIC interfaces in this group that serve the network element. The interface must be a multiservices interface.

```
[edit unified-edge tdf gateway gateway-name diameter network-element element-name session-
pics group group-name]
user@host# set session-pic ams number
user@host# set session-pic ms-fpc/pic/port
```

RELATED DOCUMENTATION

Configuring Diameter Profiles | **150**

Configuring Diameter Network Elements

A Diameter network element consists of associated functions and a list of prioritized peers. The functions associate a Diameter application with the network element. The prioritization determines failover or load-balancing behavior for peer selection.

Before you configure Diameter network elements, perform the following task:

- Define the Diameter peers. See "Configuring Diameter Peers" on page 157.

To configure a Diameter network element:

1. Specify the name of the network element.

```
[edit access diameter]
user@host# set network-element element-name
```

2. Associate one or more functions with the network element.
   All functions are associated by default.

```
[edit access diameter network-element element-name]
user@host# set function function-name
```

3. Associate a Diameter peer with the network element and set the priority for the peer.

Peers with the lower priority number have the higher priority for peer selection. Peers with the same priority are load-balancing peers so the peer selection alternates between the two peers.

```
[edit access diameter network-element element-name]
user@host# set peer peer-name priority priority-value
```

4. (Optional) Associate a Diameter peer with the network element and set the amount of time to wait for a response from this peer before retransmitting the request to another peer. The default is 4 seconds.

```
[edit access diameter network-element element-name]
user@host# set peer peer-name timeout seconds
```

RELATED DOCUMENTATION

Configuring Diameter AVPs for Gx Applications

You can exclude Diameter attribute-value pairs (AVPs) from or include in the credit control request (CCR) messages between the MX Series router and the policy and charging rules function (PCRF) server.

> **NOTE**: The configuration of the Diameter AVPs for dynamic PCEF policies is optional.

To configure Diameter AVPs for Gx applications:

1. Specify the name of the Diameter Gx profile for which you are configuring the Diameter AVPs.

```
[edit]
user@host# edit unified-edge diameter-profiles gx-profile profile-name
```

The Diameter Gx profile name can contain letters, numbers, and hyphens (-) and can be up to 128 characters long.

2. Specify the optional AVPs to be excluded from the CCR messages between the MX Series router and the PCRF. By default, all AVPs are included in the CCR messages.

```
[edit unified-edge diameter-profiles gx-profile profile-name]
user@host# set attributes exclude [attribute]
```

You can specify more than one AVP in a single line.

Table 7 on page 155 describes the AVPs that you can exclude from CCR messages.

Table 7: Diameter AVP Exclusions for Gx Applications

AVP	Information in AVP
an-gw-address	AN-GW-Address AVP, which contains the IP addresses of the access node gateway.
default-eps-bearer-qos	Default-EPS-Bearer-QoS AVP.
packet-filter-information	Packet-Filter-Information AVP.
packet-filter-operation	Packet-Filter-Operation AVP.
rat-type	RAT-Type AVP.

3. Specify the optional AVPs to be included in the CCR messages between the MX Series router and the PCRF. By default, all AVPs are included in the CCR messages.

```
[edit unified-edge diameter-profiles gx-profile profile-name]
user@host# set attributes include [attribute]
```

You can specify more than one AVP in a single line.

Table 8 on page 156 describes the AVPs that you can included in CCR messages.

Table 8: Diameter AVP Inclusions for Gx Applications

AVP	Information in AVP
`gx-capability-list`	Gx-capability-list AVP.
`rule-suggestion`	Rule-suggestion AVP.

4. (Only on MX series devices) Support for customization of external Subscription ID is activated by default.

   Customization of Calling-Station-Id in RADIUS requests.

5. (Only on MX series devices) To customize Calling-Station-Id in RADIUS requests, use the command `set remote-circuit-id-format (postpend | prepend)` under `[edit access profile <profile-name> radius options]` mode.

```
[edit access profile <profile-name> radius options]
user@host# set remote-circuit-id-format postpend 1
```

6. (Only on MX series devices) To monitor usage, view the 3GPP attribute-value pairs (AVPs) defined as Gx for subscriber services using `monitoring-key default-value premium` configuration under `dynamic-profile`.

```
user@host# show configuration dynamic-profiles filter_service
variables
{
    var-input-filter default-value upstrm-filter;
    var-output-filter default-value dwnstrm-filter;
    monitoring-key default-value premium;
}
```

RELATED DOCUMENTATION

Configuring Diameter Profiles | **150**

Configuring Diameter Peers

You can configure the remote peers to which Diameter sends messages. Port 3868 is used for active connections to peers by default.

To configure a remote peer for a Diameter instance:

1. Specify the name of the Diameter peer.

```
[edit access diameter]
user@host# set peer peer-name
```

2. Specify the address of the Diameter peer.

```
[edit access diameter peer peer-name]
user@host# set address ip-address
```

3. Specify the transport that Diameter uses for active connections to the peer.

```
[edit access diameter peer peer-name]
user@host# set connect-actively transport transport-name
```

4. (Optional) Specify the port that Diameter uses for active connections to the peer. The default is port 3868.

```
[edit access diameter peer peer-name]
user@host# set connect-actively port port-number
```

5. (Optional) Specify the time to wait for connection acknowledgment from the peer. The default is 10 seconds.

```
[edit access diameter peer peer-name]
user@host# set connect-actively timeout seconds
```

6. (Optional) Specify the time to wait before trying to reconnect to a peer after receiving a Disconnect-Peer-Request message with the DO_NOT_WANT_TO_TALK_TO_YOU value for the Disconnect-Cause AVP. If you do not set a value, no reconnection attempt is made.

```
[edit access diameter peer peer-name]
user@host# set connect-actively repeat-timeout seconds
```

7. (Optional) Specify the time to wait for a Capabilities-Exchange-Answer message from the peer. The default is 10 seconds.

```
[edit access diameter peer peer-name]
user@host# set connect-actively capabilities-exchange-timeout seconds
```

8. (Optional) Specify the time to wait between connection attempts for this peer. The default is 30 seconds.

```
[edit access diameter peer peer-name]
user@host# set connect-actively retry-timeout seconds
```

9. (Optional) Specify the time to wait for a Device-Watchdog-Answer message from the peer. The default is 30 seconds.

```
[edit access diameter peer peer-name]
user@host# set watchdog-timeout seconds
```

10. (Optional) Specify the time to wait in the Closing state while disconnecting this peer. The default is 10 seconds.

```
[edit access diameter peer peer-name]
user@host# set disconnect-peer-timeout seconds
```

11. (Optional) Specify the size of the incoming queue for the peer. The default is 6000. You can specify a smaller value if you want to throttle the peer.

```
[edit access diameter peer peer-name]
user@host# set incoming-queue size size
```

12. (Optional) Specify the size of the outgoing queue for the peer. The default is 6000. You can specify a smaller value if you want to throttle the peer.

```
[edit access diameter peer peer-name]
user@host# set outgoing-queue size size
```

13. (Optional) Specify the high watermark of the outgoing queue for the peer.

The default is 80 percent. If the queue size reaches the high watermark, the peer is marked unavailable, any new messages to the Diameter network element are not sent to this peer, and the SNMP trap **Diameter_PeerOutQHiWMarkNotif** is generated.

```
[edit access diameter peer peer-name]
user@host# set outgoing-queue high-watermark high-watermark
```

14. (Optional) Specify the low watermark of the outgoing queue for the peer.

The default is 60 percent. If the queue size descends to the low watermark after reaching the high watermark, the peer becomes available and the SNMP trap **Diameter_PeerLowQHiWMarkNotif** is generated.

```
[edit access diameter peer peer-name]
user@host# set outgoing-queue low-watermark low-watermark
```

RELATED DOCUMENTATION

Configuring Diameter Profiles | 150

Configuring the Diameter Transport

You can configure one or more transports for a Diameter instance to set the source IP address for the local connection, and optionally configure a routing instance context. The routing instance for the transport connection must match that for the peer, or a configuration error is reported. Multiple peers can share the same transport.

To configure a transport for a Diameter instance:

1. Configure the transport name.

```
[edit access diameter]
user@host# set transport transport-name
```

2. Configure the source IP address for the Diameter local transport connection.

```
[edit access diameter transport transport-name]
user@host# set address ip-address
```

3. (Optional) Configure a routing instance, to which the address is bound, for the transport.

```
[edit access diameter transport transport-name]
user@host# set routing-instance routing-instance
```

RELATED DOCUMENTATION

Configuring Diameter Profiles | 150

Configuring Advertisements in Diameter Messages

You can configure information advertised in the Capabilities-Exchange-Request or Capabilities-Exchange-Answer messages. This information includes firmware revision, product name, and vendor identification.

To configure the advertisements:

1. (Optional) Specify the value for the Firmware-Revision AVP that is advertised. 0 is the default.

```
[edit access diameter]
user@host# set firmware-revision firmware-revision
```

2. (Optional) Specify the value of the Product-Name AVP that is advertised. Juniper Diameter Client is the default.

```
[edit access diameter]
user@host# set product-name name
```

3. (Optional) Specify the value of the Vendor-Id AVP that is advertised. 2636 is the default.

```
[edit access diameter]
user@host# set vendor-id vendor-id
```

RELATED DOCUMENTATION

Configuring Diameter Profiles | **150**

Configuring Parameters for Diameter Applications

You can configure parameters for Diameter applications, including the maximum number of pending requests.

To configure the parameters for the Diameter application:

1. Specify the Gx application (pcc-gx), for which you want to configure parameters.

```
[edit access diameter]
user@host# set applications pcc-gx
```

2. (Optional) Specify the maximum number of pending requests for the Diameter application. The default is 20,000.

```
[edit access diameter applications pcc-gx]
user@host# set maximum-pending-requests requests
```

RELATED DOCUMENTATION

Configuring Diameter Profiles | **150**

Configuring the Origin Attributes of the Diameter Instance

You can configure the identifying characteristics of the endpoint node that originates Diameter messages for the Diameter instance. The hostname is supplied as the value for the Origin-Host prefix. The realm is supplied as the value for the Origin-Realm attribute-value pair (AVP).

To configure the origin attributes:

1. Specify the Origin-Host prefix that originates the Diameter message.

   ```
   [edit access diameter origin]
   user@host# set host hostname
   ```

2. Specify the realm of the host that originates the Diameter message.

   ```
   [edit access diameter origin]
   user@host# set realm realm-name
   ```

RELATED DOCUMENTATION

Configuring Diameter Profiles | 150

3
PART

Configuring Reporting for Subscriber-Aware Data Sessions

CHAPTER 9

Configuring Reporting

Logging and Reporting Function for Subscribers

The logging and reporting function (LRF) enables you to log data for subscriber application-aware policy control sessions and send that data in an IPFIX format to an external log collector using UDP-based transport. These data session logs can include subscriber information, application information, HTTP metadata, data volume, time-of-day information, and source and destination details.

Starting in Junos OS Release 16.1R4 and in Junos OS Release 17.2R1, LRF is available in Junos OS Broadband Subscriber Management. Starting in Junos OS Release 19.3R2, LRF is available in Junos OS Broadband Subscriber Management if you have enabled Next Gen Services on the MX240, MX480 or MX960 router with the MX-SPC3 card..

The external collector, which is not a Juniper Networks product, can then use this data to perform analytics that provide you with insights about subscriber and application usage, allowing you to create packages and policies that increase revenue.

Log and Report Control

A subscriber's data sessions are logged and sent to collectors based on an LRF profile that you configure and associate with the subscriber.

The LRF profile includes:

- Templates—Specify the type of data that you want sent and the trigger that causes data to be sent. You can configure a maximum of 16 templates in an LRF profile.

- Collectors—Identify the destination to send data to. You can configure a maximum of eight collectors in an LRF profile.

- LRF rules—Specify the template and collector to use and, if applicable, a data volume limit that triggers the sending of data. An LRF rule's actions are performed when the matching conditions in a static PCC rule that references the LRF rule are met. You can configure a maximum of 32 LRF rules in an LRF profile.

To associate the LRF profile with a subscriber:

- For Junos OS Subscriber Aware, assign the LRF profile to the subscriber-aware TDF service set that belongs to the TDF interface (mif) in the subscriber's TDF domain.

