Junos® OS Release 12.1X46 Feature Guide

Junos OS Release 12.1X46-D30
8 January 2015
Revision 4

This feature guide accompanies Junos OS Release 12.1X46-D30. This guide contains detailed information about new or enhanced functionality introduced in Junos OS Release 12.1X46-D30, Junos OS Release 12.1X46-D25, Junos OS Release 12.1X46-D20, and Junos OS Release 12.1X46-D15, that is summarized in the Release Notes.

Contents

New Features in Junos OS Release 12.1X46-D30

Junos OS Release 12.1X46-D30 introduces the following feature:

Application Layer Gateways (ALGs)

This topic includes the following section:

MS-RPC ALG and Sun RPC ALG Map Table Scaling

Understanding Microsoft RPC Services

MS RPC is the Microsoft implementation of the Distributed Computing Environment (DCE) RPC. Like the Sun RPC, the MS RPC provides a way for a program running on one host to call procedures in a program running on another host. The MS RPC is dynamically negotiated based on the service program's universal unique identifier (UUID). The specific UUID is mapped to a transport address.

Predefined Microsoft remote procedure call (MS RPC) services include:

- **junos-ms-rpc-epm**
- **junos-ms-rpc-tcp**
- **junos-ms-rpc-udp**

MS RPC application defaults include:

- **junos-ms-rpc-iis-com-1**
- **junos-ms-rpc-iis-com-adminbase**
- **junos-ms-rpc-msexchange-directory-nsp**
- **junos-ms-rpc-msexchange-directory-rfr**
- **junos-ms-rpc-msexchange-info-store**
- **junos-ms-rpc-uuid-any-tcp**
- **junos-ms-rpc-uuid-any-udp**
- **junos-ms-rpc-wmic-admin**

- junos-ms-rpc-wmic-admin2
- junos-ms-rpc-wmic-mgmt
- junos-ms-rpc-wmic-webm-callresult
- junos-ms-rpc-wmic-webm-classobject
- junos-ms-rpc-wmic-webm-level1login
- junos-ms-rpc-wmic-webm-login-clientid
- junos-ms-rpc-wmic-webm-login-helper
- junos-ms-rpc-wmic-webm-objectsink
- junos-ms-rpc-wmic-webm-refreshing-services
- junos-ms-rpc-wmic-webm-remote-refresher
- junos-ms-rpc-wmic-webm-services
- junos-ms-rpc-wmic-webm-shutdown

MS RPC application-set defaults include:

- junos-ms-rpc
- junos-ms-rpc-any
- junos-ms-rpc-iis-com
- junos-ms-rpc-msexchange
- junos-ms-rpc-wmic

Table 1 on page 4 lists predefined MS RPC services, UUID values associated with each service, and a description of each service.

Table 1: Predefined MS RPC services

Service	UUID	Description
EPM	e1af8308-5d1f-11c9-91a4-08002b14a0fa	MS RPC Endpoint Mapper (EPM) protocol is a TCP/UDP port-based service that includes TCP/UDP port 135.
EXCHANGE-DATABASE	1a190310-bb9c-11cd-90f8-00aa00466520	Microsoft Exchange Database service.
EXCHANGE-DIRECTORY	f5cc5a18-4264-101a-8c59-08002b2f8426 f5cc5a7c-4264-101a-8c59-08002b2f8426 f5cc59b4-4264-101a-8c59-08002b2f8426	Microsoft Exchange Directory service.
WIN-DNS	50abc2a4-574d-40b3-9d66-ee4fd5fba076	Microsoft Windows DNS server.

Table 1: Predefined MS RPC services *(continued)*

Service	UUID	Description
WINS	5f52c28-7f9f-101a-b52b-08002b2efabe 811109bf-a4e1-11d1-ab54-00a0c91e9b45	Microsoft WINS service.
WMIC-Webm-Level1Login	f309ad18-d86a-11d0-a075-00c04fb68820	This service allows users to connect to the management services interface in a particular namespace.

Related Documentation

- *RPC ALG Feature Guide for Security Devices*
- *Understanding Microsoft RPC ALGs*
- *Customizing Microsoft RPC Applications (CLI Procedure)*
- Understanding Sun RPC Services on page 5

Understanding Sun RPC Services

Sun RPC, also known as Open Network computing remote procedure call (ONC RPC), provides a way for a program running on one host to call procedures in a program running on another host. Sun RPC services are defined by a program identifier. The program identifier is independent of any transport address, and most of the Sun RPC sessions are initiated through TCP or UDP port 111. Each host links the required RPC service to a dynamic TCP or UDP port that is negotiated over the port 111 control channel, allowing the client to connect to either TCP or UDP port 111.

Predefined Sun Microsystems remote procedure call (Sun RPC) services include:

- **junos-sun-rpc-tcp**
- **junos-sun-rpc-udp**

The Sun RPC ALG can be applied by using the following methods:

- ALG default application—Use one of the following predefined applications for control and data connections in your policy:
  - **junos-sun-rpc-any-tcp**
  - **junos-sun-rpc-any-udp**
  - **junos-sun-rpc-mountd-tcp**
  - **junos-sun-rpc-mountd-udp**
  - **junos-sun-rpc-nfs-tcp**
  - **junos-sun-rpc-nfs-udp**
  - **junos-sun-rpc-nlockmgr-tcp**

- junos-sun-rpc-nlockmgr-udp
- junos-sun-rpc-portmap-tcp
- junos-sun-rpc-portmap-udp
- junos-sun-rpc-rquotad-tcp
- junos-sun-rpc-rquotad-udp
- junos-sun-rpc-ruserd-tcp
- junos-sun-rpc-ruserd-udp
- junos-sun-rpc-sadmind-tcp
- junos-sun-rpc-sadmind-udp
- junos-sun-rpc-sprayd-tcp
- junos-sun-rpc-sprayd-udp
- junos-sun-rpc-status-tcp
- junos-sun-rpc-status-udp
- junos-sun-rpc-walld-tcp
- junos-sun-rpc-walld-udp
- junos-sun-rpc-ypbind-tcp
- junos-sun-rpc-ypbind-udp
- junos-sun-rpc-ypserv-tcp
- junos-sun-rpc-ypserv-udp

- Default control application—Use the predefined control through **junos-sun-rpc**:
  - Create an application for data (**USER_DEFINED_DATA**). You can make a set of your own data (for example, **my_rpc_application_set**) and use it in the policy.
  - ALG default application set—Use the predefined application set for control and customized data application in the policy:
    - **junos-sun-rpc** (for control sessions)
    - **junos-sun-rpc-any**
    - **junos-sun-rpc-mountd**
    - **junos-sun-rpc-nfs**
    - **junos-sun-rpc-nfs-access**
    - **junos-sun-rpc-nlockmgr**
    - **junos-sun-rpc-portmap** (for data sessions)
    - **junos-sun-rpc-rquotad**
    - **junos-sun-rpc-ruserd**

- **junos-sun-rpc-sadmind**
- **junos-sun-rpc-sprayd**
- **junos-sun-rpc-status**
- **junos-sun-rpc-walld**
- **junos-sun-rpc-ypbind**
- **junos-sun-rpc-ypserv**

- Custom control and custom data application—Use a customized application:
  - Create an application for control (**USER_DEFINED_CONTROL**) and data (**USER_DEFINED_DATA**).
  - In the policy, use the user-defined application set for a control and customized data application:
    - **USER_DEFINED_CONTROL**
    - **USER_DEFINED_DATA**

Table 2 on page 7 lists predefined Sun RPC services, a program identifier associated with each service, and a description of each service.

Table 2: Predefined Sun RPC Services

Service	Program ID	Description
PORTMAP	100000	Sun RPC Portmapper protocol is a TCP or UDP port-based service that includes TCP or UDP port 111.
NFS	100003	Sun RPC Network File System.
MOUNT	100005	Sun RPC mount process.
YPBIND	100007	Sun RPC Yellow Page Bind service.
STATUS	100024	Sun RPC status.

Related Documentation

- *RPC ALG Feature Guide for Security Devices*
- *Understanding Sun RPC ALGs*
- *Customizing Sun RPC Applications (CLI Procedure)*
- Understanding Microsoft RPC Services on page 3

map-entry-timeout

Syntax map-entry-timeout *map-entry-timeout*;

Hierarchy Level [edit security alg msrpc]
[edit security alg sunrpc]

Release Information Statement introduced in Junos OS Release 12.3X48-D10.

Description Configure the mapping entry timeout value. When the incoming traffic hits the mapping entry, the timeout value has been reset to configured value. The mapping entry is removed from the table when the timeout value expires. The lifetime of the mapping entry is global and applies to all entries in the table.

Options ***map-entry-timeout***—Specify the Microsoft remote procedure call Application Layer Gateway (MS-RPC ALG) or Sun RPC ALG mapping entry timeout value in hours.

Range: 8 through 72 hours.

Default: 32 hours.

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

Related Documentation

- *RPC ALG Feature Guide for Security Devices*

msrpc

Syntax

```
msrpc {
  disable;
  map-entry-timeout map-entry-timeout;
  traceoptions {
    flag {
      all <extensive>;
    }
  }
}
```

Hierarchy Level [edit security alg]

Release Information Statement introduced in Junos OS Release 9.0.

Description Specify the Microsoft remote procedure call Application Layer Gateway (MS-RPC ALG) on the device.

Options

- **disable**—Disable the MS-RPC ALG. By default, the MS-RPC ALG is enabled.
- **map-entry-timeout** ***map-entry-timeout***—Specify the MS-RPC ALG mapping entry timeout value in hours.

  Range: 8 through 72 hours.

  Default: 32 hours.
- **traceoptions**—Configure the MS-RPC ALG tracing options.
  - **flag**—Trace operation to perform.
    - **all**—Trace all events.
    - ***extensive***—Display extensive amount of data.

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

Related Documentation

- *ALG Overview Feature Guide for Security Devices*

sunrpc

Syntax

```
sunrpc {
  disable;
  map-entry-timeout map-entry-timeout;
  traceoptions {
    flag {
      all <extensive>;
    }
  }
}
```

Hierarchy Level [edit security alg]

Release Information Statement introduced in Junos OS Release 9.0.

Description Specify the Sun Microsystems remote procedure call (RPC) ALG on the device.

Options

- **disable**—Disable the Sun RPC ALG. By default, the Sun RPC ALG is enabled.
- **map-entry-timeout** ***map-entry-timeout***—Specify the Sun RPC ALG mapping entry timeout value in hours.

  Range: 8 through 72 hours.

  Default: 32 hours.
- **traceoptions**—Configure the Sun RPC ALG tracing options.
  - **flag**—Trace operation to perform.
    - **all**—Trace all events.
    - ***extensive***—Display extensive amount of data.

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

Related Documentation

- *ALG Overview Feature Guide for Security Devices*

[edit security alg] Hierarchy Level

```
security {
  alg {
    alg-manager {
      traceoptions {
        flag {
          all <extensive>;
        }
      }
    }
    alg-support-lib {
      traceoptions {
        flag {
```

```
        all <extensive>;
      }
    }
  }
  dns {
    disable;
    doctoring (none | sanity-check);
    maximum-message-length number;
    traceoptions {
      flag {
        all <extensive>;
      }
    }
  }
  ftp {
    allow-mismatch-ip-address;
    disable;
    ftps-extension;
    line-break-extension;
    traceoptions {
      flag {
        all <extensive>;
      }
    }
  }
  h323 {
    application-screen {
      message-flood {
        gatekeeper {
          threshold rate;
        }
      }
      unknown-message {
        permit-nat-applied;
        permit-routed;
      }
    }
    disable;
    dscp-rewrite {
      code-point string;
    }
    endpoint-registration-timeout value-in-seconds;
    media-source-port-any;
    traceoptions {
      flag flag <detail | extensive | terse>;
    }
  }
  ike-esp-nat {
    enable;
    esp-gate-timeout value-in-seconds;
    esp-session-timeout value-in-seconds;
    state-timeout value-in-seconds;
    traceoptions {
      flag {
        all <extensive>;
      }
```

```
    }
  }
  mgcp {
    application-screen {
      connection-flood {
        threshold rate;
      }
      message-flood {
        threshold rate;
      }
      unknown-message {
        permit-nat-applied;
        permit-routed;
      }
    }
    disable;
    dscp-rewrite {
      code-point string;
    }
    inactive-media-timeout value-in-seconds;
    maximum-call-duration value-in-minutes;
    traceoptions {
      flag flag <extensive>;
    }
    transaction-timeout value-in-seconds;
  }
  msrpc {
    disable;
    map-entry-timeout map-entry-timeout;
    traceoptions {
      flag {
        all <extensive>;
      }
    }
  }
  pptp {
    disable;
    traceoptions {
      flag {
        all <extensive>;
      }
    }
  }
  real {
    disable;
    traceoptions {
      flag {
        all <extensive>;
      }
    }
  }
  rsh {
    disable;
    traceoptions {
      flag {
        all <extensive>;
```

```
        }
    }
}
rtsp {
    disable;
    traceoptions {
        flag {
            all <extensive>;
        }
    }
}
sccp {
    application-screen {
        call-flood {
            threshold rate;
        }
        unknown-message {
            permit-nat-applied;
            permit-routed;
        }
    }
    disable;
    dscp-rewrite {
        code-point string;
    }
    inactive-media-timeout value-in-seconds;
    traceoptions {
        flag flag <extensive>;
    }
}
sip {
    application-screen {
        protect {
            deny {
                all {
                    timeout value-in-seconds;
                }
                destination-ip address;
                timeout value-in-seconds;
            }
        }
        unknown-message {
            permit-nat-applied;
            permit-routed;
        }
    }
    c-timeout value-in-minutes;
    disable;
    dscp-rewrite {
        code-point string;
    }
    inactive-media-timeout value-in-seconds;
    maximum-call-duration value-in-minutes;
    retain-hold-resource;
    t1-interval value-in-milliseconds;
    t4-interval value-in-seconds;
```

```
        traceoptions {
          flag flag <detail | extensive | terse>;
        }
      }
      sql {
        disable;
        traceoptions {
          flag {
            all <extensive>;
          }
        }
      }
      sunrpc {
        disable;
        map-entry-timeout map-entry-timeout;
        traceoptions {
          flag {
            all <extensive>;
          }
        }
      }
      talk {
        disable;
        traceoptions {
          flag {
            all <extensive>;
          }
        }
      }
      tftp {
        disable;
        traceoptions {
          flag {
            all <extensive>;
          }
        }
      }
      traceoptions {
        file {
          filename;
          files number;
          match regular-expression;
          (no-world-readable | world-readable);
          size maximum-file-size;
        }
        level (brief | detail | extensive | verbose);
        no-remote-trace;
      }
    }
  }
```

Related Documentation

- *Security Configuration Statement Hierarchy*
- *ALG Overview Feature Guide for Security Devices*

New Features in Junos OS Release 12.1X46-D25

Junos OS Release 12.1X46-D25 introduces the following features:

General Packet Radio Service

This topic includes the following section:

GPRS Tunneling Protocol (GTP)

- show security gprs gtp gsn statistics

show security gprs gtp gsn statistics

Syntax show security gprs gtp gsn statistics

Release Information Command introduced in Junos OS Release 12.1X46-D25.

Description Display a brief summary of GPRS support node (GSN) statistics, including active GSNs, obsolete GSNs, and the usage rate of each SPU.

Required Privilege Level view

Related Documentation
- *show security gprs gtp counters path-rate-limit*
- *General Packet Radio Service Feature Guide for Security Devices*

List of Sample Output show security gprs gtp gsn statistics on page 16

Sample Output

show security gprs gtp gsn statistics

```
user@host> show security gprs gtp gsn statistics
FPC 1 PIC 0:

Active GSNs: 0 Obsolete GSNs: 0 Use rate: 0%

FPC 2 PIC 0:

Active GSNs: 0 Obsolete GSNs: 0 Use rate: 0%
```

Security Logging

TCP/TLS Support for Real-Time Logging

log (Security)

Syntax

```
log {
  cache {
    exclude exclude-name {
      destination-address destination-address;
      destination-port destination-port;
      event-id event-id;
      failure;
      interface-name interface-name;
      policy-name policy-name;
      process process-name;
      protocol protocol;
      source-address source-address;
      source-port source-port;
      success;
      user-name user-name;
    }
    limit value;
  }
  disable;
  event-rate rate;
  file {
    files max-file-number;
    name file-name;
    path binary-log-file-path;
    size maximum-file-size;
  }
  format (binary | sd-syslog | syslog);
  mode (event | stream);
  rate-cap rate-cap-value;
  (source-address source-address | source-interface interface-name);
  stream stream-name {
    category (all | content-security);
    format (binary | sd-syslog | syslog | welf);
    host {
      ip-address;
      port port-number;
    }
    severity (alert | critical | debug | emergency | error | info | notice | warning);
  }
  traceoptions {
    file {
      filename;
      files number;
      match regular-expression;
      size maximum-file-size;
      (world-readable | no-world-readable);
    }
    flag flag;
    no-remote-trace;
  }
  transport {
    protocol (udp | tcp | tls);
```

```
    tls-profile tls-profile-name;
    tcp-connections tcp-connections;
  }
  utc-time-stamp;
}
```

Hierarchy Level [edit security]

Release Information Statement introduced in Junos OS Release 9.2.
Support for the **source-interface** option added in Junos OS Release 12.1X46-D25.

Description You can set the mode of logging (event for traditional system logging or stream for streaming security logs through a revenue port to a server). You can also specify all the other parameters for security logging.

Options **disable**—Disable the security logging for the device.

event-rate *rate*—Limit the rate (0 through 1500) at which logs will be streamed per second.

rate-cap *rate-cap-value*—Works with event mode only. Limit the rate (0 through 5000) at which data plane logs will be generated per second.

source-address *source-address*—Specify a source IP address or IP address used when exporting security logs.

source-interface *interface-name*—Specify a source interface name, which is mandatory to configure **stream**.

NOTE: The **source-address** and **source-interface** are alternate values. Using one of the options is mandatory.

utc-time-stamp—Specify to use UTC time for security log timestamps.

The remaining statements are explained separately.

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

Related Documentation

- *Application Tracking Feature Guide for Security Devices*
- *Master Administrator for Logical Systems Feature Guide for Security Devices*

[edit security log] Hierarchy Level

```
security {
  log {
    cache {
      exclude exclude-name {
        destination-address destination-address;
        destination-port destination-port;
        event-id event-id;
```

```
            failure;
            interface-name interface-name;
            policy-name policy-name;
            process process-name;
            protocol protocol;
            source-address source-address;
            source-port source-port;
            success;
            user-name user-name;
          }
          limit value;
        }
        disable;
        event-rate rate;
        file {
          files max-file-number;
          name file-name;
          path binary-log-file-path;
          size maximum-file-size;
        }
        format (binary | sd-syslog | syslog);
        mode (event | stream);
        source-address source-address | source-interface interface-name;
        stream stream-name {
          category (all | content-security);
          format (binary | sd-syslog | syslog | welf);
          host {
            ip-address;
            port port-number;
          }
          severity (alert | critical | debug | emergency | error | info | notice | warning);
        }
        traceoptions {
          file {
            file-name;
            files max-file-number;
            match regular-expression;
            (no-world-readable | world-readable);
            size maximum-file-size;
          }
          flag flag;
          no-remote-trace;
        }
        transport {
          protocol (udp | tcp | tls);
          tls-profile tls-profile-name;
          tcp-connections tcp-connections;
        }
        utc-time-stamp;
      }
    }
```

Related Documentation

- *Security Configuration Statement Hierarchy*
- *Application Tracking Feature Guide for Security Devices*

- *Master Administrator for Logical Systems Feature Guide for Security Devices*

transport (Security Log)

Syntax

```
transport {
  protocol (udp | tcp | tls);
  tls-profile tls-profile-name;
  tcp-connections tcp-connections;
}
```

Hierarchy Level [edit security log]

Release Information Statement introduced in Junos OS Release 12.1X46-D25.

Description Configure security log transport options.

Options **protocol**—Specify the type of transport protocol to be used to log the data.

- **UDP**—Set the transport protocol to UDP.
- **TCP**—Set the transport protocol to TCP.
- **TLS**—Set the transport protocol to TLS.

tls-profile *tls-profile-name*—Specify the TLS profile name.

tcp-connections *tcp-connections*—Specify the number of TCP connections per SPU.

Range: 1 through 5.

Default: 1.

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

Related Documentation

- *Application Tracking Feature Guide for Security Devices*

System Log Messages

RTLOG System Log Messages

RTLOG_CONN_OPEN

System Log Message args stream-name transport-proto source-address source-port destination-address destination-port;

Description RTLOG connection was established.

Type Event: This message reports an event, not an error

RTLOG_CONN_CLOSE

System Log Message args stream-name transport-proto source-address source-port destination-address destination-port;

Description RTLOG connection was closed.

Type Event: This message reports an event, not an error

RTLOG_CONN_ERROR

System Log Message args stream-name error-message;

Description RTLOG connection was aborted.

Type Event: This message reports an event, not an error

New Features in Junos OS Release 12.1X46-D20

Junos OS Release 12.1X46-D20 introduces the following new or enhanced functionality:

- Chassis Cluster on page 22
- Public Key Infrastructure (PKI) on page 40
- Routing on page 61
- UTM on page 75
- VPNs on page 79

Chassis Cluster

This topic includes the following sections:

Autorecovery of fabric link

Understanding the Chassis Cluster Data Plane

The data plane software, which operates in active/active mode, manages flow processing and session state redundancy and processes transit traffic. All packets belonging to a particular session are processed on the same node to ensure that the same security treatment is applied to them. The system identifies the node on which a session is active and forwards its packets to that node for processing. (After a packet is processed, the Packet Forwarding Engine transmits the packet to the node on which its egress interface exists if that node is not the local one.)

To provide for session (or flow) redundancy, the data plane software synchronizes its state by sending special payload packets called runtime objects (RTOs) from one node to the other across the fabric data link. By transmitting information about a session between the nodes, RTOs ensure the consistency and stability of sessions if a failover were to occur, and thus they enable the system to continue to process traffic belonging to existing sessions. To ensure that session information is always synchronized between the two nodes, the data plane software gives RTOs transmission priority over transit traffic.

Understanding Session RTOs

The data plane software creates RTOs for UDP and TCP sessions and tracks state changes. It also synchronizes traffic for IPv4 pass-through protocols such as Generic Routing Encapsulation (GRE) and IPsec.

RTOs for synchronizing a session include:

- Session creation RTOs on the first packet
- Session deletion and age-out RTOs
- Change-related RTOs, including:
  - TCP state changes
  - Timeout synchronization request and response messages

- RTOs for creating and deleting temporary openings in the firewall (pinholes) and child session pinholes

Understanding Data Forwarding

For Junos OS, flow processing occurs on a single node on which the session for that flow was established and is active. This approach ensures that the same security measures are applied to all packets belonging to a session.

A chassis cluster can receive traffic on an interface on one node and send it out to an interface on the other node. (In active/active mode, the ingress interface for traffic might exist on one node and its egress interface on the other.)

This traversal is required in the following situations:

- When packets are processed on one node, but need to be forwarded out an egress interface on the other node
- When packets arrive on an interface on one node, but must be processed on the other node

  If the ingress and egress interfaces for a packet are on one node, but the packet must be processed on the other node because its session was established there, it must traverse the data link twice. This can be the case for some complex media sessions, such as voice-over-IP (VoIP) sessions.

Understanding Fabric Data Link Failure and Recovery

NOTE: Intrusion Detection and Prevention (IDP) services do not support failover. For this reason, IDP services are not applied for sessions that were present prior to the failover. IDP services are applied for new sessions created on the new primary node.

The fabric data link is vital to the chassis cluster. If the link is unavailable, traffic forwarding and RTO synchronization are affected, which can result in loss of traffic and unpredictable system behavior.

To eliminate this possibility, Junos OS uses fabric monitoring to check whether the fabric link, or the two fabric links in the case of a dual fabric link configuration, are alive by periodically transmitting probes over the fabric links. If Junos OS detects fabric faults, RG1+ status of the secondary node changes to ineligible. It determines that a fabric fault has occurred if a fabric probe is not received but the fabric interface is active.

To recover from this state, you must reboot the disabled node. When you reboot it, the node synchronizes its state and RTOs with the primary node.

NOTE: If you make any changes to the configuration while the secondary node is disabled, execute the commit command to synchronize the configuration after you reboot the node. If you did not make configuration changes, the configuration file remains synchronized with that of the primary node.

NOTE: Starting with Junos OS Release 12.1X46-D20, the fabric monitoring feature is enabled by default on high-end SRX Series devices.

Starting with Junos OS Release 12.1X46-D20, recovery of the fabric link and synchronization take place automatically.