- For Junos OS Broadband Subscriber Management, assign the LRF profile to the service set that is configured for application-aware policy control.

Templates

> **(i) NOTE**: If you have enabled Next Gen Services with the MX-SPC3 services card, then the DNS, IPv4 extended, IPv6 extended, mobile subscriber, video, and wireline subscriber templates are not supported.

You specify the data fields in a template by configuring one or more types for the template; for example, HTTP and IPv4. Each type represents a set of fields, and the template you configure includes fields from all the types you configure. The template is sent to the collector when you configure it, and is re-sent at a configurable interval. The template types that you can select and the fields that are included by each type are:

- Device Data—Contains data fields specific to the device collecting the logging feed:

  - DPI Engine Version

- IP address of TDF gateway (in IPv4 format)

- DNS—(Not available if Next Gen Services is enabled with the MX-SPC3 services card) Contains the DNS response time data field.

- Flow ID—Contains the Flow ID data field.

  When HTTP multiple transaction logging is enabled, FlowID is an implicit type that gets included with the HTTP template. When the consolidated session log is generated at the time of SESSION_CLOSE, LRF includes the FlowID that can be used to correlate with the HTTP transaction log records.

- HTTP—Contains data fields for the HTTP metadata from header fields:

  - User Agent

  - Content Length - Request

  - HTTP Response Code

  - Language

  - Host

  - Location

  - Http Method

  - Referer (HTTP)

  - MIME type

  - Time to First Byte

- IFL subscriber— Contains data fields specific to IFL-based subscribers:

  - Subscriber Name—Not applicable for BNG subscribers, hence this value is not be honored (is filled with zero).

  - IFL Name—Filled with default IFL name (filled with values Next Gen Services IFL)

- IPFlow—Contains data fields for the uplink and downlink octets and bytes. When a data record for volume limit is exported, these IPFlow statistics in the record are the actual data received after the last volume limit was reported in that data session and *not* cumulative data.

  - Uplink Octets

  - Downlink Octets

  - Uplink Packets

- Downlink Packets

- Ip Protocol—Protocol ID from IP header; for example, 17 (UDP), 6 (TCP).

- Record Reason—A value of `1` for the session close and a value of `2` for volume-limit.

- IPFlow Extended—Contains data fields for the service set name, routing instance, and payload timestamps. The initiator of the very first packet of a session is the client and the responder is the server.

  - Service-Set-Name—Filled with active `service-set-name` (16 byte value is filled active `service-set-name`. For example, if `service-set-name` is: bng-service-set-1, the template has a value of: bng-service-set-(16bytes)

  - Routing-Instance—Not applicable for BNG subscribers, hence this value is not be honored (is filled with zero).

- IPFlow TCP—Contains data fields for TCP-related timestamps:

  - Retransmitted TCP packets uplink

  - Retransmitted TCP packets downlink

  - TCP flow creation timestamp

- IPFlow TCP Timestamp—Contains IBM-specific data fields for TCP-related timestamps:

  - Smooth RTT uplink

  - Smooth RTT downlink

  - Client setup time

  - Server Setup time

  - First Client Payload timestamp

  - Upload time

  - First Server Payload timestamp

  - Download time

  - Acknowledged volumes uplink

  - Acknowledged volumes downlink

To use the IPFlow TCP Timestamp template when configuring an LRF profile, identify the template as vendor specific to avoid a commit warning. See *Configuring an LRF Profile for Subscribers*.

- IPFlow Timestamp—Contains data fields for the flow start and end timestamps:

  - Flow Start Time—For TCP, the flow start time is when the SYN packet is received. For UDP, it is when the first packet is sent.

  - Flow End Time

- IPv4—Contains data fields for the basic source and destination IPv4 information:

  - Source IPv4 Address

  - Destination IPv4 Address

- IPv4 Extended—(Not available if Next Gen Services with the MX-SPC3 services card are enabled) Contains data fields for the elements of IPv4 extended fields:

  - IPv4 TOS / Class of Service

  - IPv4 Source Mask

  - IPv4 Destination Mask

  - IPv4 Next Hop

- IPv6—Contains data fields for the basic source and destination IPv6 information:

  - Source IPv6 Address

  - Destination IPv6 Address

- IPv6 Extended—(Not available if Next Gen Services are enabled with the MX-SPC3 services card) Contains data fields for the elements of IPv6 extended fields:

  - IPv6 Source Mask

  - IPv6 Destination Mask

  - IPv6 Next Hop

  - Traffic Class

- L7 Application—Contains data fields for the Layer 7 application:

  - Application Protocol—Application data protocol below the classified application name; for example, `http` or `ssl`.

  - Application Name—Application name; for example, `junos:facebook` or `junos:Netflix`.

  - Host—HTTP header host when application protocol is `http`, SSL common name when application protocol is `ssl`, DNS name when application protocol is `dns`.

- Mobile Subscriber—(Not available if Next Gen Services with the MX-SPC3 services card are enabled) Contains data fields specific to mobile subscribers:

  - IMSI

  - MSISDN

  - IMEI

  - RAT-type

  - ULI

  - RADIUS Called Station ID

- PCC—Contains the PCC rule name data field.Not applicable if Next Gen Services are enabled.

- Status Code Distribution—Contains data fields for the HTTP or DNS status codes:

  - Status code 1

  - Status code 2

  - Status code 3

  - Status code 4

  - Status code 5

  - Num Instances 1

  - Num Instances 2

  - Num Instances 3

  - Num Instances 4

  - Num Instances 5

- Subscriber Data—Contains data fields for Generic Subscriber information that can be included with wireless (mobile) subscribers or wireline subscribers:

  - NAS_IP_ADDR—Not applicable for BNG subscribers, hence this value is not be honored (is filled with zero).

  - Subscriber Type—1 for IP-based subscriber, 2 for IFL-based subscriber.

  - Subscriber IP Address

  - Subscriber VRF—Not applicable for BNG subscribers, hence this value is not be honored (is filled with zero).

- NAS Port ID—Not applicable for BNG subscribers, hence this value is not be honored (is filled with zero).

- Accounting-Session-Id—Not applicable for BNG subscribers, hence this value is not be honored (is filled with zero).

- Class—Not applicable for BNG subscribers, hence this value is not be honored (is filled with zero).

- NAS Port Type—Not applicable for BNG subscribers, hence this value is not be honored (is filled with zero).

- Transport Layer—Contains data fields for the transport layer:

  - Source Transport Port

  - Destination Transport Port

- Video—(Not available if Next Gen Services with the MX-SPC3 services card are enabled) Contains data fields for video traffic:

  - Bitrate

  - Duration

- Wireline Subscriber—(Not available if Next Gen Services with the MX-SPC3 serices card are enabled) Contains the UserName data field for wireline subscribers. This is the same as RADIUS Called Station ID.

The template that is specified in an LRF rule determines the set of data fields that are included when data is sent to a collector. The data message includes a pointer to the template ID so that the collector can correlate the data contents with the data field lengths and types.

In a template, you also specify the type of trigger that determines when to send data to the collector. This trigger type can be a data volume limit, a time limit, or the closing of a data session (UDP sessions are considered closed after 60 seconds of inactivity; TCP sessions are considered closed when a FIN, FIN-ACK, or RST is received).

HTTP Transaction Logging

You may enable HTTP transaction logging in an LRF profile. This causes each HTTP transaction in a TCP session to be separately logged and sent to the collector, as shown in . This option is only relevant when the template being used includes HTTP in the template type.

By default, HTTP transaction logging is disabled, and the HTTP transaction records for a TCP session are sent together as one group of records.

Figure 13: HTTP Transaction Logging

Change History Table

Feature support is determined by the platform and release you are using. Use Feature Explorer to determine if a feature is supported on your platform.

Release	Description
19.3R1	Starting in Junos OS Release 19.3R2, LRF is available in Junos OS Broadband Subscriber Management if you have enabled Next Gen Services on the MX240, MX480 or MX960 router with the MX-SPC3 card..
16.1R4	Starting in Junos OS Release 16.1R4 and in Junos OS Release 17.2R1, LRF is available in Junos OS Broadband Subscriber Management.

RELATED DOCUMENTATION

Log Dictionary for Template Types

Configuring Logging and Reporting for Junos OS Subscriber Aware | 182

Configuring Logging and Reporting for Subscriber Management

Log Dictionary for Template Types

Table 9 on page 172 shows the logging dictionary of the template types that LRF supports. The log fields are a mix of IETF standard fields and fields that Juniper Networks defined. The IPFIX convention

for vendor-defined fields is an enterprise bit set to 1 and an enterprise ID set to the vendor-ID. (The Juniper Networks vendor-ID is 2636.) An IETF standard field has an enterprise bit set to 0 and no value for the enterprise ID.

> ⓘ **NOTE**: If you have enabled Next Gen Services with the MX-SPC3 services card, then the DNS, IPv4 extended, IPv6 extended, mobile subscriber, video, and wireline subscriber templates are not supported.

Table 9: Logging Dictionary for Template Types

Template Type	Field Name	Enterprise Bit/ID	Information Element Identifier	Data Type	Data Length (bytes)
Device Data	DPI Engine Version	1/2636	503	string	32
	IP address of TDF gateway.	1/2636	502	ipv4Address	4
DNS (Not available if Next Gen Services with the MX-SPC3 services card are enabled)	DNS response time	1/2636	876	dateTimeMilliseconds	8
Flow ID	Flow ID	1/2636	107	unsigned32	4
HTTP	User Agent	1/2636	152	string	32
	Content Length - Request	1/2636	154	unsigned32	4
	HTTP Response Code	1/2636	155	unsigned16	2

Table 9: Logging Dictionary for Template Types *(Continued)*

Template Type	Field Name	Enterprise Bit/ID	Information Element Identifier	Data Type	Data Length (bytes)
	Language	1/2636	156	string	16
	Host	1/2636	157	string	64
	Location	1/2636	158	string	64
	Http Method	1/2636	159	string	8
	Referer(HTTP)	1/2636	160	string	64
	MIME type	1/2636	161	string	32
	Http URI	1/2636	163	string	255
	Time to First Byte	1/2636	181	dateTimeMilliseconds	8
IFL Subscriber	Subscriber Name	1/2636	511	string Not applicable for BNG subscribers, hence this value is not be honored (is filled with zero).	16
	IFL Name	1/2636	512	string Filled with default IFL name (filled with values Next Gen Services IFL)	16
IPFlow	Uplink Octets	1/2636	103	unsigned32	4

Table 9: Logging Dictionary for Template Types *(Continued)*

Template Type	Field Name	Enterprise Bit/ID	Information Element Identifier	Data Type	Data Length (bytes)
	Downlink Octets	1/2636	104	unsigned32	4
	Uplink Packets	1/2636	105	unsigned32	4
	Downlink Packets	1/2636	106	unsigned32	4
	Ip Protocol	0	4	unsigned8	1
	Record Reason	1/2636	112	unsigned8	1

Table 9: Logging Dictionary for Template Types *(Continued)*

Template Type	Field Name	Enterprise Bit/ID	Information Element Identifier	Data Type	Data Length (bytes)
IPFlow Extended	Service-Set-Name	1/2636	520	string Contains data fields for the `service-set-name`, `routing-instance`, and payload timestamps. The initiator of the very first packet of a session is the client and the responder is the server. Filled with active `service-set-name` (16 byte value is filled active `service-set-name`. For example, if `service-set-name` is: bng-service-set-1, the template has a value of: bng-service-set-(16bytes)	16
	Routing-Instance	1/2636	521	string Not applicable for BNG subscribers, hence this value is not be honored (is filled with zero).	16

Table 9: Logging Dictionary for Template Types *(Continued)*

Template Type	Field Name	Enterprise Bit/ID	Information Element Identifier	Data Type	Data Length (bytes)
IPFlow TCP Timestamp	Retransmitted TCP packets uplink	1/2636	115	unsigned32	4
	Retransmitted TCP packets downlink	1/2636	116	unsigned32	4
	Smooth RTT uplink	1/2636	117	dateTimeMilliseconds	8
	Smooth RTT downlink	1/2636	118	dateTimeMilliseconds	8
	Client setup Time	1/2636	119	dateTimeMilliseconds	8
	Server Setup time	1/2636	120	dateTimeMilliseconds	8
	TCP flow creation timestamp	1/2636	121	dateTimeMilliseconds	8
	First Client Payload TS	1/2636	108	dateTimeMilliseconds	8
	Upload time	1/2636	113	dateTimeMilliseconds	8
	First Server Payload TS	1/2636	110	dateTimeMilliseconds	8