When both the primary and secondary nodes are healthy (that is, there are no failures) and the fabric link goes down, RG1+ redundancy group(s) on the secondary node becomes ineligible. When one of the nodes is unhealthy (that is, there is a failure), RG1+ redundancy group(s) on this node (either the primary or secondary node) becomes ineligible. When both nodes are unhealthy and the fabric link goes down, RG1+ redundancy group(s) on the secondary node becomes ineligible. When the fabric link comes up, the node on which RG1+ became ineligible performs a cold synchronization on all Services Processing Units and transitions to active standby.

NOTE: Only RG1+ transitions to an ineligible state. RG0 continues to be in either a primary or secondary state.

Use the **show chassis cluster interfaces** CLI command to verify the status of the fabric link.

Related Documentation

- *Chassis Cluster Feature Guide for Security Devices*
- *Understanding Chassis Cluster Dual Fabric Links*
- *Example: Configuring the Chassis Cluster Fabric*
- *Verifying Chassis Cluster Data Plane Interfaces*
- *Verifying Chassis Cluster Data Plane Statistics*
- *Clearing Chassis Cluster Data Plane Statistics*
- *Understanding Chassis Cluster Formation*

show chassis cluster interfaces

Syntax **show chassis cluster interfaces**

Release Information Command modified in Junos OS Release 9.0. Output changed to support dual control ports in Junos OS Release 10.0. Output changed to support control interfaces in Junos OS Release 11.2. Output changed to support redundant pseudointerfaces in Junos OS Release 12.1X44-D10. For high-end SRX Series devices, output changed to support the internal security association (SA) option in Junos OS Release 12.1X45-D10.

Description Display the status of the control interface in a chassis cluster configuration.

Required Privilege Level view

Related Documentation
- *Chassis Cluster Feature Guide for Security Devices*

List of Sample Output
show chassis cluster interfaces on page 26
show chassis cluster interfaces (SRX3000 and SRX5000 line devices) on page 27
show chassis cluster interfaces on page 27

Output Fields Table 3 on page 25 lists the output fields for the **show chassis cluster interfaces** command. Output fields are listed in the approximate order in which they appear.

Table 3: show chassis cluster interfaces Output Fields

Field Name	Field Description
Control link status	State of the chassis cluster control interface: **up** or **down**.
Control interfaces	• **Index**—Index number of the chassis cluster control interface. • **Name**—Name of the chassis cluster control interface. • **Monitored-Status**—Monitored state of the interface: **up** or **down**. • **Internal SA**—State of the internal SA option on the chassis cluster control link: **enabled** or **disabled**. NOTE: This field is available only on high-end SRX Series devices.
Fabric link status	State of the fabric interface: **up** or **down**.
Fabric interfaces	• **Name**—Name of the fabric interface. • **Child-interface**—Name of the child fabric interface. • **Status**—State of the interface: **up** or **down**.
Redundant-ethernet Information	• **Name**—Name of the redundant Ethernet interface. • **Status**—State of the interface: **up** or **down**. • **Redundancy-group**—Identification number (1–255) of the redundancy group associated with the redundant Ethernet interface.

Table 3: show chassis cluster interfaces Output Fields *(continued)*

Field Name	Field Description
Redundant-pseudo-interface Information	• **Name**—Name of the redundant pseudo interface. • **Status**—State of the redundant pseudo interface: **up** or **down**. • **Redundancy-group**—Identification number (1–255) of the redundancy group associated with the redundant pseudo interface.
Interface Monitoring	• **Interface**—Name of the interface to be monitored. • **Weight**—Relative importance of the interface to redundancy group operation. • **Status**—State of the interface: **up** or **down**. • **Redundancy-group**—Identification number of the redundancy group associated with the interface.

Sample Output

show chassis cluster interfaces

```
user@host> show chassis cluster interfaces
Control link status: Up

Control interfaces:
      Index   Interface        Monitored-Status
      0       em0              Up
      1       em1              Down

Fabric link status: Up

Fabric interfaces:
    Name    Child-interface    Status
    fab0    ge-0/1/0           Up
    fab0
    fab1    ge-6/1/0           Up
    fab1

Redundant-ethernet Information:
    Name         Status      Redundancy-group
    reth0        Up          1
    reth1        Up          2
    reth2        Down        Not configured
    reth3        Down        Not configured
    reth4        Down        Not configured
    reth5        Down        Not configured
    reth6        Down        Not configured
    reth7        Down        Not configured
    reth8        Down        Not configured
    reth9        Down        Not configured
    reth10       Down        Not configured
    reth11       Down        Not configured

Redundant-pseudo-interface Information:
    Name        Status       Redundancy-group
     lo0          Up               1

Interface Monitoring:
    Interface         Weight    Status    Redundancy-group
```

```
    ge-0/1/9          100       Up        0
    ge-0/1/9          100       Up
```

Sample Output

show chassis cluster interfaces (SRX3000 and SRX5000 line devices)

```
user@host> show chassis cluster interfaces
Control link status: Up

Control interfaces:
      Index   Interface    Monitored-Status   Internal SA
      0       em0          Up                 enabled
      1       em1          Down               enabled

Fabric link status: Up

Fabric interfaces:
    Name    Child-interface    Status
    fab0    ge-0/1/0           Up
    fab0
    fab1    ge-6/1/0           Up
    fab1

Redundant-ethernet Information:
    Name         Status      Redundancy-group
    reth0        Up          1
    reth1        Up          2
    reth2        Down        Not configured
    reth3        Down        Not configured
    reth4        Down        Not configured
    reth5        Down        Not configured
    reth6        Down        Not configured
    reth7        Down        Not configured
    reth8        Down        Not configured
    reth9        Down        Not configured
    reth10       Down        Not configured
    reth11       Down        Not configured

Redundant-pseudo-interface Information:
    Name        Status       Redundancy-group
     lo0          Up              1

Interface Monitoring:
    Interface         Weight    Status    Redundancy-group
    ge-0/1/9          100       Up        0
    ge-0/1/9          100       Up
```

Sample Output

show chassis cluster interfaces

```
user@host> show chassis cluster interfaces
```

The below output is specific to fabric monitoring failure.

```
Control link status: Up

Control interfaces:
    Index   Interface    Monitored-Status   Internal-SA
    0       fxp1         Up                 Disabled
```

```
Fabric link status: Down

Fabric interfaces:
    Name    Child-interface    Status
                               (Physical/Monitored)
    fab0    ge-0/0/2           Down / Down
    fab0
    fab1    ge-9/0/2           Up   / Up
    fab1

Redundant-pseudo-interface Information:
    Name         Status      Redundancy-group
    lo0          Up          0
```

Enhanced Debugging Support for Chassis Cluster

- show chassis cluster information
- show chassis cluster ip-monitoring status redundancy-group
- show chassis cluster status

show chassis cluster information

Syntax show chassis cluster information

Release Information Command introduced in Junos OS Release 12.1X46-D20.

Description Display chassis cluster messages. The messages indicate each node's health condition and details of the monitored failure.

Required Privilege Level view

Related Documentation
- show chassis cluster status on page 37

List of Sample Output
show chassis cluster information on page 30
show chassis cluster information on page 31

Output Fields Table 4 on page 30 lists the output fields for the **show chassis cluster information** command. Output fields are listed in the approximate order in which they appear.

Table 4: show chassis cluster information Output Fields

Field Name	Field Description
Node	Node (device) in the chassis cluster (**node0** or **node1**).
Redundancy Group Information	• Redundancy Group—ID number (0 - 255) of a redundancy group in the cluster. • Current State—State of the redundancy group: **primary**, **secondary**, **hold**, or **secondary-hold**. • Weight—Relative importance of the redundancy group. • Time—Time when the redundancy group changed the state. • From—State of the redundancy group before the change. • To—State of the redundancy group after the change. • Reason—Reason for the change of state of the redundancy group.
Chassis cluster LED information	• Current LED color—Current color state of the LED. • Last LED change reason—Reason for change of state of the LED.

Sample Output

show chassis cluster information

```
user@host> show chassis cluster information

node0:
--------------------------------------------------------------------------
Redundancy Group Information:

    Redundancy Group 0 , Current State: primary, Weight: 255

        Time            From           To             Reason
        Mar 27 17:44:19 hold           secondary      Hold timer expired
```

```
        Mar 27 17:44:27 secondary      primary        Better priority (200/200)

    Redundancy Group 1 , Current State: primary, Weight: 255

        Time            From           To             Reason
        Mar 27 17:44:19 hold           secondary      Hold timer expired
        Mar 27 17:44:27 secondary      primary        Remote yield (0/0)

    Redundancy Group 2 , Current State: secondary, Weight: 255

        Time            From           To             Reason
        Mar 27 17:44:19 hold           secondary      Hold timer expired
        Mar 27 17:44:27 secondary      primary        Remote yield (0/0)
        Mar 27 17:50:24 primary        secondary-hold Preempt/yield(100/200)
        Mar 27 17:50:25 secondary-hold secondary      Ready to become secondary

Chassis cluster LED information:
    Current LED color: Green
    Last LED change reason: No failures

node1:
--------------------------------------------------------------------------
Redundancy Group Information:

    Redundancy Group 0 , Current State: secondary, Weight: 255

        Time            From           To             Reason
        Mar 27 17:44:27 hold           secondary      Hold timer expired

    Redundancy Group 1 , Current State: secondary, Weight: 255

        Time            From           To             Reason
        Mar 27 17:44:27 hold           secondary      Hold timer expired
        Mar 27 17:50:23 secondary      primary        Remote yield (100/0)
        Mar 27 17:50:24 primary        secondary-hold Preempt/yield(100/200)
        Mar 27 17:50:25 secondary-hold secondary      Ready to become secondary

    Redundancy Group 2 , Current State: primary, Weight: 255

        Time            From           To             Reason
        Mar 27 17:44:27 hold           secondary      Hold timer expired
        Mar 27 17:50:23 secondary      primary        Remote yield (200/0)

Chassis cluster LED information:
    Current LED color: Green
    Last LED change reason: No failures
```

Sample Output

show chassis cluster information

`user@host>` **show chassis cluster information**

The following output is specific to monitoring abnormal (unhealthy) case.

```
node0:
--------------------------------------------------------------------------
Redundancy Group Information:

    Redundancy Group 0 , Current State: secondary, Weight: 255
```

```
        Time            From           To             Reason
        Apr  1 11:07:38 hold           secondary      Hold timer expired
        Apr  1 11:07:41 secondary      primary        Only node present
        Apr  1 11:29:20 primary        secondary-hold Manual failover
        Apr  1 11:34:20 secondary-hold secondary      Ready to become secondary

    Redundancy Group 1 , Current State: primary, Weight: 0

        Time            From           To             Reason
        Apr  1 11:07:38 hold           secondary      Hold timer expired
        Apr  1 11:07:41 secondary      primary        Only node present

    Redundancy Group 2 , Current State: primary, Weight: 255

        Time            From           To             Reason
        Apr  1 11:07:38 hold           secondary      Hold timer expired
        Apr  1 11:07:41 secondary      primary        Only node present

Chassis cluster LED information:
    Current LED color: Amber
    Last LED change reason: Monitored objects are down

Failure Information:

    IP Monitoring Failure Information:
    Redundancy Group 1, Monitoring Status: Failed
    IP Address        Status       Reason
    1.1.1.1           Unreachable  redundancy-group state unknown

node1:
--------------------------------------------------------------------------
Redundancy Group Information:

    Redundancy Group 0 , Current State: primary, Weight: 255

        Time            From           To             Reason
        Apr  1 11:08:40 hold           secondary      Hold timer expired
       Apr  1 11:29:20 secondary      primary       Remote is in secondary hold


    Redundancy Group 1 , Current State: secondary, Weight: 0

        Time            From           To             Reason
        Apr  1 11:08:40 hold           secondary      Hold timer expired

    Redundancy Group 2 , Current State: secondary, Weight: 255

        Time            From           To             Reason
        Apr  1 11:08:40 hold           secondary      Hold timer expired

Chassis cluster LED information:
    Current LED color: Amber
    Last LED change reason: Monitored objects are down

Failure Information:

    IP Monitoring Failure Information:
    Redundancy Group 1, Monitoring Status: Failed
    IP Address        Status       Reason
    1.1.1.1           Unreachable  redundancy-group state unknown
```

show chassis cluster ip-monitoring status redundancy-group

Syntax **show chassis cluster ip-monitoring status**
<redundancy-group *group-number*>

Release Information Command introduced in Junos OS Release 9.6. Support for global threshold, current threshold, and weight of each monitored IP address added in Junos OS Release 12.1X46-D20.

Description Display the status of all monitored IP addresses for a redundancy group.

Options **none**—Display the status of monitored IP addresses for all redundancy groups on the node.

redundancy-group *group-number* —Display the status of monitored IP addresses under the specified redundancy group.

Required Privilege Level view

Related Documentation

- *IP Monitoring Feature Guide for Security Devices*
- *Chassis Cluster Feature Guide for Security Devices*
- *redundancy-group (Interfaces)*
- *clear chassis cluster failover-count*
- *request chassis cluster failover node*
- *request chassis cluster failover reset*

List of Sample Output
show chassis cluster ip-monitoring status on page 35
show chassis cluster ip-monitoring status redundancy-group on page 36

Output Fields Table 5 on page 34 lists the output fields for the **show chassis cluster ip-monitoring status** command.

Table 5: show chassis cluster ip-monitoring status Output Fields

Field Name	Field Description
Redundancy-group	ID number (0 - 255) of a redundancy group in the cluster.
Global threshold	Failover value for all IP addresses monitored by the redundancy group.
Current threshold	Value equal to the global threshold minus the total weight of the unreachable IP address.
IP Address	Monitored IP address in the redundancy group.

Table 5: show chassis cluster ip-monitoring status Output Fields *(continued)*

Field Name	Field Description
Status	Current reachability state of the monitored IP address. Values for this field are: **reachable**, **unreachable**, and **unknown**. The status is "unknown" if packet forwarding engines (PFEs) are not yet up and running.
Failure count	Number of attempts to reach an IP address.
Reason	Explanation for the reported status. See Table 6 on page 35.
Weight	Combined weight (0 - 255) assigned to all monitored IP addresses. A higher weight value indicates greater importance.

Expanded reason output fields for unreachable IP addresses added in Junos OS Release 10.1. You might see any of the following reasons displayed.

Table 6: show chassis cluster ip-monitoring status redundancy group Reason Fields

Reason	Reason Description
No route to host	The router could not resolve the ARP, which is needed to send the ICMP packet to the host with the monitored IP address.
No auxiliary IP found	The redundant Ethernet interface does not have an auxiliary IP address configured.
Reth child not up	A child interface of a redundant Ethernet interface is down.
redundancy-group state unknown	Unable to obtain the state (primary, secondary, secondary-hold, disable) of a redundancy-group.
No reth child MAC address	Could not extract the MAC address of the redundant Ethernet child interface.
Secondary link not monitored	The secondary link may be down (the secondary child interface of a redundant Ethernet interface is either down or non-functional).
Unknown	The IP address has just been configured and the router still does not know the status of this IP. or Do not know the exact reason for the failure.

Sample Output

show chassis cluster ip-monitoring status

```
user@host> show chassis cluster ip-monitoring status
node0:
--------------------------------------------------------------------------

Redundancy group: 1
```

```
Global threshold: 200
Current threshold: -120

IP address               Status        Failure count  Reason    Weight
10.254.5.44              reachable       0            n/a       220
2.2.2.1                  reachable       0            n/a       100

node1:
--------------------------------------------------------------------------

Redundancy group: 1
Global threshold: 200
Current threshold: -120

IP address               Status        Failure count  Reason    Weight
10.254.5.44              reachable       0            n/a       220
2.2.2.1                  reachable       0            n/a       100
```

Sample Output

show chassis cluster ip-monitoring status redundancy-group

```
user@host> show chassis cluster ip-monitoring status redundancy-group 1
node0:
--------------------------------------------------------------------------

Redundancy group: 1

IP address               Status        Failure count  Reason
10.254.5.44              reachable       0            n/a
2.2.2.1                  reachable       0            n/a
1.1.1.5                  reachable       0            n/a
1.1.1.4                  reachable       0            n/a
1.1.1.1                  reachable       0            n/a

node1:
--------------------------------------------------------------------------

Redundancy group: 1

IP address               Status        Failure count  Reason
10.254.5.44              reachable       0            n/a
2.2.2.1                  reachable       0            n/a
1.1.1.5                  reachable       0            n/a
1.1.1.4                  reachable       0            n/a
1.1.1.1                  reachable       0            n/a
```

show chassis cluster status

Syntax

show chassis cluster status
<redundancy-group ***group-number*** **>**

Release Information

Command modified in Junos OS Release 9.2. Support for dual control ports added in Junos OS Release 10.0. Support for monitoring failures added in Junos OS Release 12.1X46-D20.

Description

Display the failover status of a chassis cluster.

Options

none—Display the status of all redundancy groups in the chassis cluster.

redundancy-group ***group-number***—(Optional) Display the status of the specified redundancy group.

Required Privilege Level

view

Related Documentation

- *Chassis Cluster Feature Guide for Security Devices*
- *redundancy-group (Chassis Cluster)*
- *clear chassis cluster failover-count*
- *request chassis cluster failover node*
- *request chassis cluster failover reset*
- *Master Administrator for Logical Systems Feature Guide for Security Devices*

List of Sample Output

show chassis cluster status on page 38
show chassis cluster status redundancy-group 2 on page 38
show chassis cluster status on page 38
show chassis cluster status on page 39
show chassis cluster status on page 39

Output Fields

Table 7 on page 37 lists the output fields for the **show chassis cluster status** command. Output fields are listed in the approximate order in which they appear.

Table 7: show chassis cluster status Output Fields

Field Name	Field Description
Cluster ID	ID number (1-15) of a cluster is applicable for releases upto 12.1X45-D10. ID number (1-255) is applicable for releases 12.1X45-D10 and later. Setting a cluster ID to 0 is equivalent to disabling a cluster.
Redundancy-Group	ID number (1-128) of a redundancy group in the chassis cluster.
Node name	Node (device) in the chassis cluster (**node0** or **node1**).
Priority	Assigned priority for the redundancy group on that node.

Table 7: show chassis cluster status Output Fields *(continued)*

Field Name	Field Description
Status	State of the redundancy group (**Primary**, **Secondary**, **Lost**, or **Unavailable**). • **Primary**—Redundancy group is active and passing traffic. • **Secondary**—Redundancy group is passive and not passing traffic. • **Lost**—Node loses contact with the other node through the control link. Most likely to occur when both nodes are in a cluster and due to control link failure, one node cannot exchange heartbeats, or when the other node is rebooted. • **Unavailable**—Node has not received a single heartbeat over the control link from the other node since the other node booted up. Most likely to occur when one node boots up before the other node, or if only one node is present in the cluster.
Preempt	• **Yes**: Mastership can be preempted based on priority. • **No**: Change in priority will not preempt the mastership.
Manual failover	• **Yes**: If the Mastership is set manually through the CLI with the **request chassis cluster failover node** or **request chassis cluster failover redundancy-group** command. This overrides **Priority** and **Preempt**. • **No**: Mastership is not set manually through the CLI.

Sample Output

show chassis cluster status

```
user@host> show chassis cluster status
Cluster ID: 3
   Node name               Priority    Status    Preempt  Manual failover
 Redundancy-group: 0, Failover count: 1
    node0                   254         primary   no       no
    node1                   2           secondary no       no
 Redundancy-group: 1, Failover count: 1
    node0                   254         primary   no       no
    node1                   1           secondary no       no
```

Sample Output

show chassis cluster status redundancy-group 2

```
user@host> show chassis cluster status redundancy-group 2
Cluster ID: 14
   Node name          Priority    Status    Preempt  Manual failover
 Redundancy-Group: 2, Failover count: 1
    node0              50          secondary no       no
    node1              100         primary   no       no
```

Sample Output

show chassis cluster status

```
user@host> show chassis cluster status
Cluster ID: 15
Node                  Priority          Status    Preempt  Manual failover
Redundancy group: 0 , Failover count: 5
    node0                   200         primary        no       no
```

```
    node1                   0           lost          n/a      n/a
Redundancy group: 1 , Failover count: 41
    node0                   101         primary       no       no
    node1                   0           lost          n/a      n/a
```

Sample Output

show chassis cluster status

```
user@host> show chassis cluster status
Cluster ID: 15
Node                  Priority          Status    Preempt  Manual failover
Redundancy group: 0 , Failover count: 5
    node0                   200         primary       no       no
    node1                   0           unavailable   n/a      n/a
Redundancy group: 1 , Failover count: 41
    node0                   101         primary       no       no
    node1                   0           unavailable   n/a      n/a
```

Sample Output

show chassis cluster status

```
user@host> show chassis cluster status
```

The following output is specific to monitoring failures.

```
Monitor Failure codes:
    CS  Cold Sync monitoring        FL  Fabric connection down
    GR  GRES monitoring             HW  Hardware monitoring
    IF  Interface monitoring        IP  IP monitoring
    LB  Loopback monitoring         MB  Mbuf monitoring
    NH  Nexthop monitoring          NP  NPC monitoring
    SP  SPU monitoring              SM  Schedule monitoring

Cluster ID: 3
Node   Priority Status         Preempt Manual   Monitor-failures

Redundancy group: 0 , Failover count: 2
node0 0        secondary      no      no       FL
node1 200      primary        no      no       None

Redundancy group: 1 , Failover count: 6
node0 0        ineligible     no      no       FL
node1 200      primary        no      no       None
```

Public Key Infrastructure (PKI)

This topic includes the following sections:

Online Certificate Status Protocol (OCSP)

Understanding Online Certificate Status Protocol

OCSP is used to check the revocation status of X509 certificates. OCSP provides revocation status on certificates in real time and is useful in time-sensitive situations such as bank transactions and stock trades.

The revocation status of a certificate is checked by sending a request to an OCSP server that resides outside of an SRX Series device. Based on the response from the server, the VPN connection is allowed or denied. OCSP responses are not cached on SRX Series devices.

The OCSP server can be the certificate authority (CA) that issues a certificate or a designated authorized responder. The location of the OCSP server can be configured manually or extracted from the certificate that is being verified. Requests are sent first to OCSP server locations that are manually configured in CA profiles with the **ocsp url** statement at the [**edit security pki ca-profile** *profile-name* **revocation-check**] hierarchy level; up to two locations can be configured for each CA profile. If the first configured OCSP server is not reachable, the request is sent to the second OCSP server. If the second OCSP server is not reachable, the request is then sent to the location in the certificate's AuthorityInfoAccess extension field. The **use-ocsp** option must also be configured, as certificate revocation list (CRL) is the default checking method.

SRX Series devices accept only signed OCSP responses from the CA or authorized responder. The response received is validated using trusted certificates. The response is validated as follows:

1. The CA certificate enrolled for the configured CA profile is used to validate the response.
2. The OCSP response might contain a certificate to validate the OCSP response. The received certificate must be signed by a CA certificate enrolled in the SRX Series device. After the received certificate is validated by the CA certificate, it is used to validate the OCSP response.

The response from the OCSP server can be signed by different CAs. The following scenarios are supported:

- The CA server that issues the end entity certificate for a device also signs the OCSP revocation status response. The SRX Series device verifies the OCSP response signature using the CA certificate enrolled in the SRX Series device. After the OCSP response is validated, the certificate revocation status is checked.
- An authorized responder signs the OCSP revocation status response. The certificate for the authorized responder and the end entity certificate being verified must be issued by the same CA. The authorized responder is first verified using the CA certificate enrolled in the SRX Series device. The OCSP response is validated using the responder's CA certificate. The SRX Series device then uses the OCSP response to check the revocation status of the end entity certificate.
- There are different CA signers for the end entity certificate being verified and the OCSP response. The OCSP response is signed by a CA in the certificate chain for the end entity certificate being verified. (All peers participating in an IKE negotiation need to have at least one common trusted CA in their respective certificate chains.) The OCSP responder's CA is verified using a CA in the certificate chain. After validating the responder CA certificate, the OCSP response is validated using the responder's CA certificate.

To prevent replay attacks, a nonce payload can be sent in an OCSP request. Nonce payloads are sent by default unless it is explicitly disabled. If enabled, the SRX Series device expects the OCSP response to contain a nonce payload, otherwise the revocation check fails. If OCSP responders are not capable of responding with a nonce payload, then the nonce payload must be disabled on the SRX Series device.

Related Documentation

- Comparison of Online Certificate Status Protocol and Certificate Revocation List on page 41
- Example: Configuring OCSP for Certificate Revocation Status on page 47
- *Public Key Infrastructure Feature Guide for Security Devices*

Comparison of Online Certificate Status Protocol and Certificate Revocation List

Online Certificate Status Protocol (OCSP) and certificate revocation list (CRL) can both be used to check the revocation status of a certificate. There are advantages and disadvantages to each method.