Table 9: Logging Dictionary for Template Types *(Continued)*

Template Type	Field Name	Enterprise Bit/ID	Information Element Identifier	Data Type	Data Length (bytes)
	Download time	1/2636	114	dateTimeMilliseconds	8
	Acknowledged volumes uplink	1/2636	122	unsigned64	8
	Acknowledged volumes downlink	1/2636	123	unsigned64	8
IPFlow Timestamp	Flow Start Time	1/2636	101	dateTimeMilliseconds	8
	Flow End Time	1/2636	102	dateTimeMilliseconds	8
IPv4	Source IPv4 Address	0	8	ipv4Address	4
	Destination IPv4 Address	0	12	ipv4Address	4
IPv4 Extended (Not available if Next Gen Services with the MX-SPC3 services card are enabled)	IPv4 TOS/Class of Service	0	5	unsigned8	1
	IPv4 Source Mask	0	9	unsigned8	1
	IPv4 Destination Mask	0	13	unsigned8	1

Table 9: Logging Dictionary for Template Types *(Continued)*

Template Type	Field Name	Enterprise Bit/ID	Information Element Identifier	Data Type	Data Length (bytes)
	IPv4 Next Hop	0	15	ipv4Address	4
IPv6	Source IPv6 Address	0	27	ipv6Address	16
	Destination IPv6 Address	0	28	ipv6Address	16
IPv6 Extended (Not available if Next Gen Services are enabled on the MX-SPC3 services card)	IPv6 Source Mask	0	29	unsigned8	1
	IPv6 Destination Mask	0	30	unsigned8	1
	IPv6 Next hop	0	62	ipv6Address	16
	Traffic Class	1/2636	126	unsigned8	1
L7 Application	Application Protocol	1/2636	151	string	32
	Application Name	1/2636	170	string	32
	Host	1/2636	157	string	64
Mobile Subscriber (Not available if Next Gen Services are enabled on the	IMSI	1/2636	504	string	16
	MSISDN	1/2636	505	string	16

Table 9: Logging Dictionary for Template Types *(Continued)*

Template Type	Field Name	Enterprise Bit/ID	Information Element Identifier	Data Type	Data Length (bytes)
MX-SPC3 services card)	IMEI	1/2636	506	string	16
	RAT-type	1/2636	507	unsigned8	1
	ULI	1/2636	508	string	13
	RADIUS Called Station ID	1/2636	509	string	32
PCC	PCC rule name	1/2636	901	string Not applicable if Next Gen Services are enabled.	64
Status Code Distribution	Status code 1	1/2636	171	unsigned16	2
	Status code 2	1/2636	172	unsigned16	2
	Status code 3	1/2636	173	unsigned16	2
	Status code 4	1/2636	174	unsigned16	2
	Status code 5	1/2636	175	unsigned16	2
	Num Instances 1	1/2636	176	unsigned16	2
	Num Instances 2	1/2636	177	unsigned16	2

Table 9: Logging Dictionary for Template Types *(Continued)*

Template Type	Field Name	Enterprise Bit/ID	Information Element Identifier	Data Type	Data Length (bytes)
	Num Instances 3	1/2636	178	unsigned16	2
	Num Instances 4	1/2636	179	unsigned16	2
	Num Instances 5	1/2636	180	unsigned16	2
Subscriber Data	NAS_IP_ADDR	1/2636	519	ipv4Address Not applicable for BNG subscribers, hence this value is not be honored (is filled with zero).	4
	Subscriber Type	1/2636	515	unsigned8 1 for IP-based subscriber, 2 for IFL-based subscriber	1
	Subscriber IP address	1/2636	516	ipv4Address	4
	Subscriber VRF	1/2636	517	unsigned32 Not applicable for BNG subscribers, hence this value is not be honored (is filled with zero).	4

Table 9: Logging Dictionary for Template Types *(Continued)*

Template Type	Field Name	Enterprise Bit/ID	Information Element Identifier	Data Type	Data Length (bytes)
	NAS Port ID	1/2636	518	string Not applicable for BNG subscribers, hence this value is not be honored (is filled with zero).	32
	Accounting-Session-Id	1/2636	514	string Not applicable for BNG subscribers, hence this value is not be honored (is filled with zero).	32
	Class	1/2636	522	String Not applicable for BNG subscribers, hence this value is not be honored (is filled with zero).	32
	NAS Port Type	1/2636	523	unsigned32 Not applicable for BNG subscribers, hence this value is not be honored (is filled with zero).	4
Transport Layer	Source Transport Port	0	7	unsigned16	2
	Destination Transport Port	0	11	unsigned16	2

Table 9: Logging Dictionary for Template Types *(Continued)*

Template Type	Field Name	Enterprise Bit/ID	Information Element Identifier	Data Type	Data Length (bytes)
Video (Not available if Next Gen Services are enabled on the MX-SPC3 services card)	Bitrate	1/2636	851	unsigned32	2
	Duration	1/2636	852	unsigned32	4
Wireline Subscriber (Not available if Next Gen Services are enabled on the MX-SPC3 services card)	UserName	1/2636	513	string	32

Configuring Logging and Reporting for Junos OS Subscriber Aware

To configure logging and reporting for traffic belonging to a set of subscribers, you configure LRF rules, collectors, and templates in an LRF profile; assign that LRF profile to the TDF service set associated with the subscribers' TDF domain; and assign each LRF rule to a PCC rule to activate it.

Before you begin to configure logging and reporting, you must:

- Configure the TDF domain for the subscriber.

- Configure the subscriber-aware service set for those subscribers.

To configure logging and reporting:

1. Configure an LRF profile to specify a set of logging and reporting parameters, which includes data templates, collectors, and LRF rules.

   See .

2. Assign the LRF profile to a set of subscribers.

   See .

3. Configure activation of an LRF rule with a static PCC rule.

   See "Configuring the Activation of an LRF Rule by a PCC Rule" on page 192.

RELATED DOCUMENTATION

| Logging and Reporting Function for Subscribers | **164**

Configuring an LRF Profile for Subscribers

> (*i*) **NOTE**: Starting in Junos OS Release 19.3R1, LRF profiles are also supported for Broadband Subscriber Management if Next Gen Services are enabled on the MX-SPC3 services card.

Configure an LRF profile to specify a set of logging and reporting parameters, which includes data templates, collectors, and LRF rules.

To configure an LRF profile:

Configuring the LRF Profile Name

An LRF profile is identified by a name, which you later specify in the service set for the subscribers.

- Configure a name for the LRF profile.

```
[edit services lrf]
user@host# set profile profile-name
```

For example:

```
[edit services lrf]
user@host# set profile lrf_profile1
```

Configuring Policy-Based Logging

Policy-based logging causes the LRF rules to be activated by PCC rules in a static PCEF profile.

- Configure policy-based logging in the LRF profile.

```
[edit services lrf profile profile-name]
user@host# set policy-based-logging
```

For example:

```
[edit services lrf profile  lrf_profile1]
user@host# set policy-based-logging
```

(Optional) Configuring HTTP Transaction Logging

Configure HTTP transaction logging if you want the HTTP metadata generated and sent separately for each transaction of a data session. This option is only relevant if the template specified in an LRF rule includes http in the template-type.

- Configure HTTP transaction logging in the LRF profile.

```
[edit services lrf profile profile-name]
user@host# set  http-log-multiple-transactions
```

For example:

```
[edit services lrf profile  lrf_profile1]
user@host# set  http-log-multiple-transactions
```

Configuring Collectors

Configure one or more collectors that you want to receive logging and reporting data when an LRF rule is activated. You can configure up to eight collectors for an LRF profile. For each collector:

1. Configure a name for the collector.

```
[edit services lrf profile profile-name]
user@host# set collector collector-name
```

For example:

```
[edit services lrf profile lrf_profile1]
user@host# set collector collector1
```

2. Specify the destination IP address of the collector.

```
[edit services lrf profile profile-name collector collector-name destination]
user@host# set address collector-address
```

For example:

```
[edit services lrf profile  lrf_profile1 collector collector1 destination]
user@host# set address 192.0.2.5
```

3. Specify the destination port of the collector.

```
[edit services lrf profile profile-name collector collector-name destination]
user@host# set port collector-port-number
```

For example:

```
[edit services lrf profile  lrf_profile1 collector collector1 destination]
user@host# set port 4739
```

4. Configure the source address to be used when exporting data to the collector.

```
[edit services lrf profile profile-name collector collector-name]
user@host# set source-address source-address
```

For example:

```
[edit services lrf profile lrf_profile1 collector collector1]
user@host# set source-address 10.1.1.1
```

Configuring Templates

Configure one or more templates, each of which specifies a set of data to be transmitted when an LRF rule is activated. You can configure up to 16 templates for an LRF profile. For each template:

1. Configure a name for the template.

```
[edit services lrf profile profile-name]
user@host# set template template-name
```

For example:

```
[edit services lrf profile lrf_profile1]
user@host# set template template1
```

2. Configure a format for the template. Only the IPFIX format is supported for this release.

```
[edit services lrf profile profile-name template template-name]
user@host# set format ipfix
```

For example:

```
[edit services lrf profile lrf_profile1 template template1]
user@host# set format ipfix
```

3. Configure the template types, which specify the data fields to include. You must configure at least one type, and you can configure multiple types.

```
[edit services lrf profile profile-name template template-name]
user@host# set template-type template-type
```

For example:

```
[edit services lrf profile lrf_profile1 template template1]
user@host# set template-type http ipv4
```

This example results in a template that includes fields from both the HTTP and IPv4 templates.

> **NOTE**: If you have enabled Next Gen Services on the MX-SPC3 services card, then the DNS, IFL subscriber, IPv4 extended, IPv6 extended, mobile subscriber, video, and wireline subscriber templates are not supported.

4. If you used the `ipflow-tcp-ts` template type, identify it as an IBM template to avoid a commit warning.

```
[edit services lrf profile profile-name]
user@host# set vendor-support ibm
```

5. Configure the interval, in seconds, at which you want the template to be retransmitted to the collector. The interval can be from 10 through 600, and the default is 60.

```
[edit services lrf profile profile-name template template-name]
user@host# set template-tx-interval tx-time
```

For example:

```
[edit services lrf profile lrf_profile1 template template1]
user@host# set template-tx-interval 100
```

6. Configure the type of trigger that causes the generation of data records and transmission to the collector. You can specify the trigger type as either the closing of the data session (default) or a data volume limit. The data volume limit value is specified within an LRF rule.

```
[edit services lrf profile profile-name template template-name]
user@host# set trigger-type (session-close | volume)
```

For example:

```
[edit services lrf profile lrf_profile1 template template1]
user@host# set trigger-type volume
```

Configuring Logging and Reporting Rules

Configure one or more LRF rules, which control how data sessions are logged and reported. You can configure up to 32 LRF rules for an LRF profile. For each LRF rule:

1. Configure a name for the LRF rule.

```
[edit services lrf profile profile-name]
user@host# set rule lrf-rule-name
```

For example:

```
[edit services lrf profile lrf_profile1]
user@host# set rule rule1
```

You cannot use the same LRF rule name in multiple LRF profiles.

2. Specify the collector that you want to receive the data if this rule is matched.

```
[edit services lrf profile profile-name rule lrf-rule-name ]
user@host# set then report collector collector-name
```

For example:

```
[edit services lrf profile lrf_profile1 rule rule1]
user@host# set then report collector collector1
```

3. Specify the template that identifies the type of data to report if this rule is matched.

```
[edit services lrf profile profile-name rule lrf-rule-name]
user@host# set then report template template-name
```

For example:

```
[edit services lrf profile lrf_profile1 rule rule1]
user@host# set then report template template1
```

4. If you specified volume for the template's trigger type in Step 6 of "Configuring Templates" on page 186, configure the data volume limit to be used for reporting by this rule.

```
[edit services lrf profile profile-name rule lrf-rule-name]
user@host# set then report volume-limit volume
```

The data volume, in megabytes, can be from 1 through 1024.