- OCSP provides certificate status in real time, while CRL uses cached data. For time-sensitive applications, OCSP is the preferred approach.
- CRL checking is faster because lookup for certificate status is done on information cached on the VPN device. OCSP requires time to obtain the revocation status from an external server.
- CRL requires additional memory to store the revocation list received from a CRL server. OCSP does not require additional memory to save the revocation status of certificates.
- OCSP requires that the OCSP server be available at all times. CRL can use cached data to check the revocation status of certificates when the server is unreachable.

NOTE: On SRX Series devices, CRL is the default method used to check the revocation status of a certificate.

Related Documentation

- Understanding Online Certificate Status Protocol on page 40
- *Understanding Certificate Revocation Lists*
- *Public Key Infrastructure Feature Guide for Security Devices*

ocsp (Security PKI)

Syntax

```
ocsp {
  connection-failure (disable | fallback-crl);
  disable-responder-revocation-check;
  nonce-payload (enable | disable);
  url ocsp-url;
}
```

Hierarchy Level [edit security pki ca-profile *ca-profile-name* revocation-check]

Release Information Statement introduced in Junos OS Release 12.1X46-D20.

Description Configure Online Certificate Status Protocol (OCSP) to check the revocation status of a certificate.

Options

connection-failure—(Optional) Specify action to take if there is a connection failure to the OCSP responder. If this option is not configured and there is no response from the OCSP responder, certificate validation will fail.

disable—Skip the revocation check if the OCSP responder is not reachable.

fallback-crl—Use CRL to check the revocation status of the certificate.

disable-responder-revocation-check —(Optional) Disable revocation check for the CA certificate received in an OCSP response. The certificates received in an OCSP response generally have shorter lifetimes and revocation check is not required.

nonce-payload—(Optional) Send a nonce payload to prevent replay attack. A nonce payload is sent by default unless it is explicitly disabled. If enabled, the SRX Series device expects OCSP responses to contain a nonce payload, otherwise the revocation check will fail. If OCSP responders are not capable of responding with a nonce payload, disable this option.

disable—Explicitly disable the sending of a nonce payload.

enable—Enable the sending of a nonce payload. This is the default.

url *ocsp-url*—Specify HTTP addresses for OCSP responders. A maximum of two HTTP URL addresses can be configured. If the configured URLs are not reachable, or URLs are not configured, the URL from the certificate being verified is checked.

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

Related Documentation

- *Public Key Infrastructure Feature Guide for Security Devices*

use-ocsp (Security PKI)

Syntax use-ocsp;

Hierarchy Level [edit security pki ca-profile *ca-profile-name*]

Release Information Statement introduced in Junos OS Release 12.1X46-D20.

Description Specify the Online Certificate Status Protocol (OCSP) as the method to check the revocation status of a certificate. CRL is the default method.

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

Related Documentation
- *Public Key Infrastructure Feature Guide for Security Devices*

ca-profile (Security PKI)

Syntax

```
ca-profile ca-profile-name {
  administrator {
    e-mail-address e-mail-address;
  }
  ca-identity ca-identity ;
  enrollment {
      retry number;
      retry-interval seconds;
      url url-name;
  }
  revocation-check {
    crl {
      disable {
        on-download-failure;
      }
      refresh-interval hours;
      url url-name;
    }
    disable;
    ocsp {
      connection-failure (disable | fallback-crl);
      disable-responder-revocation-check;
      nonce-payload (enable | disable);
      url ocsp-url;
    }
    use-ocsp;
  }
  routing-instance routing-instance-name ;
}
```

Hierarchy Level [edit security pki]

Release Information Statement modified in Junos OS Release 8.5. Support for **ocsp** and **use-ocsp** options added in Junos OS Release 12.1X46-D20.

Description Configure certificate authority (CA) profile.

Options ***ca-profile-name*** —Name of a trusted CA.

The remaining statements are explained separately.

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

Related Documentation

- *Public Key Infrastructure Feature Guide for Security Devices*

revocation-check (Security PKI)

Syntax

```
revocation-check {
  crl {
    disable {
      on-download-failure;
    }
    refresh-interval hours;
    url url-name;
  }
  disable;
  ocsp {
    connection-failure (disable | fallback-crl);
    disable-responder-revocation-check;
    nonce-payload (enable | disable);
    url ocsp-url;
  }
  use-ocsp;
}
```

Hierarchy Level [edit security pki ca-profile *ca-profile-name*]

Release Information Statement modified in Junos OS Release 8.5. Support for **ocsp** and **use-ocsp** options added in Junos OS Release 12.1X46-D20.

Description Specify the method the device uses to verify the revocation status of digital certificates.

Options The remaining statements are explained separately.

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

Related Documentation

- *Public Key Infrastructure Feature Guide for Security Devices*

Example: Configuring OCSP for Certificate Revocation Status

This example shows how to configure two peers using the Online Certificate Status Protocol (OCSP) to check the revocation status of the certificates used in Phase 1 negotiations for the IPsec VPN tunnel.

Requirements

On each device:

- Obtain and enroll a local certificate. This can be done either manually or by using the Simple Certificate Enrollment Protocol (SCEP).
- Optionally, enable automatic renewal of the local certificate.
- Configure security policies to permit traffic to and from the peer device.

Overview

On both peers, a certificate authority (CA) profile OCSP-ROOT is configured with the following options:

- CA name is OCSP-ROOT.
- Enrollment URL is http://1.1.1.1:8080/scep/OCSP-ROOT/. This is the URL where SCEP requests to the CA are sent.
- The URL for the OCSP server is http://10.157.88.56:8210/OCSP-ROOT/.
- OCSP is used first to check the certificate revocation status. If there is no response from the OCSP server, then the certificate revocation list (CRL) is used to check the status. The CRL URL is http://1.1.1.1:8080/crl-as-der/currentcrl-45.crlid=45.
- The CA certificate received in an OCSP response is not checked for certificate revocation. Certificates received in an OCSP response generally have shorter lifetimes and a revocation check is not required.

Table 8 on page 47 shows the Phase 1 options used in this example.

Table 8: Phase 1 Options for OCSP Configuration Example

Option	Peer A	Peer B
IKE proposal	ike_prop	ike_prop
Authentication method	RSA signatures	RSA signatures
DH group	group2	group2
Authentication algorithm	SHA 1	SHA 1

Table 8: Phase 1 Options for OCSP Configuration Example *(continued)*

Option	Peer A	Peer B
Encryption algorithm	3DES CBC	3DES CBC
IKE policy	ike_policy	ike_policy
Mode	aggressive	aggressive
Proposal	ike_prop	ike_prop
Certificate	local-certificate localcert1	local-certificate localcert1
IKE gateway	jsr_gateway	jsr_gateway
Policy	ike_policy	ike_policy
Gateway address	101.10.2.50	100.10.1.50
Remote identity	localcert11.juniper.net	-
Local identity	-	localcert11.juniper.net
External interface	reth1	ge-0/0/2.0
Version	v2	v2

Table 9 on page 48 shows the Phase 2 options used in this example.

Table 9: Phase 2 Options for OCSP Configuration Example

Option	Peer A	Peer B
IPsec proposal	ipsec_prop	ipsec_prop
Protocol	ESP	ESP
Authentication algorithm	HMAC SHA1-96	HMAC SHA1-96
Encryption algorithm	3DES CBC	3DES CBC
Lifetime seconds	1200	1200
Lifetime kilobytes	150,000	150,000
IPsec policy	ipsec_policy	ipsec_policy
PFC keys	group2	group2
Proposal	ipsec_prop	ipsec_prop

Table 9: Phase 2 Options for OCSP Configuration Example *(continued)*

Option	Peer A	Peer B
VPN	test_vpn	test_vpn
Bind interface	st0.1	st0.1
IKE gateway	jsr_gateway	jsr_gateway
Policy	ipsec_policy	ipsec_policy
Establish tunnels	-	immediately

Figure 1 on page 49 shows the peer devices that are configured in this example.

Figure 1: OCSP Configuration Example

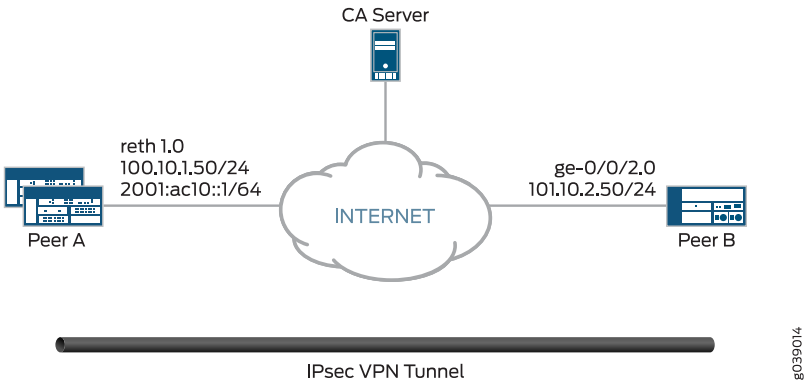

Configuration

- Configuring Peer A on page 49
- Configuring Peer B on page 53

Configuring Peer A

CLI Quick Configuration

To quickly configure VPN peer A to use OCSP, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the [**edit**] hierarchy level, and then enter **commit** from configuration mode.

```
set interfaces ge-0/0/3 gigether-options redundant-parent reth1
set interfaces ge-9/0/3 gigether-options redundant-parent reth1
set interfaces lo0 unit 0 family inet address 100.100.1.100/24
set interfaces lo0 redundant-pseudo-interface-options redundancy-group 1
set interfaces reth1 redundant-ether-options redundancy-group 1
set interfaces reth1 unit 0 family inet address 100.10.1.50/24
set interfaces st0 unit 1 family inet address 202.2.1.100/24
set security pki ca-profile OCSP-ROOT ca-identity OCSP-ROOT
```

```
set security pki ca-profile OCSP-ROOT enrollment url
  http://1.1.1.1:8080/scep/OCSP-ROOT/
set security pki ca-profile OCSP-ROOT revocation-check ocsp url
  http://10.157.88.56:8210/OCSP-ROOT/
set security pki ca-profile OCSP-ROOT revocation-check use-ocsp
set security pki ca-profile OCSP-ROOT revocation-check ocsp
  disable-responder-revocation-check
set security pki ca-profile OCSP-ROOT revocation-check ocsp connection-failure
  fallback-crl
set security pki ca-profile OCSP-ROOT revocation-check crl url
  http://1.1.1.1:8080/crl-as-der/currentcrl-45.crlid=45
set security ike proposal ike_prop authentication-method rsa-signatures
set security ike proposal ike_prop dh-group group2
set security ike proposal ike_prop authentication-algorithm sha1
set security ike proposal ike_prop encryption-algorithm 3des-cbc
set security ike policy ike_policy mode aggressive
set security ike policy ike_policy proposals ike_prop
set security ike policy ike_policy certificate local-certificate localcert1
set security ike gateway jsr_gateway ike-policy ike_policy
set security ike gateway jsr_gateway address 101.10.2.50
set security ike gateway jsr_gateway remote-identity hostname localcert11.juniper.net
set security ike gateway jsr_gateway external-interface reth1
set security ike gateway jsr_gateway version v2-only
set security ipsec proposal ipsec_prop protocol esp
set security ipsec proposal ipsec_prop authentication-algorithm hmac-sha1-96
set security ipsec proposal ipsec_prop encryption-algorithm 3des-cbc
set security ipsec proposal ipsec_prop lifetime-seconds 1200
set security ipsec proposal ipsec_prop lifetime-kilobytes 150000
set security ipsec policy ipsec_policy perfect-forward-secrecy keys group2
set security ipsec policy ipsec_policy proposals ipsec_prop
set security ipsec vpn test_vpn bind-interface st0.1
set security ipsec vpn test_vpn ike gateway jsr_gateway
set security ipsec vpn test_vpn ike ipsec-policy ipsec_policy
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure VPN peer A to use OCSP:

1. Configure interfaces.

```
[edit interfaces]
set ge-0/0/3 gigether-options redundant-parent reth1
set ge-9/0/3 gigether-options redundant-parent reth1
set lo0 unit 0 family inet address 100.100.1.100/24
set lo0 redundant-pseudo-interface-options redundancy-group 1
set reth1 redundant-ether-options redundancy-group 1
set reth1 unit 0 family inet address 100.10.1.50/24
set st0 unit 1 family inet address 202.2.1.100/24
```

2. Configure the CA profile.

```
[edit security pki ca-profile OCSP-ROOT]
set ca-identity OCSP-ROOT
set enrollment url http://1.1.1.1:8080/scep/OCSP-ROOT/
```

```
set revocation-check ocsp url http://10.157.88.56:8210/OCSP-ROOT/
set revocation-check use-ocsp
set revocation-check ocsp disable-responder-revocation-check
set revocation-check ocsp connection-failure fallback-crl
set revocation-check crl url http://1.1.1.1:8080/crl-as-der/currentcrl-45.crlid=45
```

3. Configure Phase 1 options.

```
[edit security ike proposal ike_prop]
set authentication-method rsa-signatures
set dh-group group2
set authentication-algorithm sha1
set encryption-algorithm 3des-cbc

[edit security ike policy ike_policy]
set mode aggressive
set proposals ike_prop
set certificate local-certificate localcert1

[edit security ike gateway jsr_gateway]
set ike-policy ike_policy
set address 101.10.2.50
set remote-identity hostname localcert11.juniper.net
set external-interface reth1
set version v2-only
```

4. Configure Phase 2 options.

```
[edit security ipsec proposal ipsec_prop]
set protocol esp
set authentication-algorithm hmac-sha1-96
set encryption-algorithm 3des-cbc
set lifetime-seconds 1200
set lifetime-kilobytes 150000

[edit security ipsec policy ipsec_policy]
set perfect-forward-secrecy keys group2
set proposals ipsec_prop

[edit security ipsec vpn test_vpn]
set bind-interface st0.1
set ike gateway jsr_gateway
set ike ipsec-policy ipsec_policy
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show security pki ca-profile OCSP-ROOT**, **show security ike**, and **show security ipsec** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show interfaces
ge-0/0/3 {
  gigether-options {
    redundant-parent reth1;
  }
```



```
}
ge-9/0/3 {
  gigether-options {
    redundant-parent reth1;
  }
}
lo0 {
  unit 0 {
    family inet {
      address 100.100.1.100/24;
    }
  }
  redundant-pseudo-interface-options {
    redundancy-group 1;
  }
}
reth1 {
  redundant-ether-options {
    redundancy-group 1;
  }
  unit 0 {
    family inet {
      address 100.10.1.50/24;
    }
  }
}
st0 {
  unit 1 {
    family inet {
      address 202.2.1.100/24;
    }
  }
}
[edit]
user@host# show security pki ca-profile OCSP-ROOT
ca-identity OCSP-ROOT;
enrollment {
  url http://1.1.1.1:8080/scep/OCSP-ROOT/;
}
revocation-check {
  crl {
    url http://1.1.1.1:8080/crl-as-der/currentcrl-45.crlid=45;
  }
  ocsp {
    disable-responder-revocation-check;
    url http://10.157.88.56:8210/OCSP-ROOT/;
  }
  use-ocsp;
}
[edit]
user@host# show security ike
proposal ike_prop {
  authentication-method rsa-signatures;
  dh-group group2;
  authentication-algorithm sha1;
  encryption-algorithm 3des-cbc;
```

```
}
policy ike_policy {
    mode aggressive;
    proposals ike_prop;
    certificate {
        local-certificate localcert1;
    }
}
gateway jsr_gateway {
    ike-policy ike_policy;
    address 101.10.2.50;
    remote-identity hostname localcert11.juniper.net;
    external-interface reth1;
    version v2-only;
}
[edit]
user@host# show security ipsec
proposal ipsec_prop {
    protocol esp;
    authentication-algorithm hmac-sha1-96;
    encryption-algorithm 3des-cbc;
    lifetime-seconds 1200;
    lifetime-kilobytes 150000;
}
policy ipsec_policy {
    perfect-forward-secrecy {
        keys group2;
    }
    proposals ipsec_prop;
}
vpn test_vpn {
    bind-interface st0.1;
    ike {
        gateway jsr_gateway;
        ipsec-policy ipsec_policy;
    }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Configuring Peer B

CLI Quick Configuration

To quickly configure VPN peer B to use OCSP, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the [**edit**] hierarchy level, and then enter **commit** from configuration mode.

```
set interfaces ge-0/0/2 unit 0 family inet address 101.10.2.50/24
set interfaces lo0 unit 0 family inet address 102.100.1.100/24
set interfaces st0 unit 1 family inet address 202.2.1.1/24
set security pki ca-profile OCSP-ROOT ca-identity OCSP-ROOT
set security pki ca-profile OCSP-ROOT enrollment url
    http://1.1.1.1:8080/scep/OCSP-ROOT/
set security pki ca-profile OCSP-ROOT revocation-check ocsp url
    http://10.157.88.56:8210/OCSP-ROOT/
```

```
set security pki ca-profile OCSP-ROOT revocation-check use-ocsp
set security pki ca-profile OCSP-ROOT revocation-check ocsp
  disable-responder-revocation-check
set security pki ca-profile OCSP-ROOT revocation-check ocsp connection-failure
  fallback-crl
set security pki ca-profile OCSP-ROOT revocation-check crl url
  http://1.1.1.1:8080/crl-as-der/currentcrl-45.crlid=45
set security ike proposal ike_prop authentication-method rsa-signatures
set security ike proposal ike_prop dh-group group2
set security ike proposal ike_prop authentication-algorithm sha1
set security ike proposal ike_prop encryption-algorithm 3des-cbc
set security ike policy ike_policy mode aggressive
set security ike policy ike_policy proposals ike_prop
set security ike policy ike_policy certificate local-certificate localcert11
set security ike gateway jsr_gateway ike-policy ike_policy
set security ike gateway jsr_gateway address 100.10.1.50
set security ike gateway jsr_gateway local-identity hostname localcert11.juniper.net
set security ike gateway jsr_gateway external-interface ge-0/0/2.0
set security ike gateway jsr_gateway version v2-only
set security ipsec proposal ipsec_prop protocol esp
set security ipsec proposal ipsec_prop authentication-algorithm hmac-sha1-96
set security ipsec proposal ipsec_prop encryption-algorithm 3des-cbc
set security ipsec proposal ipsec_prop lifetime-seconds 1200
set security ipsec proposal ipsec_prop lifetime-kilobytes 150000
set security ipsec policy ipsec_policy perfect-forward-secrecy keys group2
set security ipsec policy ipsec_policy proposals ipsec_prop
set security ipsec vpn test_vpn bind-interface st0.1
set security ipsec vpn test_vpn ike gateway jsr_gateway
set security ipsec vpn test_vpn ike ipsec-policy ipsec_policy
set security ipsec vpn test_vpn establish-tunnels immediately
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure VPN peer B to use OCSP:

1. Configure interfaces.

   ```
   [edit interfaces]
   set ge-0/0/2 unit 0 family inet address 101.10.2.50/24
   set lo0 unit 0 family inet address 102.100.1.100/24
   set st0 unit 1 family inet address 202.2.1.1/24
   ```

2. Configure the CA profile.

   ```
   [edit security pki ca-profile OCSP-ROOT]
   set ca-identity OCSP-ROOT
   set enrollment url http://1.1.1.1:8080/scep/OCSP-ROOT/
   set revocation-check ocsp url http://10.157.88.56:8210/OCSP-ROOT/
   set revocation-check use-ocsp
   set revocation-check ocsp disable-responder-revocation-check
   set revocation-check ocsp connection-failure fallback-crl
   set revocation-check crl url http://1.1.1.1:8080/crl-as-der/currentcrl-45.crlid=45
   ```

3. Configure Phase 1 options.

```
[edit security ike proposal ike_prop]
set authentication-method rsa-signatures
set dh-group group2
set authentication-algorithm sha1
set encryption-algorithm 3des-cbc

[edit security ike policy ike_policy]
set mode aggressive
set proposals ike_prop
set certificate local-certificate localcert1

[edit security ike gateway jsr_gateway]
set ike-policy ike_policy
set address 100.10.1.50
set local-identity hostname localcert11.juniper.net
set external-interface ge-0/0/2.0
set version v2-only
```

4. Configure Phase 2 options.

```
[edit security ipsec proposal ipsec_prop]
set protocol esp
set authentication-algorithm hmac-sha1-96
set encryption-algorithm 3des-cbc
set lifetime-seconds 1200
set lifetime-kilobytes 150000

[edit security ipsec policy ipsec_policy]
set perfect-forward-secrecy keys group2
set proposals ipsec_prop

[edit security ipsec vpn test_vpn]
set bind-interface st0.1
set ike gateway jsr_gateway
set ike ipsec-policy ipsec_policy
set establish-tunnels immediately
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show security pki ca-profile OCSP-ROOT**, **show security ike**, and **show security ipsec** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show interfaces
ge-0/0/2 {
   unit 0 {
      family inet {
         address 101.10.2.50/24;
      }
   }
}
lo0 {
   unit 0 {
      family inet {
```

```
        address 102.100.1.100/24;
      }
    }
  }
  st0 {
    unit 1 {
      family inet {
        address 202.2.1.1/24;
      }
    }
  }
  [edit]
  user@host# show security pki ca-profile OCSP-ROOT
  ca-identity OCSP-ROOT;
  enrollment {
    url http://1.1.1.1:8080/scep/OCSP-ROOT/;
  }
  revocation-check {
    crl {
      url http://1.1.1.1:8080/crl-as-der/currentcrl-45.crlid=45;
    }
    ocsp {
      disable-responder-revocation-check;
      url http://10.157.88.56:8210/OCSP-ROOT/;
    }
    use-ocsp;
  }
  [edit]
  user@host# show security ike
  proposal ike_prop {
    authentication-method rsa-signatures;
    dh-group group2;
    authentication-algorithm sha1;
    encryption-algorithm 3des-cbc;
  }
  policy ike_policy {
    mode aggressive;
    proposals ike_prop;
    certificate {
      local-certificate localcert11;
    }
  }
  gateway jsr_gateway {
    ike-policy ike_policy;
    address 100.10.1.50;
    local-identity hostname localcert11.juniper.net;
    external-interface ge-0/0/2.0;
    version v2-only;
  }
  [edit]
  user@host# show security ipsec
  proposal ipsec_prop {
    protocol esp;
    authentication-algorithm hmac-sha1-96;
    encryption-algorithm 3des-cbc;
    lifetime-seconds 1200;
```

```
    lifetime-kilobytes 150000;
}
policy ipsec_policy {
    perfect-forward-secrecy {
        keys group2;
    }
    proposals ipsec_prop;
}
vpn test_vpn {
    bind-interface st0.1;
    ike {
        gateway jsr_gateway;
        ipsec-policy ipsec_policy;
    }
    establish-tunnels immediately;
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly.

- Verifying CA Certificates on page 57
- Verifying Local Certificates on page 58
- Verifying IKE Phase 1 Status on page 59
- Verifying IPsec Phase 2 Status on page 60

Verifying CA Certificates

Purpose Verify the validity of a CA certificate on each peer device.