For example:

```
[edit services lrf profile lrf_profile1 rule rule1]
user@host# set then report volume-limit 4
```

5. If you specified time for the template's trigger type in Step 6 of "Configuring Templates" on page 186, configure the time limit to be used for reporting by this rule.

```
[edit services lrf profile profile-name rule lrf-rule-name]
user@host# set then report time-limit time-interval
```

The time limit, in seconds, can be from 60 through 1800. The default is 300.

For example:

```
[edit services lrf profile lrf_profile1 rule rule1]
user@host# set then report time-limit 360
```

Applying Logging and Reporting Configuration to a Subscriber Management Service Set

Assigning an LRF Profile to Subscribers

Before you can assign an LRF profile to a set of subscribers, you must:

- Configure the LRF profile.

- Configure the TDF interface (mif).

- Configure the TDF domain for the set of subscribers.

- Configure the service set for the TDF domain's TDF interface (mif).

Assign the LRF profile to a set of subscribers to apply the profile's logging and reporting configuration to the subscribers' traffic. You accomplish this by assigning the LRF profile to the subscriber-aware TDF service set associated with the TDF interface (mif) in the subscribers' TDF domain.

To assign an LRF profile to subscribers:

1. Identify the mif interface in the subscribers' TDF domain.

```
[edit unified-edge gateways tdf]
user@host# show domains domain-name
```

For example:

```
[edit unified-edge gateways tdf]
user@host# show domains domain1
```

```
pcef-profile pcef-prof-static;
tdf-interface mif.0;
access-interfaces {
    ge-1/0/1.0;
```

```
    }
    ...
```

2. Identify the service set or sets assigned to the mif interface.

```
[edit interfaces]
user@host# show mif.number
```

For example:

```
[edit interfaces]
user@host# show mif.0
```

```
family inet {
    service {
        input {
            service-set sset1;
        }
        output {
            service-set sset1;
        }
    }
}
```

3. Assign the LRF profile to the service set or sets.

```
[edit services service-set service-set-name]
user@host# set lrf-profile profile-name
```

For example:

```
[edit services service-set sset1]
user@host# set lrf-profile lrf_profile1
```

RELATED DOCUMENTATION

Configuring the Activation of an LRF Rule by a PCC Rule

> ⓘ **NOTE**: Starting in Junos OS Release 19.3R1, LRF rules are also supported for Broadband Subscriber Management if Next Gen Services are enabled on the MX-SPC3 services card.

> ⓘ **NOTE**: If you are using Junos OS Subscriber Aware, you must be in maintenance mode to make a change to a PCC action profile. (See "Changing PCEF Profiles, PCC Rules, PCC Rulebases, Diameter Profiles, Flow Descriptions, and PCC Action Profiles" on page 217).

> ⓘ **NOTE**: If you are using Junos OS Broadband Subscriber Management, you cannot make a change to a PCC action profile that is being used by subscribers. To modify the PCC action profile, you must first log off the subscribers that are using the PCC action profile.

Before you configure activation of an LRF rule by a PCC rule, you must:

- Configure the LRF rule in an LRF profile.

- Configure policy-based logging in the LRF profile.

- Configure the PCC rule.

You use a PCC rule's matching conditions to activate an LRF rule, which controls how data sessions are logged and reported. You identify the LRF rule in the PCC rule's action profile.

You can configure a PCC rule to activate an LRF rule for Junos OS Subscriber Aware or for Junos OS Broadband Subscriber Management, but you use a different CLI hierarchy level for each product.

- If you are using Junos OS Subscriber Aware, configure PCC rules at the `[edit unified-edge pcef]` hierarchy level.

- If you are using Junos OS Broadband Subscriber Management, configure PCC rules at the `[edit services pcef]` hierarchy level.

To configure a PCC rule to activate an LRF rule:

1. Identify the PCC action profile that is used in the PCC rule.

   For Junos OS Subscriber Aware:

   ```
   [edit unified-edge pcef]
   user@host# show pcc-rules rule-name
   ```

   For Junos OS Broadband Subscriber Management:

   ```
   [edit services pcef]
   user@host# show pcc-rules rule-name
   ```

   For example:

   For Junos OS Subscriber Aware:

   ```
   [edit unified-edge pcef]
   user@host# show pcc-rules all-traffic
   ```

   ```
   from {
       flows {
           all;
       }
   }
   then {
       pcc-action-profile all-traffic-action;
   }
   ```

   For Junos OS Broadband Subscriber Management:

> **NOTE**: The `from` statement is not applicable for Next Gen Services MX-SPC3 services card.

```
[edit services pcef]
user@host# show pcc-rules all-traffic
```

```
from {
    flows {
        all;
    }
}
then {
    pcc-action-profile all-traffic-action;
}
```

2. Assign the LRF rule to the PCC action profile.

   For Junos OS Subscriber Aware:

   ```
   [edit unified-edge pcef pcc-action-profiles profile-name]
   user@host# set logging-rule lrf-rule-name
   ```

   For Junos OS Broadband Subscriber Management:

   ```
   [edit services pcef pcc-action-profiles profile-name]
   user@host# set logging-rule lrf-rule-name
   ```

   For example:

   For Junos OS Subscriber Aware:

   ```
   [edit unified-edge pcef pcc-action-profiles all-traffic-action]
   user@host# set logging-rule rule1
   ```

For Junos OS Broadband Subscriber Management:

```
[edit services pcef pcc-action-profiles all-traffic-action]
user@host# set logging-rule rule1
```

RELATED DOCUMENTATION

Logging and Reporting Function for Subscribers | **164**

Configuring an LRF Profile for Subscribers | **183**

Configuring Policy and Charging Control Rules | **83**

4
PART

Modifying Subscriber-Aware Configuration

Modifying Subscriber-Aware Configuration in Maintenance Mode

IN THIS CHAPTER

Maintenance Mode Overview for Subscriber Aware Policy Enforcement

With Junos OS maintenance mode, you can take certain network functionality offline to perform specific maintenance tasks without disrupting service. When the traffic detection function (TDF) domains, TDF gateways, TDF subscribers, TDF interfaces, subscriber polices, or service PICs need maintenance, entering maintenance mode prevents these subscriber services elements from accepting new requests. You have the option of allowing all existing services to complete, or clear them. When ready, you can proceed with critical maintenance functions with a minimum of service disruption.

Subscribers who attempt to access a gateway that is in maintenance mode receive a message that the service is not supported.

If you want to perform any of the following operations, you must do so in maintenance mode:

- Delete or modify the addresses of certain TDF (mif) interfaces

- Delete or change the type of a TDF domain

- Change TDF interface configuration parameters

- Change a TDF interface for a TDF domain

- Change a static time-of-day profile

- Delete or modify a policy and charging enforcement function (PCEF) profile (However, maintenance mode is not required to add PCC rules or rulebases to a dynamic PCEF profile.)

- Delete or modify a PCC rule

- Delete or modify a PCC rulebase

- Delete or modify a Diameter profile

- Delete or modify a flow description

- Delete an address pool or modify its parameters

You can perform all other maintenance tasks outside of maintenance mode.

The maintenance mode procedures listed do not include adding elements. New elements carry no traffic and thus do not need to be gracefully halted. However, you can create new network elements in maintenance mode as an environment in which to test configurations before deploying them.

RELATED DOCUMENTATION

Changing Address Attributes in the Address Pool

This procedure describes how to place an address pool of a virtual routing and forwarding (VRF) instance in maintenance mode, allow all existing sessions using this pool to gracefully terminate, and then delete or modify pool attributes (for example, change address ranges in a pool).

To change address attributes in the address pool:

1. From configuration mode, activate maintenance mode for an address pool.

   ```
   [edit]
   user@host# set routing-instance vrf-name access address-assignment address-pools juniper-pool
   service-mode maintenance
   user@host# commit
   ```

2. Verify that all subscriber sessions have ended.

   ```
   user@host# run show unified-edge tdf address-assignment pool brief
   ```

   The service mode shows `Maintenance - Active Phase` if all the sessions are cleared. The service mode shows `Maintenance - In Phase` if some sessions are active. The service mode shows `Maintenance - Out Phase` if maintenance mode is not configured (that is, it is in operational mode).

3. (Optional) Terminate existing sessions using the **clear** command.

   ```
   [edit]
   user@host#  run clear unified-edge tdf subscribers routing-instance juniper-vrf
   ```

   When the subscriber count is zero and all sessions have terminated, the service mode status indicates `Maintenance - Active phase`. In this state, you can modify address pool attributes and commit changes.

4. Make changes to the pool.

5. Verify that changes were properly saved.

```
[edit]
user@host# run show configuration routing-instance access address-assignment address-pools
pool-name detail
```

> **NOTE**: These modifications, if made outside of active maintenance mode, fail.

6. Exit maintenance mode to return to normal operational mode.

```
[edit]
user@host# delete routing-instance juniper-vrf access address-assignment address-pools pool-
name service-mode
```

7. Return the gateway to operational state.

```
[edit]
user@host# run show unified-edge tdf gateway service-mode
```

RELATED DOCUMENTATION

Deleting an Address Pool

This procedure describes how to delete an address pool. You must first halt new sessions from being started and verify that no active sessions remain. The steps are similar to those described in "Changing Address Attributes in the Address Pool" on page 199.

To delete an address from an address pool:

1. From configuration mode, activate maintenance mode for an address pool.

```
[edit]
user@host# set routing-instance  juniper-vrf access address-assignment address-pools  pool-
```

```
name service-mode maintenance
user@host# commit
```

2. Verify that all subscriber sessions have ended.

```
[edit]
user@host# run show unified-edge tdf address-assignment pool brief
```

The service mode shows `Maintenance - Active Phase` if all the sessions are cleared. The service mode shows `Maintenance - In Phase` if some sessions are active. The service mode shows `Maintenance - Out Phase` if maintenance mode is not configured (that is, it is in operational mode).

3. (Optional) Terminate sessions that are using an address pool using the **clear** command.

```
[edit]
user@host#   run clear unified-edge tdf subscribers routing-instance juniper-vrf
```

When the subscriber count is zero and all sessions have terminated, the service mode status indicates `Maintenance - Active phase`. In this state, you can modify pool attributes and commit changes.

4. When the subscriber count is zero and all sessions have ended, modify address pool attributes and commit changes.

> ⓘ **NOTE**: These modifications, if made outside of active maintenance mode, fail.

5. Delete the address pool and commit the change.

```
[edit]
user@host# delete routing-instance juniper-vrf access address-assignment address-pools
juniper-pool
user@host# commit
```

6. Verify that the address pool has been deleted (that is, it is not listed in the output).

```
[edit]
user@host# run show configuration routing-instance juniper-vrf access address-assignment
address-pools juniper-pool
```

Changing AMS Interface Parameters on a TDF Gateway

This procedure shows how to change the parameters for an aggregated multiservices (AMS) interface on a TDF gateway using maintenance mode at the `[edit interfaces]` hierarchy level. If an AMS interface is configured under a gateway's session PICs or services PICs, and you change any load-balancing options such as membership of AMS interfaces (mams), then the AMS interface must be in maintenance mode.

Before you change AMS parameters using maintenance mode:

- Make sure that this change has been coordinated with affected groups and users.