Action From operational mode, enter the **show security pki ca-certificate ca-profile OCSP-ROOT** or **show security pki ca-certificate ca-profile OCSP-ROOT detail** command.

```
user@host> show security pki ca-certificate ca-profile OCSP-ROOT
  Certificate identifier: OCSP-ROOT
    Issued to: OCSP-ROOT, Issued by: C = US, O = Juniper, CN = OCSP-ROOT
    Validity:
      Not before: 11-15-2013 22:26 UTC
      Not after: 11-14-2016 22:26 UTC
    Public key algorithm: rsaEncryption(2048 bits)

user@host> show security pki ca-certificate ca-profile OCSP-ROOT detail
  Certificate identifier: OCSP-ROOT
    Certificate version: 3
    Serial number: 0000a17f
    Issuer:
      Organization: Juniper, Country: US, Common name: OCSP-ROOT
    Subject:
      Organization: Juniper, Country: US, Common name: OCSP-ROOT
    Subject string:
      C=US, O=Juniper, CN=OCSP-ROOT
    Validity:
```

```
    Not before: 11-15-2013 22:26 UTC
    Not after: 11-14-2016 22:26 UTC
  Public key algorithm: rsaEncryption(2048 bits)
    30:82:01:0a:02:82:01:01:00:c6:38:e9:03:69:5e:45:d8:a3:ea:3d
    2e:e3:b8:3f:f0:5b:39:f0:b7:35:64:ed:60:a0:ba:89:28:63:29:e7
    27:82:47:c4:f6:41:53:c8:97:d7:1e:3c:ca:f0:a0:b9:09:0e:3d:f8
    76:5b:10:6f:b5:f8:ef:c5:e8:48:b9:fe:46:a3:c6:ba:b5:05:de:2d
    91:ce:20:12:8f:55:3c:a6:a4:99:bb:91:cf:05:5c:89:d3:a7:dc:a4
    d1:46:f2:dc:36:f3:f0:b5:fd:1d:18:f2:e6:33:d3:38:bb:44:8a:19
    ad:e0:b1:1a:15:c3:56:07:f9:2d:f6:19:f7:cd:80:cf:61:de:58:b8
    a3:f5:e0:d1:a3:3a:19:99:80:b0:63:03:1f:25:05:cc:b2:0c:cd:18
    ef:37:37:46:91:20:04:bc:a3:4a:44:a9:85:3b:50:33:76:45:d9:ba
    26:3a:3b:0d:ff:82:40:36:64:4e:ea:6a:d8:9b:06:ff:3f:e2:c4:a6
    76:ee:8b:58:56:a6:09:d3:4e:08:b0:64:60:75:f3:e2:06:91:64:73
    d2:78:e9:7a:cb:8c:57:0e:d1:9a:6d:3a:4a:9e:5b:d9:e4:a2:ef:31
    5d:2b:2b:53:ab:a1:ad:45:49:fd:a5:e0:8b:4e:0b:71:52:ca:6b:fa
    8b:0e:2c:7c:7b:02:03:01:00:01
  Signature algorithm: sha1WithRSAEncryption
  Distribution CRL:
    http://1.1.1.1:8080/crl-as-der/currentcrl-45.crl?id=45
  Authority Information Access OCSP:
    http://1.1.1.1:8090/OCSP-ROOT/
  Use for key: CRL signing, Certificate signing, Key encipherment, Digital signature
  Fingerprint:
    ed:ce:ec:13:1a:d2:ab:0a:76:e5:26:6d:2c:29:5d:49:90:57:f9:41 (sha1)
    af:87:07:69:f0:3e:f7:c6:b8:2c:f8:df:0b:ae:b0:28 (md5)
```

In this example, IP addresses are used in the URLs in the CA profile configuration. If IP addresses are not used with CA-issued certificates or CA certificates, DNS must be configured in the device's configuration. DNS must be able to resolve the host in the distribution CRL and in the CA URL in the CA profile configuration. Additionally, you must have network reachability to the same host to receive revocation checks.

Verifying Local Certificates

Purpose Verify the validity of a local certificate on each peer device.

Action From operational mode, enter the **show security pki local-certificate certificate-id localcert1 detail** command.

```
user@host> show security pki local-certificate certificate-id localcert1 detail
  Certificate identifier: localcert1
    Certificate version: 3
    Serial number: 013e3f1d
    Issuer:
      Organization: Juniper, Country: US, Common name: OCSP-ROOT
    Subject:
      Organization: juniper1, Organizational unit: sltqa1, State: california1, Locality: sunnyvale1,
    Common name: localcert1, Domain component: domain_component1
    Subject string:
      DC=domain_component1, CN=localcert1, OU=sltqa1, O=juniper1, L=sunnyvale1,
    ST=california1, C=us1
    Alternate subject: "localcert1@juniper.net", localcert1.juniper.net, 100.10.1.50
    Validity:
      Not before: 01-28-2014 22:23 UTC
```

```
  Not after: 03-29-2014 22:53 UTC
 Public key algorithm: rsaEncryption(1024 bits)
  30:81:89:02:81:81:00:a6:df:c1:57:59:f8:4d:0f:c4:a8:96:25:97
  03:c4:a0:fb:df:d5:f3:d5:56:b6:5a:26:65:b8:1a:ec:be:f6:c6:5f
  b3:d7:d3:59:39:48:52:4a:e3:1b:e4:e0:6d:24:c3:c1:50:8c:55:3b
  c0:c1:29:a0:45:29:8e:ec:3e:52:2f:84:b3:e8:89:9a:0f:8b:7d:e8
  90:4b:c1:28:48:95:b3:aa:11:ab:b4:8c:a8:80:ce:90:07:2a:13:a2
  2f:84:44:92:3b:be:7d:39:5b:2f:9a:4c:7a:2f:2d:31:8b:12:6d:52
  34:7d:6b:e4:69:7e:f3:86:55:e2:89:31:98:c9:15:02:03:01:00:01
 Signature algorithm: sha1WithRSAEncryption
 Distribution CRL:
  http://1.1.1.1:8080/crl-as-der/currentcrl-45.crl?id=45
 Authority Information Access OCSP:
  http://1.1.1.1/:8090/OCSP-ROOT/
 Fingerprint:
  00:c6:56:64:ad:e3:ce:8e:26:6b:df:17:1e:de:fc:14:a4:bb:8c:e4 (sha1)
  7f:43:c6:ed:e4:b3:7a:4f:9a:8c:0b:61:95:01:c9:52 (md5)
 Auto-re-enrollment:
  Status: Disabled
  Next trigger time: Timer not started
```

Verifying IKE Phase 1 Status

Purpose Verify the IKE Phase 1 status on each peer device.

Action From operational mode, enter the **show security ike security-associations** command.

```
user@host> show security ike security-associations
  Index   State  Initiator cookie  Responder cookie  Mode        Remote Address
  6534660 UP    3e62e05abd6a703f  c552b238e8a26668  IKEv2       101.10.2.50
```

From operational mode, enter the **show security ike security-associations detail** command.

```
user@host> show security ike security-associations detail
  IKE peer 101.10.2.50, Index 6534660, Gateway Name: jsr_gateway
   Role: Responder, State: UP
   Initiator cookie: 3e62e05abd6a703f, Responder cookie: c552b238e8a26668
   Exchange type: IKEv2, Authentication method: RSA-signatures
   Local: 100.10.1.50:500, Remote: 101.10.2.50:500
   Lifetime: Expires in 26906 seconds
   Peer ike-id: localcert11.juniper.net
   Xauth assigned IP: 0.0.0.0
   Algorithms:
    Authentication      : hmac-sha1-96
    Encryption          : 3des-cbc
    Pseudo random function: hmac-sha1
    Diffie-Hellman group  : DH-group-2
   Traffic statistics:
    Input  bytes :             2152
    Output bytes :             2097
    Input  packets:               4
    Output packets:               4
   Flags: IKE SA is created
   IPSec security associations: 4 created, 0 deleted
   Phase 2 negotiations in progress: 0
```

```
Negotiation type: Quick mode, Role: Responder, Message ID: 0
Local: 100.10.1.50:500, Remote: 101.10.2.50:500
Local identity: 100.10.1.50
Remote identity: localcert11.juniper.net
Flags: IKE SA is created
```

Verifying IPsec Phase 2 Status

Purpose Verify the IPsec Phase 2 status on each peer device.

Action From operational mode, enter the **show security ipsec security-associations** command.

```
user@host> show security ipsec security-associations
  Total active tunnels: 1
  ID   Algorithm     SPI     Life:sec/kb  Mon lsys Port  Gateway
  <131073 ESP:3des/sha1 9d1066e2 252/  150000 -   root 500   101.10.2.50
  >131073 ESP:3des/sha1 82079c2c 252/  150000 -   root 500   101.10.2.50
```

From operational mode, enter the **show security ipsec security-associations detail** command.

```
user@host> show security ipsec security-associations detail
  ID: 131073 Virtual-system: root, VPN Name: test_vpn
  Local Gateway: 100.10.1.50, Remote Gateway: 101.10.2.50
  Local Identity: ipv4_subnet(any:0,[0..7]=0.0.0.0/0)
  Remote Identity: ipv4_subnet(any:0,[0..7]=0.0.0.0/0)
  Version: IKEv2
   DF-bit: clear
   Bind-interface: st0.1

  Port: 500, Nego#: 2, Fail#: 0, Def-Del#: 0 Flag: 0x600a29
  Last Tunnel Down Reason: Delete payload received
   Direction: inbound, SPI: 9d1066e2, AUX-SPI: 0
                    , VPN Monitoring: -
   Hard lifetime: Expires in 249 seconds
   Lifesize Remaining:  150000 kilobytes
   Soft lifetime: Expires in 10 seconds
   Mode: Tunnel(0 0), Type: dynamic, State: installed
   Protocol: ESP, Authentication: hmac-sha1-96, Encryption: 3des-cbc
   Anti-replay service: counter-based enabled, Replay window size: 64

   Direction: outbound, SPI: 82079c2c, AUX-SPI: 0
                    , VPN Monitoring: -
   Hard lifetime: Expires in 249 seconds
   Lifesize Remaining:  150000 kilobytes
   Soft lifetime: Expires in 10 seconds
   Mode: Tunnel(0 0), Type: dynamic, State: installed
   Protocol: ESP, Authentication: hmac-sha1-96, Encryption: 3des-cbc
   Anti-replay service: counter-based enabled, Replay window size: 64
```

Related Documentation

- Understanding Online Certificate Status Protocol on page 40
- *Public Key Infrastructure Feature Guide for Security Devices*

Routing

This topic includes the following sections:

OSPF and OSPFv3 IPsec Authentication and Confidentiality

Understanding OSPF and OSPFv3 Authentication on SRX Series Devices

OSPFv3 does not have a built-in authentication method and relies on the IP Security (IPsec) suite to provide this functionality. IPsec provides authentication of origin, data integrity, confidentiality, replay protection, and nonrepudiation of source. You can use IPsec to secure specific OSPFv3 interfaces and virtual links and to provide encryption for OSPF packets.

OSPFv3 uses the IP authentication header (AH) and the IP Encapsulating Security Payload (ESP) portions of the IPsec protocol to authenticate routing information between peers. AH can provide connectionless integrity and data origin authentication. It also provides protection against replays. AH authenticates as much of the IP header as possible, as well as the upper-level protocol data. However, some IP header fields might change in transit. Because the value of these fields might not be predictable by the sender, they cannot be protected by AH. ESP can provide encryption and limited traffic flow confidentiality or connectionless integrity, data origin authentication, and an anti-replay service.

IPsec is based on security associations (SAs). An SA is a set of IPsec specifications that are negotiated between devices that are establishing an IPsec relationship. This simplex connection provides security services to the packets carried by the SA. These specifications include preferences for the type of authentication, encryption, and IPsec protocol to be used when establishing the IPsec connection. An SA is used to encrypt and authenticate a particular flow in one direction. Therefore, in normal bidirectional traffic, the flows are secured by a pair of SAs. An SA to be used with OSPFv3 must be configured manually and use transport mode. Static values must be configured on both ends of the SA.

To configure IPsec for OSPF or OSPFv3, first define a manual SA with the **security-association** ***sa-name*** option at the [**edit security ipsec**] hierarchy level. This feature

only supports bidirectional manual key SAs in transport mode. Manual SAs require no negotiation between the peers. All values, including the keys, are static and specified in the configuration. Manual SAs statically define the security parameter index (SPI) values, algorithms, and keys to be used and require matching configurations on both endpoints (OSPF or OSPFv3 peers). As a result, each peer must have the same configured options for communication to take place.

The actual choice of encryption and authentication algorithms is left to your IPsec administrator; however, we have the following recommendations:

- Use ESP with null encryption to provide authentication to protocol headers but not to the IPv6 header, extension headers, and options. With null encryption, you are choosing not to provide encryption on protocol headers. This can be useful for troubleshooting and debugging purposes. For more information about null encryption, see RFC 2410, *The NULL Encryption Algorithm and Its Use with IPsec*.
- Use ESP with DES or 3DES for full confidentiality.
- Use AH to provide authentication to protocol headers, immutable fields in IPv6 headers, and extension headers and options.

The configured SA is applied to the OSPF or OSPFv3 configurations as follows:

- For an OSPF or OSPFv3 interface, include the **ipsec-sa** *name* statement at the [**edit protocols ospf area** *area-id* **interface** *interface-name*] or [**edit protocols ospf3 area** *area-id* **interface** *interface-name*] hierarchy level. Only one IPsec SA name can be specified for an OSPF or OSPFv3 interface; however, different OSPF/OSPFv3 interfaces can specify the same IPsec SA.
- For an OSPF or OSPFv3 virtual link, include the **ipsec-sa** *name* statement at the [**edit protocols ospf area** *area-id* **virtual-link neighbor-id** *router-id* **transit-area** *area-id*] or [**edit protocols ospf3 area** *area-id* **virtual-link neighbor-id** *router-id* **transit-area** *area-id*] hierarchy level. You must configure the same IPsec SA for all virtual links with the same remote endpoint address.