To configure maintenance mode and AMS parameter change:

1. Verify the current status of maintenance mode for the AMS.

```
[edit]
user@host> show unified-edge tdf gateway-name system interfaces service-mode
```

The `service-mode` option displays the information details about maintenance mode as well as status.

```
Maintenance Mode
    MM Active Phase - System is ready to accept configuration changes for all
                 attributes of this object and its sub-hierarchies.
    MM In/Out Phase - System is ready to accept configuration changes only for
                 non-maintenance mode attributes of this object and
                 its sub-hierarchies.

Interface Name      Gateway Name         Service Mode
ms-1/0/0              SCG1                 Operational
ms-1/1/0              SCG1                 Operational
ms-2/0/0              SCG1                 Operational
ms-2/1/0              SCG1                 Operational
pfe-0/0/0             SCG1                  Operational
pfe-0/1/0             SCG1                  Operational
pfe-0/2/0             SCG1                  Operational
```

```
pfe-0/3/0              SCG1                    Operational
ams1                   SCG1                    Operational
```

2. From configuration mode, show the current configuration for the AMS interface.

```
user@host# show interfaces interface-name
load-balancing-options {
    member-interface mams-4/1/0;
    member-interface mams-5/1/0;
     member-failure-options {
        redistribute-all-traffic {
            enable-rejoin;
        }
    }
    high-availability-options {
        many-to-one {
            preferred-backup mams-5/1/0;
        }
    }
}
unit 1 {
    family inet;
}
unit 2 {
    family inet;
}
```

3. On the gateway, place the interface in maintenance mode.

```
[edit]
user@host# set unified-edge tdf gateway-name system interface interface-name service-mode
maintenance
user@host# commit
```

4. Verify that the AMS interface is in active maintenance mode where configuration changes are accepted for this object and all of its subhierarchies, after you commit the configuration.

```
user@host> show unified-edge tdf gateway-name system interfaces service-mode
```

```
Maintenance Mode
    MM Active Phase - System is ready to accept configuration changes for all
                  attributes of this object and its sub-hierarchies.
    MM In/Out Phase - System is ready to accept configuration changes only for
                  non-maintenance mode attributes of this object and
                  its sub-hierarchies.

Interface Name      Gateway Name        Service Mode
ms-1/0/0               SCG1                Operational
ms-1/1/0               SCG1                Operational
ms-2/0/0               SCG1                Operational
ms-2/1/0               SCG1                Operational
pfe-0/0/0             SCG1                 Operational
pfe-0/1/0             SCG1                 Operational
pfe-0/2/0             SCG1                 Operational
pfe-0/3/0             SCG1                 Operational
ams1                  SCG1                Maintenance - Active Phase
```

> (i) **NOTE**: All subscribers serviced by the AMS interface must go to zero. You can wait for these conditions to be met, or use the `clear` command for the interface (or gateway) to force these conditions.

5. Delete or change AMS member interfaces and parameters.

```
user@host> show unified-edge tdf gateway-name system interfaces service-mode
[edit unified-edge]
user@host# delete unified-edge tdf gateway-name system interface interface-name load-
balancing-options member-interface mams-interface-name
[edit interfaces]
user@host# set interfaces interface-name load-balancing-options member-interface mams-
interface-name
user@host# delete interfaces interface-name load-balancing-options high-availability-options
many-to-one preferred-backup mams-interface-name
```

```
user@host# set interfaces interface-name load-balancing-options high-availability-options
many-to-one preferred-backup mams-interface-name
```

6. Exit maintenance mode and commit the changes.

```
user@host# delete unified-edge tdf gateway-name system interface interface-name service-mode
maintenance
user@host# commit
```

RELATED DOCUMENTATION

Modifying a TDF Domain

This procedure describes how to use maintenance mode to modify a TDF domain. Options include modifying such parameters as TDF domain, mobile-interface, address filtering, AAA parameters, session characteristics, and access interfaces. You must first halt new sessions from being started and verify that there are no active sessions remaining.

To change a TDF domain for a group of subscribers that belong to that domain:

1. From configuration mode, activate maintenance mode for an TDF domain.

```
[edit]
user@host# set unified-edge gateways tdf  gateway-name domains domain-name service-mode
maintenance
user@host# commit
```

2. Verify that the TDF domain is in maintenance mode.

```
[edit]
user@host# run show unified-edge tdf domains service-mode
```

This command displays the service-mode status for all the TDF domains. You can verify the status for the specific TDF domain and take action accordingly.

The service mode for the TDF domain shows **Maintenance – Active Phase** if all the sessions using this TDF domain are cleared. The service mode for the TDF domain shows **Maintenance - In Phase** if some sessions are actively using this TDF domain.

3. Verify that no subscribers are active on the TDF domain.

```
[edit]
user@host# run show unified-edge tdf subscribers | match domain-name
```

4. (Optional) Terminate sessions on a TDF domain using the **clear** command.

```
[edit]
user@host# run clear  unified-edge tdf subscribers domain  domain-name gateway gateway-name
```

5. When the subscriber count is zero and all sessions have ended, make and commit changes to the TDF domain in active maintenance mode.

> **NOTE**: These modifications must be made in active maintenance mode or they fail.

6. Modify the TDF domain and commit the changes.

7. Exit maintenance mode and commit the changes.

```
[edit]
user@host# delete unified-edge gateways tdf gateway-name domains domain-name service-mode
user@host# commit
```

8. Verify that changes were properly committed.

```
[edit]
user@host# run show configuration unified-edge gateways tdf gateway-name tdf-services
domains  domain-name
```

The command output displays the configuration changes you made to the TDF domain.

9. Return the gateway to operational state.

```
[edit]
user@host# run show unified-edge tdf gateway service-mode
```

> ⓘ **NOTE**: Although maintenance mode does not explicitly include AAA options, certain AAA changes require you to place affected TDF domains in maintenance mode first. These changes include changing an AAA profile name and changing authorization or accounting elements. If you attempt to make AAA changes that affect a TDF domain that is not in maintenance mode, you are prompted to place the appropriate TDF domain into maintenance mode before proceeding with AAA profile name or element changes.

RELATED DOCUMENTATION

Modifying the TDF Interface of a TDF Domain

This procedure describes how to use maintenance mode to modify attributes of the TDF interface for a TDF domain. You must first halt new sessions from being started and verify that no active sessions remain.

To configure the mobile interface of a TDF domain:

1. From configuration mode, activate maintenance mode for the TDF domain using the mobile interface to be modified.

```
[edit]
user@host# set unified-edge gateways tdf  gateway-name domains  domain-name service-mode
maintenance
user@host# commit
```

2. Verify that the TDF domain of this mobile interface is in maintenance mode.

```
[edit]
user@host# run show unified-edge tdf domain service-mode
```

From the gateway hierarchy, the service mode for the gateway shows `Maintenance - Active Phase` if all the sessions using this TDF domain are cleared. The service mode for the gateway shows `Maintenance - In Phase` if some sessions are actively using this TDF domain. The service mode for the TDF domain shows `Maintenance - Out Phase` if maintenance mode is not configured (that is, it is in operational mode).

You cannot make and commit changes to a mobile interface unless the TDF domain to which it is attached is in maintenance mode.

3. Verify that no subscribers are active on the TDF domain.

```
[edit]
user@host# run show unified-edge tdf subscribers | match domain-name
```

4. (Optional) Terminate sessions that are using an address pool using the **clear** command.

```
[edit]
user@host# run clear  unified-edge tdf subscribers domain  domain-name gateway gateway-name
```

5. When the subscriber count is zero and all sessions have ended, make and commit changes to the TDF domain interface in active maintenance mode.

> ⓘ **NOTE**: These modifications must be made in active maintenance mode or they fail.

6. Modify the interface.

7. Exit maintenance mode and commit the changes.

```
[edit]
user@host# delete unified-edge gateways tdf gateway-name domain domain-name service-mode
user@host# commit
```

8. Verify that changes were properly committed.

```
[edit]
user@host# run show configuration unified-edge gateways tdf gateway-name domain domain-name
```

9. Return the gateway to operational state.

```
[edit]
user@host# run show unified-edge tdf service-mode
```

RELATED DOCUMENTATION

Deleting a TDF Domain

This procedure describes how to use maintenance mode to delete a TDF domain. You must first halt new sessions from being started and verify that there no active sessions remain.

To delete a TDF domain name:

1. From configuration mode, activate maintenance mode for a TDF domain.

```
[edit]
user@host# set unified-edge gateways tdf  gateway-name domains domain-name service-mode
maintenance
user@host#  commit
```

2. Verify that the TDF domain is in maintenance mode.

```
[edit]
user@host# run show unified-edge tdf domains service-mode
```

The service mode shows `Maintenance - Active Phase` if all the sessions are cleared. The service mode shows `Maintenance - In Phase` if some sessions are active. The service mode shows `Maintenance - Out Phase` if maintenance mode is not configured (that is, it is in operational mode).

3. Verify that no subscribers are active on the TDF domain.

```
user@host# run show unified-edge tdf domain domain-name gateway gateway-name
```

4. (Optional) Terminate sessions that are using a TDF domain using the **clear** command.

```
user@host#  run clear  unified-edge tdf subscribers domain domain-name gateway gateway-name
```

5. When the subscriber count is zero and all sessions have ended, delete the TDF domain in active maintenance mode.

> **NOTE**: These modifications must be made in active maintenance mode or they fail.

6. Delete the TDF domain and commit the changes.

```
user@host# delete unified-edge gateways tdf  gateway-name tdf-services domains domain-name
user@host# commit
```

7. Verify that changes were properly committed by showing the configuration for the entire unified edge to make sure the TDF domain is deleted.

```
[edit]
user@host# run show configuration unified-edge gateways tdf gateway-name domain domain-name
```

8. Return the gateway to the operational state.

```
[edit]
user@host# run show unified-edge tdf gateway service-mode
```

RELATED DOCUMENTATION

Changing a TDF Interface

This procedure describes how to use maintenance mode to halt new sessions from being started and to verify that no active sessions remain before making changes to a TDF interface address.

211

1. From configuration mode, activate maintenance mode for a gateway.

```
[edit]
user@host# set unified-edge gateways tdf gateway-name service-mode maintenance
user@host# commit
```

2. Verify that the TDF gateway is in maintenance mode.

```
[edit]
user@host# run show unified-edge tdf gateway service-mode
```

From the gateway hierarchy, the service mode for the TDF gateway shows Maintenance - Active Phase if all the sessions using this pool are cleared. The service mode for the gateway shows Maintenance - In Phase if some sessions are actively using this pool.

3. Verify that no subscribers are active on this gateway.

```
[edit]
user@host# run show unified-edge tdf subscribers gateway gateway-name
```

> *i* **NOTE**: If a large number of subscribers use this gateway, the preceding command can be process intensive, in which case you can use the following command to show the active contexts across all of the gateway instances:
>
> ```
> [edit]
> user@host# run show unified-edge tdf status
> ```

4. (Optional) Terminate sessions that are using the gateway using the following **clear** command:

```
[edit]
user@host# run clear  unified-edge tdf subscribers gateway  gateway-name
```

> ⚠ **CAUTION**: This clear command deletes all of the existing subscribers on the gateway. Only issue these commands if you intend to disconnect service to all these subscribers.

5. When the subscriber count is zero, and all sessions have ended, modify the TDF interface in active maintenance mode.

```
[edit]
user@host# set unified-edge gateways tdf gateway-name domains domain-name
user@host# commit
```

> **NOTE**: These modifications must be made in active maintenance mode or they fail.

6. Verify that changes were properly committed.

```
[edit]
user@host# run show configuration unified-edge tdf gateway  gateway-name
```

7. Exit maintenance mode and commit the changes.

```
[edit]
user@host# delete unified-edge gateways tdf gateway-name gateway gateway-name service-mode
user@host# commit
```

8. Return the gateway to operational state.

```
[edit]
user@host# run show unified-edge tdf gateway service-mode
```

RELATED DOCUMENTATION

Deleting a TDF Interface

This procedure describes how to use maintenance mode to delete a TDF interface. You must first halt new sessions from being started and verify that no active sessions are remaining.

You can use maintenance mode to remove any of the TDF interfaces.

You can also enter maintenance mode to delete control and data portions of these interface configurations.

1. From configuration mode, activate maintenance mode for a gateway.

```
[edit]
user@host# set unified-edge gateways tdf gateway-name service-mode maintenance
user@host# commit
```

2. Verify that the TDF gateway is in maintenance mode.

```
[edit]
user@host# run show unified-edge tdf gateway service-mode
```

From the gateway hierarchy, the service mode for the gateway shows `Maintenance - Active Phase` if all the sessions using this pool are cleared. The service mode for the gateway shows `Maintenance - In Phase` if some sessions are actively using this pool. The service mode for the gateway shows `Maintenance - Out Phase` if maintenance mode is not configured (that is, the gateway is in operational mode).

3. Verify that no subscribers are active on this gateway.

```
[edit]
user@host# run show unified-edge tdf subscriber gateway gateway-name
```

4. (Optional) Terminate sessions that are using the gateway and clear CDRs using the following **clear** command.

```
[edit]
user@host# run clear  unified-edge tdf subscribers gateway  gateway-name
```

5. When the subscriber count is zero, and all sessions have ended, delete the TDF interface in active maintenance mode.

> *(i)* **NOTE**: These modifications must be made in active maintenance mode or they fail.

6. Delete the TDF interface.

```
[edit]
user@host# delete unified-edge gateways tdf gateway-name domains domain-name tdf-interface
mif  interface-name
```

7. Exit maintenance mode and commit the changes.

```
user@host# delete unified-edge gateways tdf  gateway-name gateway gateway-name service-mode
user@host# commit
```

8. Verify that changes were properly committed.

```
user@host# run show configuration unified-edge tdf gateway gateway-name
```

RELATED DOCUMENTATION

Changing TDF Gateway Parameters with Maintenance Mode

This procedure shows how to change the parameters for a TDF gateway using maintenance mode at the `[edit unified-edge gateways tdf gateway-name]` hierarchy level.

The gateway must be in maintenance mode to change:

- Maximum number of sessions

- Maximum amount of memory and CPU utilization.

Before you change these gateway parameters using maintenance mode:

- Make sure that this change has been coordinated with affected groups and users.

- Make sure that this change is applied to the correct gateway type and name.

To configure maintenance mode for a gateway parameter change:

1. Verify the current status of maintenance mode for the gateway.

   Under normal operating conditions, the service mode is Operational (that is, not in maintenance mode).

   ```
   user@host> show unified-edge tdf gateway-name service-mode
   ```

   The service-mode option displays the information details about maintenance mode as well as status.

   ```
   Maintenance Mode
       MM Active Phase - System is ready to accept configuration changes for all
                     attributes of this object and its sub-hierarchies.
       MM In/Out Phase - System is ready to accept configuration changes only for
                     non-maintenance mode attributes of this object and
                     its sub-hierarchies.


   Gateway Name         Service Mode


   <gateway-name>       Operational
   ```

2. From configuration mode, place the gateway in maintenance mode.

   ```
   [edit]
   user@host# set unified-edge tdf gateway-name service-mode maintenance
   user@host# commit
   ```

3. Verify that the gateway is in active maintenance mode where configuration changes are accepted for this object.

   ```
   [edit]
   user@host> show unified-edge tdf gateway-name service-mode
   ```

   The service-mode option displays the information details about maintenance mode as well as status.

   ```
   Maintenance Mode
       MM Active Phase - System is ready to accept configuration changes for all
                     attributes of this object and its sub-hierarchies.
       MM In/Out Phase - System is ready to accept configuration changes only for
                     non-maintenance mode attributes of this object and
   ```

```
                  its sub-hierarchies.


Gateway Name        Service Mode


<gateway-name>    Maintenance - Active Phase
```

> **NOTE**: All subscribers serviced by the gateway must go to zero. You can wait for these conditions to be met, or use the `clear` command for the gateway to force these conditions.

4. Configure the threshold for the maximum amount of CPU that the TDF gateway can use as a percentage from 1 through 90.

```
[edit unified-edge gateways tdf gateway-name]
user@host# set cac cpu cpu-pct
```

5. Configure the maximum number of TDF subscriber sessions that may be running, expressed in thousands of sessions.

```
[edit unified-edge gateways tdf gateway-name]
user@host# set cac maximum-sessions max-sessions
```

6. Configure the threshold for the maximum amount of memory that the TDF gateway can use as a percentage from 1 through 90.

```
[edit unified-edge gateways tdf gateway-name]
user@host# set cac memory memory-pct
```

7. Exit maintenance mode and commit the changes.

```
[edit]
user@host# delete unified-edge tdf gateway-name service-mode maintenance
user@host# commit
```

RELATED DOCUMENTATION

Changing PCEF Profiles, PCC Rules, PCC Rulebases, Diameter Profiles, Flow Descriptions, and PCC Action Profiles

IN THIS SECTION

- Changing PCEF Profiles, PCC Rules, PCC Rulebases, Diameter Profiles, Flow Descriptions, and PCC Action Profiles with the TDF Domain in Maintenance Mode | 218
- Changing PCEF Profiles, PCC Rules, PCC Rulebases, Diameter Profiles, Flow Descriptions, and PCC Action Profiles with the TDF Gateway in Maintenance Mode | 220

These procedures show how to enter maintenance mode to halt new sessions from being started and verify that no sessions remain on either the gateway or TDF domain before making changes to the following:

- PCEF profiles (However, maintenance mode is not required to add PCC rules or rulebases to a dynamic PCEF profile.)

- PCC rules

- PCC rulebases

- Diameter profiles

- Flow descriptions

- PCC action profiles

> **NOTE**: Even when a PCEF profile is not associated with a TDF domain or a TDF domain-selection term, configuration changes or deletion of the PCEF profile and any referenced objects of the profile require you to activate maintenance mode for the TDF gateway.

Changing PCEF Profiles, PCC Rules, PCC Rulebases, Diameter Profiles, Flow Descriptions, and PCC Action Profiles with the TDF Domain in Maintenance Mode

This procedure shows operators how to enter maintenance mode to halt new sessions from being started and to verify that no sessions remain on the TDF domain before making changes to PCEF profiles, PCC rules, PCC rulebases, Diameter profiles, flow descriptions, and PCC action profiles for a TDF domain.

To activate maintenance mode for the TDF domain and make changes:

1. From configuration mode, activate maintenance mode for the TDF domain.

```
[edit]
user@host# set unified-edge gateways tdf gateway-name domain domain-name service-mode
maintenance
user@host# commit
```

2. Verify that the TDF domain is in maintenance mode.

```
[edit]
user@host# run show unified-edge tdf domain service-mode
```

The service mode for the TDF domain shows `Maintenance-Active Phase` if all the sessions using this TDF domain are cleared. The service mode for the TDF domain shows `Maintenance - In Phase` if some sessions are actively using this TDF domain.

3. Verify that no subscribers are active on the TDF domain.

```
[edit]
user@host# run show unified-edge tdf subscribers | match domain-name
```

4. (Optional) Terminate any remaining sessions on the TDF domain by using the `clear` command.

```
[edit]
user@host# run clear unified-edge tdf subscribers | match domain-name
```

5. Verify that the TDF domain is in Active Phase.

```
[edit]
user@host# run show unified-edge tdf domain service-mode
```

6. Make the configuration changes and commit the changes.

7. Exit maintenance mode.

```
[edit]
user@host# delete unified-edge gateways tdf gateway-name domain domain-name service-mode
user@host# commit
```

8. Verify that changes were properly committed.

- To view a PCEF profile configuration:

```
[edit]
user@host# run show configuration unified-edge pcef profiles profile-name
```

- To view a PCC rulebase configuration:

```
[edit]
user@host# run show configuration unified-edge pcef pcc-rulebases rulebase-name
```

- To view a PCC rules configuration:

```
[edit]
user@host# run show configuration unified-edge pcef pcc-rules rule-name
```

- To view a flow description configuration:

```
[edit]
user@host# run show configuration unified-edge pcef flow-description flow-identifier
```

- To view a PCC action profile configuration:

```
[edit]
user@host# run show configuration unified-edge pcef pcc-action-profiles profile-name
```

9. Return the gateway to operational state.

```
[edit]
user@host# run show unified-edge tdf gateway service-mode
```

Changing PCEF Profiles, PCC Rules, PCC Rulebases, Diameter Profiles, Flow Descriptions, and PCC Action Profiles with the TDF Gateway in Maintenance Mode

This procedure shows how to enter maintenance mode to halt new sessions from being started and to verify that no sessions remain on the TDF gateway before making changes to PCEF profiles, PCC rules, PCC rulebases, Diameter profiles, flow descriptions, and PCC action profiles across multiple TDF domains on the gateway.

To activate maintenance mode for the gateway and make changes:

1. From configuration mode, activate maintenance mode for the gateway.

```
[edit]
user@host# set unified-edge gateways tdf gateway-name service-mode maintenance
user@host# commit
```

2. Verify that the TDF gateway is in maintenance mode.

```
[edit]
user@host# run show unified-edge tdf service-mode
```

From the gateway hierarchy, the service mode shows Maintenance-Active Phase if all the sessions are cleared. The service mode shows Maintenance-In Phase if some sessions are active. The service mode shows Maintenance-Out Phase if maintenance mode is not configured, and the gateway is in operational mode.

3. Make the configuration changes.

You can modify a PCEF profile by making changes to the PCC rules, PCC rulebases, or flow identifiers that the PCEF profile references or by specifying a different PCC rule, rule precedence, PCC rulebase, or Diameter profile in the PCEF profile.

4. Exit maintenance mode and commit the changes.

```
[edit]
user@host# delete unified-edge gateways tdf gateway-name service-mode
user@host# commit
```

5. Verify that changes were properly committed.

- To view a PCEF profile configuration:

```
[edit]
user@host# run show configuration unified-edge pcef profiles profile-name
```

- To view a PCC rulebase configuration:

```
[edit]
user@host# run show configuration unified-edge pcef pcc-rulebases rulebase-name
```

- To view a PCC rules configuration:

```
[edit]
user@host# run show configuration unified-edge pcef pcc-rules rule-name
```

- To view a flow description configuration:

```
[edit]
user@host# run show configuration unified-edge pcef flow-description flow-identifier
```

- To view a PCC action profile configuration:

```
[edit]
user@host# run show configuration unified-edge pcef pcc-action-profiles profile-name
```

6. Return the gateway to operational state.

```
[edit]
user@host# run show unified-edge tdf gateway service-mode
```

SEE ALSO

Deleting a PCEF Profile

IN THIS SECTION

These procedures show how to enter maintenance mode to halt new sessions from being started and verify that no sessions remain on the TDF domain or gateway before removing a policy and charging enforcement function (PCEF) profile from the TDF domain or service-selection profile configurations.

> ⓘ **NOTE**: Regardless of whether a PCEF profile is associated within a TDF domain or not, or whether a PCEF profile is associated with a TDF domain-selection term or not, configuration changes and deletion of a PCEF profile (and other referenced objects of the profile) require that the TDF gateway be placed in maintenance mode. However, you need not activate maintenance mode for the gateway if you are adding a new PCEF profile.

Deleting a PCEF Profile with the TDF Domain in Maintenance Mode

This procedure shows how to enter maintenance mode to halt new sessions from being started and to verify that there are no sessions remaining on the TDF domain before removing a PCEF profile configuration that a TDF domain or service-selection profile references.

To activate maintenance mode for the TDF domain and make changes to a PCEF profile:

1. From configuration mode, activate maintenance mode for the TDF domain that references the PCEF profile.

   ```
   [edit]
   user@host# set unified-edge gateways tdf gateway-name  domains domain-name service-mode
   maintenance
   user@host# commit
   ```

2. Verify that the TDF domain is in maintenance mode.

```
[edit]
user@host# run show unified-edge tdf domain service-mode
```

The service mode for the TDF domain shows `Maintenance-Active Phase` if all the sessions using this TDF domain are cleared. The service mode for the TDF domain shows `Maintenance-In Phase` if some sessions are actively using this TDF domain.

3. Verify that no subscribers are active on the TDF domain.

```
[edit]
user@host# run show unified-edge tdf subscribers | match domain-name
```

4. (Optional) Terminate any remaining sessions on the TDF domain.

```
[edit]
user@host# run clear unified-edge tdf subscribers domain domain-name
```

5. Verify that the TDF domain is in an active phase.

```
[edit]
user@host# run show unified-edge tdf domain service-mode
```

6. In the TDF domain or service-selection profile configuration, remove the referenced PCEF profile and commit the changes.

```
user@host# delete unified-edge gateways tdf gateway-name domains domain-name pcef-profile
pcef-profile-name
```

```
user@host# delete unified-edge gateways tdf gateway-name domain-selection term term-name
then pcef-profile pcef-profile-name
```

7. Verify that the changes were properly committed by showing the configuration for the entire TDF domain or service-selection profile to make sure the PCEF profile is deleted.

- To view a PCEF profile configuration:

```
[edit]
user@host# run show configuration unified-edge pcef profiles profile-name
```

- To view a PCC rulebase configuration:

```
[edit]
user@host# run show configuration unified-edge pcef pcc-rulebases rulebase-name
```

- To view a PCC rules configuration:

```
[edit]
user@host# run show configuration unified-edge pcef pcc-rules rule-name
```

- To view a flow description configuration:

```
[edit]
user@host# run show configuration unified-edge pcef flow-description flow-identifier
```

- To view a PCC action profile configuration:

```
[edit]
user@host# run show configuration unified-edge pcef pcc-action-profiles profile-name
```

8. (Optional) If the PCEF profile is not used in other TDF domain or service-selection profile configurations, you can delete the PCEF profile configuration and commit the changes.

```
[edit]
user@host# delete unified-edge gateways tdf gateway-name  domains domain-name service-mode
user@host# commit
```

9. Exit maintenance mode.

```
[edit]
user@host# delete unified-edge gateways tdf gateway-name service-mode
user@host# commit
```

10. Return the gateway to operational state.

```
user@host# run show unified-edge tdf gateway service-mode
```

Deleting a PCEF Profile with the Gateway in Maintenance Mode

This procedure shows how to enter maintenance mode to halt new sessions from being started and to verify that no sessions remain on the TDF gateway before deleting PCEF profiles that are referenced by one or more TDF domains on a gateway.

To activate maintenance mode for the gateway and make changes to a PCEF profile:

1. From configuration mode, activate maintenance mode for the gateway.

```
[edit]
user@host# set unified-edge gateways tdf gateway-name service-mode maintenance
user@host# commit
```

2. Verify that the TDF gateway is in maintenance mode.

```
[edit]
user@host# run show unified-edge tdf service-mode
```

From the gateway hierarchy, the service mode shows `Maintenance-Active Phase` if all the sessions are cleared. The service mode shows `Maintenance-In Phase` if some sessions are active. The service mode shows `Maintenance-Out Phase` if maintenance mode is not configured, and the gateway is in operational mode.

3. Verify that no subscribers are active on the gateway.

```
[edit]
user@host# run show unified-edge tdf subscribers gateway gateway-name
```

4. (Optional) Terminate any remaining sessions on the gateway.

```
[edit]
user@host# run clear unified-edge tdf subscribers gateway gateway-name
```

5. Verify that the gateway is in an active phase.

```
[edit]
user@host# run show unified-edge tdf gateway service-mode
```

6. For each applicable TDF domain, delete the PCEF profile from the TDF domain configuration and commit the changes.

```
user@host# delete unified-edge gateways tdf gateway-name domains domain-name pcef-profile
pcef-profile-name
user@host# commit
```

7. Verify that the changes were properly committed by showing the configuration for each TDF domain to make sure the PCEF profile is deleted.

- To view a PCEF profile configuration:

```
[edit]
user@host# run show configuration unified-edge pcef profiles profile-name
```

- To view a PCC rulebase configuration:

```
[edit]
user@host# run show configuration unified-edge pcef pcc-rulebases rulebase-name
```

- To view a PCC rules configuration:

```
[edit]
user@host# run show configuration unified-edge pcef pcc-rules rule-name
```

- To view a flow description configuration:

```
[edit]
user@host# run show configuration unified-edge pcef flow-description flow-identifier
```

- To view a PCC action profile configuration:

```
[edit]
user@host# run show configuration unified-edge pcef pcc-action-profiles profile-name
```

8. Exit maintenance mode.

```
[edit]
user@host# delete unified-edge gateways tdf gateway-name service-mode
user@host# commit
```

9. Return the gateway to operational state.

```
[edit]
user@host# run show unified-edge tdf gateway service-mode
```

SEE ALSO

Maintenance Mode Overview for Subscriber Aware Policy Enforcement | **197**

Changing PCEF Profiles, PCC Rules, PCC Rulebases, Diameter Profiles, Flow Descriptions, and PCC Action Profiles | **217**

Changing Static Time-of-Day Settings for PCC Rules

This procedure shows how to enter maintenance mode to make changes to static time-of-day activation and deactivation settings or to apply those settings to PCC rules and rulebases.

To make changes to the static time-of-day activation and deactivation configuration:

1. From configuration mode, activate maintenance mode for the gateway.

```
[edit unified-edge gateways]
user@host# set tdf gateway-name service-mode maintenance
user@host# commit
```

2. Verify that the gateway is in maintenance mode.

```
[edit unified-edge gateways]
user@host# run show unified-edge tdf service-mode
```

The service mode shows `Maintenance-Active Phase` if all the sessions are cleared. The service mode shows `Maintenance-In Phase` if some sessions are active. The service mode shows `Maintenance-Out Phase` if maintenance mode is not configured, and the gateway is in operational mode.

3. Modify the time-of-day profile settings, the assignment of time-of-day profiles to rules and rulebases within a PCEF profile, or both, and commit the changes. See "Configuring Static Time-of-Day PCC Rule Activation and Deactivation in a Junos OS Subscriber Aware PCEF Profile" on page 100.

4. Exit maintenance mode.

```
[edit unified-edge gateways]
user@host# delete tdf gateway-name service-mode
user@host# commit
```

5. Verify that changes were properly committed.

   - To view a PCEF profile configuration:

     ```
     [edit]
     user@host# run show configuration unified-edge pcef profiles profile-name
     ```

   - To view a time-of-day profile configuration:

     ```
     [edit]
     user@host# run show configuration unified-edge pcef pcc-time-of-day-profiles profile-name
     ```

RELATED DOCUMENTATION

Maintenance Mode Overview for Subscriber Aware Policy Enforcement | 197

Deleting a Services PIC

This procedure shows how to delete a services PIC using maintenance mode at the `[edit unified-edge gateways tdf `*`gateway-name`*` system session-pics interface]` hierarchy level. The services PIC can be an aggregated multiservices (AMS) interface. Services PICs perform packet-related services on a broadband gateway.

Before you delete a services PIC using maintenance mode:

- Make sure that this change has been coordinated with affected groups and users.

To configure maintenance mode and services PIC deletion:

1. Verify the current status of maintenance mode for this services PIC.

```
user@host> show unified-edge tdf gateway-name system interfaces service-mode
```

The `service-mode` option displays the information details about maintenance mode as well as status.

```
Maintenance Mode
    MM Active Phase - System is ready to accept configuration changes for all
                   attributes of this object and its sub-hierarchies.
    MM In/Out Phase - System is ready to accept configuration changes only for
                   non-maintenance mode attributes of this object and
                   its sub-hierarchies.

Interface Name     Gateway Name      Service Mode
ms-1/0/0            SCG1              Operational
ms-1/1/0            SCG1              Operational
ms-2/0/0            SCG1              Operational
ms-2/1/0            SCG1              Operational
pfe-0/0/0          SCG1              Operational
pfe-0/1/0          SCG1              Operational
pfe-0/2/0          SCG1              Operational
pfe-0/3/0          SCG1              Operational
ams1               SCG1              Operational
```

2. From configuration mode, place the interface in maintenance mode.

```
[edit]
user@host# set unified-edge gateways tdf gateway-name system session-pics interface interface-
```

```
name service-mode maintenance
user@host# commit
```

3. Verify that the services PIC is in active maintenance mode where configuration changes are accepted for this object and all of its subhierarchies.

```
[edit]
user@host> show unified-edge tdf gateway-name system interfaces service-mode
```

```
Maintenance Mode
   MM Active Phase - System is ready to accept configuration changes for all
                 attributes of this object and its sub-hierarchies.
   MM In/Out Phase - System is ready to accept configuration changes only for
                 non-maintenance mode attributes of this object and
                 its sub-hierarchies.

Interface Name      Gateway Name          Service Mode
ms-1/0/0             SCG1                  Operational
ms-1/1/0             SCG1                  Operational
ms-2/0/0             SCG1                  Maintenance - Active Phase
ms-2/1/0             SCG1                  Operational
pfe-0/0/0            SCG1                  Operational
pfe-0/1/0            SCG1                  Operational
pfe-0/2/0            SCG1                  Operational
pfe-0/3/0            SCG1                  Operational
ams1                 SCG1                  Operational
```

NOTE: All subscribers serviced by the services PIC must go to zero. You can wait for these conditions to be met, or use the `clear` command for the interface (or gateway) to force these conditions.

4. Delete the services PIC, exit maintenance mode, and commit the changes.

> **NOTE**: Deletion of a services PIC automatically exits maintenance mode for the deleted PIC.

```
[edit]
user@host# delete unified-edge gateways tdf gateway-name system interface interface-name
user@host#  commit
```

RELATED DOCUMENTATION

Deleting a Session PIC

This procedure shows how to delete a session PIC using maintenance mode at the `[edit unified-edge gateways tdf gateway-name system session-pics interface]` hierarchy level. The session PIC can be an aggregated multiservices (AMS) interface. Session PICs process control plane messages on a broadband gateway.

Before you delete a session PIC using maintenance mode:

- Make sure that this change has been coordinated with affected groups and users.

To configure maintenance mode and session PIC deletion:

1. Verify the current status of maintenance mode for this session PIC.

   ```
   user@host> show unified-edge tdf gateway-name system interfaces service-mode
   ```

   The `service-mode` option displays the information details about maintenance mode as well as status.

   ```
   Maintenance Mode
      MM Active Phase - System is ready to accept configuration changes for all
   ```

```
                    attributes of this object and its sub-hierarchies.
    MM In/Out Phase - System is ready to accept configuration changes only for
                non-maintenance mode attributes of this object and
                its sub-hierarchies.


Interface Name      Gateway Name         Service Mode
ms-1/0/0              SCG1                 Operational
ms-1/1/0              SCG1                 Operational
ms-2/0/0              SCG1                 Operational
ms-2/1/0              SCG1                 Operational
pfe-0/0/0            SCG1                  Operational
pfe-0/1/0            SCG1                  Operational
pfe-0/2/0            SCG1                  Operational
pfe-0/3/0            SCG1                  Operational
ams1                 SCG1                 Operational
```

2. From configuration mode on the TDF gateway, place the interface in maintenance mode.

```
[edit]
user@host# set unified-edge gateways tdf gateway-name system session-pics interface interface-
name service-mode maintenance
user@host# commit
```

3. Verify that the session PIC is in active maintenance mode where configuration changes are accepted for this object and all of its subhierarchies.

```
user@host> show unified-edge tdf gateway-name system interfaces service-mode
```

```
Maintenance Mode
    MM Active Phase - System is ready to accept configuration changes for all
                attributes of this object and its sub-hierarchies.
    MM In/Out Phase - System is ready to accept configuration changes only for
                non-maintenance mode attributes of this object and
                its sub-hierarchies.


Interface Name      Gateway Name         Service Mode
ms-1/0/0              SCG1                 Operational
ms-1/1/0              SCG1                 Maintenance - Active Phase
ms-2/0/0              SCG1                 Operational
```

```
ms-2/1/0              SCG1                Operational
pfe-0/0/0             SCG1                 Operational
pfe-0/1/0             SCG1                 Operational
pfe-0/2/0             SCG1                 Operational
pfe-0/3/0             SCG1                 Operational
ams1                  SCG1                Operational
```

> **NOTE**: All subscribers serviced by the session PIC must go to zero. You can wait for these conditions to be met, or use the `clear` command for the interface (or gateway) to force these conditions.

4. Delete the session PIC.

```
[edit]
user@host# delete unified-edge gateways tdf gateway-name system interface interface-name
```

5. Exit maintenance mode after committing the changes.

> **NOTE**: Deletion of a session PIC automatically exits maintenance mode for the deleted PIC.

```
[edit]
user@host#  commit
```

RELATED DOCUMENTATION

5
PART

Monitoring and Troubleshooting

CHAPTER 11

Monitoring and Troubleshooting

Configuring Tracing for PCEF Operations

To configure tracing operations for the policy and charging enforcement function (PCEF):

1. Specify that you want to configure tracing options for PCEF.

```
[edit unified-edge pcef]
user@host# edit traceoptions
```

2. (Optional) Configure the name of the file used for the trace output.

```
[edit unified-edge pcef traceoptions]
user@host# set file file-name
```

3. (Optional) Configure the maximum size of each trace file.

```
[edit unified-edge pcef traceoptions]
user@host# set file size size
```

4. (Optional) Configure the maximum number of trace files.

```
[edit unified-edge pcef traceoptions]
user@host# set file files number
```

5. (Optional) Configure the read permissions for the log file.

```
[edit unified-edge pcef traceoptions]
user@host# set file (no-world-readable | world-readable)
```

6. (Optional) Configure flags to filter the operations to be logged.

```
[edit unified-edge pcef traceoptions]
user@host# set flag flag
```

Table 10 on page 236 describes the flags that you can include.

Table 10: Trace Flags

Flag	Description
all	Trace all operations.
config	Trace configuration events.
debug	Trace the debug internal events.
fsm	Trace finite state machine events.
general	Trace general events that do not fit in any specific traces.
high-availability	Trace high availability events.
init	Trace initialization events.
tftmgr	Trace traffic flow manager events.

7. (Optional) Configure the level of tracing.

```
[edit unified-edge pcef traceoptions]
user@host# set level (all | critical | error | info | notice | verbose | warning)
```

traceoptions

Configuring Call-Rate Statistics Collection

You can configure the collection of statistics for the rate of calls for a TDF gateway and for a TDF domain. You configure the length of the interval for statistics collection and the number of call-records to keep.

To configure call-rate statistics collection for the TDF gateway or TDF domain:

1. Configure the length of the interval for statistics collection:

   • For a TDF gateway:

   ```
   [edit unified-edge gateways tdf gateway-name]
   user@host# set call-rate-statistics interval minutes
   ```

   • For a TDF domain:

   ```
   [edit unified-edge gateways tdf gateway-name domains domain-name]
   user@host# set call-rate-statistics interval minutes
   ```

2. Configure the number of call-rate records to save.

   • For a TDF gateway:

   ```
   [edit unified-edge gateways tdf gateway-name]
   user@host# set call-rate-statistics history records
   ```

   • For a TDF domain:

   ```
   [edit unified-edge gateways tdf gateway-name domains domain-name]
   user@host# set call-rate-statistics history records
   ```

   When the number of call-rate records equals the `history` value and a new record is received, the oldest record is replaced by the new record.

show unified-edge tdf call-rate statistics

Using the Enterprise-Specific Utility MIB

Using the Enterprise-Specific Utility MIB

The enterprise-specific Utility MIB enables you to add SNMP-compliant applications information to the enterprise-specific Utility MIB. The application information includes:

- NAT mappings

- Carrier-grade NAT (CGNAT) pools

- Service set CPU utilization

- Service set memory usage

- Service set summary information

- Service set packet drop information

- Service set memory zone information

- Multiservices PIC CPU and memory utilization

- Stateful firewall flow counters

- Session application connection information

- Session analysis information

- Subscriber analysis information

- Traffic Load Balancer information

You use a delivered Stylesheet Language Alternative Syntax (SLAX) script to place applications information into the enterprise-specific Utility MIB. The script is invoked based on event policies (such as reboot of the router or switchover of Routing Engines) defined in an event script. The script can also be invoked from the command line as an op script. The script only runs on the primary Routing Engine. After the script is invoked, it polls data from the specified components at regular intervals using the XML-RPC API and writes the converted data to the Utility MIB as SNMP variables. The script automatically restarts after a configured polling cycle elapses.

Populating the Enterprise-Specific Utility MIB with Information

To use a SLAX script to populate the enterprise-specific Utility MIB with information:

1. Enable the **services-oids-slax** script.

```
user@host# set system scripts op file services-oids.slax
```

2. Configure the maximum amount of memory for the data segment during the execution of the script.

```
user@host# set event-options event-script max-database 512m
```

3. Enable the script.

```
user@host# set event-options event-script file services-oids-ev-policy.slax
```

4. (Optional) Enable the **log-stats** argument to allow sys logging of stateful firewall rate statistics when the event-script is run.

   a. Display the event policies and the arguments that can be used.

   ```
   user@host> show event-options event-scripts polices
   ```

   ```
   event-options {
       policy services-oids-done {
           events system;
           attributes-match {
               system.message matches "Completed polling cycle normally. Exiting";
   ```

```
        }
        then {
            event-script services-oids.slax {
                arguments {
                    max-polls 30;
                    interval 120;
                }
            }
        }
    }
    policy system-started {
        events system;
        attributes-match {
            system.message matches "Starting of initial processes complete";
        }
        then {
            event-script services-oids.slax {
                arguments {
                    max-polls 30;
                    interval 120;
                }
            }
        }
    }
}
event-options {
    policy services-oids-done {
        events system;
        attributes-match {
            system.message matches "Completed polling cycle normally. Exiting";
        }
        then {
            event-script services-oids.slax {
                arguments {
                    max-polls 30;
                    interval 120;
                }
            }
        }
    }
    policy system-started {
        events system;
        attributes-match {
```

```
                system.message matches "Starting of initial processes complete";
        }
        then {
            event-script services-oids.slax {
                arguments {
                    max-polls 30;
                    interval 120;
                }
            }
        }
    }
}
```

The `log-stats` argument does not appear, so you must enable it.

b. Start the Linux shell.

```
user@host>  start shell
```

```
%
```

c. Open the **/var/db/scripts/event/services-oids-eve-policy.slax** file for editing.

```
<event-options> {
        /*
         * This policy detects when the services-oids.slax script ends, then restarts it.
         */
        <policy> {
            <name> "services-oids-done";
            <events> "system";
            <attributes-match> {
                <from-event-attribute> "system.message";
                <condition> "matches";
                <to-event-attribute-value> "Completed polling cycle normally. Exiting";
            }
            <then> {
                <event-script> {
                    <name> "services-oids.slax";
                    <arguments> {
                        <name>"max-polls";
```

```
                                    <value>"30";
                                }
                                <arguments> {
                                    <name>"interval";
                                    <value>"120";
                                }
                                /*
                                <arguments> {
                                    <name>"log-stats";
                                    <value>"yes";
                                }
                                */
                            }
                        }
                    }


        /*
         * This policy detects when the system has booted and kicks off the services-
oids.slax script.
         * This policy hooks the 'system started' event
         */
        <policy> {
            <name> "system-started";
            <events> "system";
            <attributes-match> {
                <from-event-attribute> "system.message";
                <condition> "matches";
                <to-event-attribute-value> "Starting of initial processes complete";
            }
            <then> {
                <event-script> {
                    <name> "services-oids.slax";
                    <arguments> {
                        <name>"max-polls";
                        <value>"30";
                    }
                    <arguments> {
                        <name>"interval";
                        <value>"120";
                    }
                    /*
                    <arguments> {
                        <name>"log-stats";
```

```
                        <value>"yes";
                }
                */
            }
        }
    }


}
```

d. Remove the comment enclosures (/* and */) surrounding the `<arguments>` tags containing "`log-stats`".

e. Exit the Linux shell and return to the CLI.

```
% exit
```

f. Load the changes you made to the event script file.

```
user@host>request system scripts event-scripts reload
```

The `log-stats` argument is available the next time the event script restarts.

5. Set up the script logging file **services-oids.log**.

```
user@host# set system syslog file services-oids.log any info
user@host# set system syslog file services-oids.log match cscript
```

6. Synchronize scripts between Routing Engines so that when a switchover of Routing Engine occurs, the event policy starts on the new primary.

- To synchronize on a per-commit basis:

```
user@host# commit synchronize scripts
```

- To synchronize scripts every time you execute a **commit synchronize**:

```
[edit system scripts]
user@host# set synchronize
user@host# commit synchronize
```

7. The script starts automatically at system boot, but you can manually start it with the CLI.

```
user@host> op services-oids arguments
```

describes the arguments that you can use.

Table 11: Arguments for services-oids.slax Script

Argument	Description
clean	A value of **1** clears all Utility MIB OIDs. Use this only to clean OID tables.
clear-semaphore	A value of **1** resets the semaphore in the Utility MIB to recover from an abnormal script exit or from a manual script exit.
debug	Prints debug messages on console.
detail	Displays detailed output.
interval	Sets the number of seconds between poll cycles (default is 120).
invoke-debugger	Invokes script in debugger mode.
log-stats	**Yes** value enables sys logging of stateful firewall rate statistics (default is no).
max-polls	Sets the number of poll cycles before exiting the script (default is 30).

Table 11: Arguments for services-oids.slax Script *(Continued)*

Argument	Description
one-cycle-only	Value of **1** quits after one cycle of polling. Event policy does not restart the script. Use this option for testing only. The default is **0**.
signal-stop	A value of **1** stops the script and sets the semaphore, which causes the next iteration to exit.
silent	Prints trace messages on console if it is unset. Set it to zero-length string (" ") to unset it. Default is 1.
|	Pipes through a command.

8. Check the status of the script from the log file.

```
router> show /var/log/services-oids.log | no-more
```

```
Jun 27 19:51:47  wf-cheesypoofs cscript: services-oids.slax(v0.14):[info] Beginning polling
cycle.
Jun 27 19:51:47  wf-cheesypoofs cscript: services-oids.slax(v0.14):[info]  processing traffic
load-balance statistics
Jun 27 19:51:48  wf-cheesypoofs cscript: services-oids.slax(v0.14):[info]  processing cgnat
pool detail
Jun 27 19:51:48  wf-cheesypoofs cscript: services-oids.slax(v0.14):[info]  processing cgnat
mappings summary
Jun 27 19:51:48  wf-cheesypoofs cscript: services-oids.slax(v0.14):[info]  processing service-
sets summary
Jun 27 19:51:48  wf-cheesypoofs cscript: services-oids.slax(v0.14):[info]  processing service-
sets cpu-usage
Jun 27 19:51:48  wf-cheesypoofs cscript: services-oids.slax(v0.14):[info]  processing service-
sets mem-usage
Jun 27 19:51:49  wf-cheesypoofs cscript: services-oids.slax(v0.14):[info]  processing
stateful firewall statistics
Jun 27 19:51:49  wf-cheesypoofs cscript: services-oids.slax(v0.14):[info]  processing
stateful firewall flow-analysis
Jun 27 19:51:49  wf-cheesypoofs cscript: services-oids.slax(v0.14):[info]  processing
```

```
stateful firewall flows counts
Jun 27 19:51:49  wf-cheesypoofs cscript: services-oids.slax(v0.14):[info]  processing FW
policy connections/second
Jun 27 19:51:49  wf-cheesypoofs cscript: services-oids.slax(v0.14):[info]  processing FW/NAT
app connections
Jun 27 19:51:51  wf-cheesypoofs cscript: services-oids.slax(v0.14):[info]  processing service-
set packet-drops
Jun 27 19:51:51  wf-cheesypoofs cscript: services-oids.slax(v0.14):[info]  processing service-
set memory-usage zone
Jun 27 19:51:51  wf-cheesypoofs cscript: services-oids.slax(v0.14):[info]  processing service-
set policy throughput stats
Jun 27 19:51:52  wf-cheesypoofs cscript: services-oids.slax(v0.14):[info]  processing ms-pic
CPU amd Memory utilization stats
Jun 27 19:51:52  wf-cheesypoofs cscript: services-oids.slax(v0.14):[info] 1/30 Sleeping for
110 seconds.
```

9. Verify that you are getting Utility MIB OID updates.

```
router> show snmp mib walk jnxUtil ascii
```

```
. . .
jnxUtilCounter64Value."services10tcp-errors09CGN-SET-1" = 0
jnxUtilCounter64Value."services10tcp-errors09CGN-SET-2" = 0
jnxUtilCounter64Value."services10tcp-errors09CGN-SET-3" = 0
jnxUtilCounter64Value."services10udp-errors09CGN-SET-1" = 1119
jnxUtilCounter64Value."services10udp-errors09CGN-SET-2" = 0
. . .
```

To exclude the timestamp information, use

```
router> show snmp mib walk jnxUtil ascii | match Value
```

Stopping the SLAX Script with the CLI

To stop the SLAX script from the CLI:

- Issue the stop argument.

```
user@host> op services-oids signal-stop 1
```

Clearing the Utility MIB

To clear all the utility MIB OIDs:

- Issue the clean argument.

```
user@host> op services-oids clean 1
```

Recovering from an Abnormal SLAX Script Exit or a SLAX Script Exit with the CLI

To recover from an abnormal SLAX script exit or an SLAX script exit with the CLI:

- Issue the clear semaphore argument.

```
user@host> op services-oids clear-semaphore 1
```

RELATED DOCUMENTATION

SLAX Overview

6
PART

Configuration Statements and Operational Commands

Junos CLI Reference Overview

We've consolidated all Junos CLI commands and configuration statements in one place. Read this guide to learn about the syntax and options that make up the statements and commands. Also understand the contexts in which you'll use these CLI elements in your network configurations and operations.

- Junos CLI Reference

Click the links to access Junos OS and Junos OS Evolved configuration statement and command summary topics.

- Configuration Statements

- Operational Commands