The following restrictions apply to IPsec authentication for OSPF or OSPFv3 on SRX Series devices:

- Manual VPN configurations that are configured at the [**edit security ipsec vpn** *vpn-name* **manual**] hierarchy level cannot be applied to OSPF or OSPFv3 interfaces or virtual links to provide IPsec authentication and confidentiality.
- You cannot configure IPsec for OSPF or OSPFv3 authentication if there is an existing IPsec VPN configured on the device with the same local and remote addresses.
- IPsec for OSPF or OSPFv3 authentication is not supported over secure tunnel st0 interfaces.
- Rekeying of manual keys is not supported.
- Dynamic Internet Key Exchange (IKE) SAs are not supported.

- Only IPsec transport mode is supported. In transport mode, only the payload (the data you transfer) of the IP packet is encrypted, authenticated, or both. Tunnel mode is not supported.
- Because only bidirectional manual SAs are supported, all OSPFv3 peers must be configured with the same IPsec SA. You configure a manual bidirectional SA at the [**edit security ipsec**] hierarchy level.
- You must configure the same IPsec SA for all virtual links with the same remote endpoint address.

Related Documentation

- Example: Configuring IPsec Authentication for an OSPF Interface on an SRX Series Device on page 70

security-association

Syntax

```
security-association sa-name {
  manual {
    direction bidirectional {
      authentication {
        algorithm (hmac-md5-96 | hmac-sha1-96);
        key {
          ascii-text key;
          hexadecimal key;
        }
      }
      auxiliary-spi auxiliary-spi-value;
      encryption {
        algorithm (3des-cbc | des-cbc | null);
        key {
          ascii-text key;
          hexadecimal key;
        }
      }
      protocol (ah | esp);
      spi spi-value;
    }
  }
  mode transport;
}
```

Hierarchy Level [edit security ipsec]

Release Information Statement introduced in Junos OS Release 12.1X46-D20.

Description Configure a manual IPsec security association (SA) to be applied to an OSPF or OSPFv3 interface or virtual link. IPsec can provide authentication and confidentiality to OSPF or OSPFv3 routing packets.

Options *sa-name*—Name of the SA.

mode—SA mode. For this feature, the mode must be **transport**.

direction—Direction of the manual SA. For this feature, the direction must be **bidirectional**.

The remaining statements are explained separately.

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

Related Documentation

- Understanding OSPF and OSPFv3 Authentication on SRX Series Devices on page 61

authentication (IPsec SA for OSPF)

Syntax

```
authentication {
  algorithm (hmac-md5-96 | hmac-sha1-96);
  key {
    ascii-text key;
    hexadecimal key;
  }
}
```

Hierarchy Level [edit security ipsec security-association *sa-name* manual direction bidirectional]

Release Information Statement introduced in Junos OS Release 12.1X46-D20.

Description Configure authentication parameters for a manual IPsec security association (SA) to be applied to an OSPF or OSPFv3 interface or virtual link.

Options **algorithm**—Hash algorithm that authenticates packet data. It can be one of the following:

- **hmac-md5-96**—Produce a 128-bit digest. This is the default.
- **hmac-sha1-96**—Produce a 160-bit digest.

key—Type of authentication key. It can be one of the following:

- **ascii-text *key***—ASCII text key. For **hmac-md5-96**, the key is 16 ASCII characters; for **hmac-sha1-96**, the key is 20 ASCII characters.
- **hexadecimal *key***—Hexadecimal key. For **hmac-md5-96**, the key is 32 hexadecimal characters; for **hmac-sha1-96**, the key is 40 hexadecimal characters.

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

Related Documentation

- Understanding OSPF and OSPFv3 Authentication on SRX Series Devices on page 61

auxiliary-spi (IPsec SA for OSPF)

Syntax auxiliary-spi *auxiliary-spi-value*;

Hierarchy Level [edit security ipsec security-association *sa-name* mode transport manual direction bidirectional]

Release Information Statement introduced in Junos OS Release 12.1X46-D20.

Description Configure an auxiliary security parameter index (SPI) for a manual IPsec security association (SA) to be applied to an OSPF or OSPFv3 interface or virtual link.

Options **auxiliary-spi**—Auxiliary SPI for the manual IPsec SA. The SPI uniquely identifies the SA to use at the receiving host (the destination address in the packet).

Range: 256 through 16639

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

Related Documentation

- Understanding OSPF and OSPFv3 Authentication on SRX Series Devices on page 61

encryption (IPsec SA for OSPF)

Syntax

```
encryption {
  algorithm (3des-cbc | des-cbc | null);
  key {
    ascii-text key;
    hexadecimal key;
  }
}
```

Hierarchy Level [edit security ipsec security-association *sa-name* manual direction bidirectional]

Release Information Statement introduced in Junos OS Release 12.1X46-D20.

Description Configure encryption parameters for a manual IPsec security association (SA) to be applied to an OSPF or OSPFv3 interface or virtual link.

Options **algorithm**—Type of encryption algorithm. It can be one of the following:

- **3des-cbc**—Has block size of 8 bytes (64 bits); its key size is 192 bits long.
- **des-cbc**—Has a block size of 8 bytes (64 bits); its key size is 48 bits long.
- **null**—With null encryption, you are choosing not to provide encryption on OSPFv3 headers.

key—Type of encryption key. It can be one of the following:

- **ascii-text *key***—ASCII text key. For the **des-cbc** option, the key contains 8 ASCII characters; for **3des-cbc**, the key contains 24 ASCII characters.
- **hexadecimal *key***—Hexadecimal key. For the **des-cbc** option, the key contains 16 hexadecimal characters; for the **3des-cbc** option, the key contains 48 hexadecimal characters.

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

Related Documentation

- Understanding OSPF and OSPFv3 Authentication on SRX Series Devices on page 61

protocol (IPsec SA for OSPF)

Syntax protocol (ah | esp);

Hierarchy Level [edit security ipsec security-association *sa-name* mode transport manual direction bidirectional]

Release Information Statement introduced in Junos OS Release 12.1X46-D20.

Description Configure the IPsec protocol for a manual IPsec security association (SA) to be applied to an OSPF or OSPFv3 interface or virtual link.

Options **protocol**—Define the IPsec protocol for the manual SA. The protocol can be one of the following:

- **ah**—Authentication Header (AH) protocol.
- **esp**—Encapsulating Security Payload (ESP) protocol. This is the default.

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

Related Documentation

- Understanding OSPF and OSPFv3 Authentication on SRX Series Devices on page 61

spi (IPsec SA for OSPF)

Syntax spi *spi-value*;

Hierarchy Level [edit security ipsec security-association *sa-name* mode transport manual direction bidirectional]

Release Information Statement introduced in Junos OS Release 12.1X46-D20.

Description Configure a security parameter index (SPI) for a manual IPsec security association (SA) to be applied to an OSPF or OSPFv3 interface or virtual link.

Options **spi**—SPI for the manual SA. The SPI uniquely identifies the SA to use at the receiving host (the destination address in the packet).

Range: 256 through 16,639

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

Related Documentation

- Understanding OSPF and OSPFv3 Authentication on SRX Series Devices on page 61

show security ipsec control-plane-security-associations

Syntax
show security ipsec control-plane-security-associations
<brief | detail>
<sa-name *sa-name*>

Release Information
Command introduced in Junos OS Release 12.1X46-D20.

Description
Display information about manual IPsec security associations (SAs) applied to OSPF or OSPFv3 interfaces or virtual links.

Options
brief | detail—(Optional) Display the specified level of output.

sa-name *sa-name*—Name of the manual SA.

Required Privilege Level
view

Related Documentation
- Understanding OSPF and OSPFv3 Authentication on SRX Series Devices on page 61

List of Sample Output
show security ipsec control-plane-security-associations on page 69
show security ipsec control-plane-security-associations sa-name on page 70
show security ipsec control-plane-security-associations detail on page 70

Output Fields
Table 10 on page 69 lists the output fields for the **show security ipsec control-plane-security-associations** command. Output fields are listed in the approximate order in which they appear.

Table 10: show security ipsec control-plane-security-associations Output Fields

Field Name	Field Description
Name	Name of the SA.
Algorithm	IPsec protocol followed by encryption algorithm and authentication algorithm.
SPI	SPI value.
Total active security-associations	Total number of active manual SAs for application to OSPF or OSPFv3 interfaces or virtual links.

Sample Output

show security ipsec control-plane-security-associations

```
user@host> show security ipsec control-plane-security-associations
Name      Algorithm       SPI
test_sa   ESP:3des/md5   3e8
test_sa   ESP:3des/md5   3e8
test_sa2  ESP:3des/sha1  7d1
test_sa2  ESP:3des/sha1  7d1
Total active security-associations: 2
```

show security ipsec control-plane-security-associations sa-name

```
user@host> show security ipsec control-plane-security-associations sa-name test_sa
Name      Algorithm       SPI
test_sa   ESP:3des/md5    3e8
test_sa   ESP:3des/md5    3e8
Total active security-associations: 1
```

show security ipsec control-plane-security-associations detail

```
user@host> show security ipsec control-plane-security-associations detail
Direction: inbound, SA Name: test_sa,
Protocol: ESP:, Authentication: md5
SPI: 3e8, AUX-SPI: 0,
Mode: transport, Type: manual,
ID: 1,

Direction: outbound, SA Name: test_sa,
Protocol: ESP:, Authentication: md5
SPI: 3e8, AUX-SPI: 0,
Mode: transport, Type: manual,
ID: 2,

Direction: inbound, SA Name: test_sa2,
Protocol: ESP:, Authentication: sha1
SPI: 7d1, AUX-SPI: 0,
Mode: transport, Type: manual,
ID: 3,

Direction: outbound, SA Name: test_sa2,
Protocol: ESP:, Authentication: sha1
SPI: 7d1, AUX-SPI: 0,
Mode: transport, Type: manual,
ID: 4,
```

Example: Configuring IPsec Authentication for an OSPF Interface on an SRX Series Device

This example shows how to configure and apply a manual security association (SA) to an OSPF interface.

Requirements

Before you begin:

- Configure the device interfaces. See the *Junos OS Interfaces Library for Security Devices*.
- Configure the router identifiers for the devices in your OSPF network. See *Example: Configuring an OSPF Router Identifier*.
- Control OSPF designated router election. See *Example: Controlling OSPF Designated Router Election*

- Configure a single-area OSPF network. See *Example: Configuring a Single-Area OSPF Network*.
- Configure a multiarea OSPF network. See *Example: Configuring a Multiarea OSPF Network*.

Overview

You can use IPsec authentication for both OSPF and OSPFv3. You configure the manual SA separately and apply it to the applicable OSPF configuration. Table 11 on page 71 lists the parameters and values configured for the manual SA in this example.

Table 11: Manual SA for IPsec OSPF Interface Authentication

Parameter	Value
SA name	sa1
Mode	transport
Direction	bidirectional
Protocol	AH
SPI	256
Authentication algorithm	hmac-md5-96
Key	(ASCII) 123456789012abc
Encryption algorithm	des
Key	(ASCII) cba210987654321

Configuration

- Configuring a Manual SA on page 71
- Enabling IPsec Authentication for an OSPF Interface on page 73

Configuring a Manual SA

CLI Quick Configuration

To quickly configure a manual SA to be used for IPsec authentication on an OSPF interface, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the [**edit**] hierarchy level, and then enter **commit** from configuration mode.

```
[edit]
set security ipsec security-association sa1
set security ipsec security-association sa1 mode transport
set security ipsec security-association sa1 manual direction bidirectional
set security ipsec security-association sa1 manual direction bidirectional protocol ah
set security ipsec security-association sa1 manual direction bidirectional spi 256
```

```
set security ipsec security-association sa1 manual direction bidirectional authentication
    algorithm hmac-md5-96 key ascii-text 123456789012abc
set security ipsec security-association sa1 manual direction bidirectional encryption
    algorithm des key ascii-text cba210987654321
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure a manual SA:

1. Specify a name for the SA.

   ```
   [edit]
   user@host# edit security ipsec security-association sa1
   ```

2. Specify the mode of the manual SA.

   ```
   [edit security ipsec security-association sa1]
   user@host# set mode transport
   ```

3. Configure the direction of the manual SA.

   ```
   [edit security ipsec security-association sa1]
   user@host# set manual direction bidirectional
   ```

4. Configure the IPsec protocol to use.

   ```
   [edit security ipsec security-association sa1]
   user@host# set manual direction bidirectional protocol ah
   ```

5. Configure the value of the SPI.

   ```
   [edit security ipsec security-association sa1]
   user@host# set manual direction bidirectional spi 256
   ```

6. Configure the authentication algorithm and key.

   ```
   [edit security ipsec security-association sa1]
   user@host# set manual direction bidirectional authentication algorithm
     hmac-md5-96 key ascii-text 123456789012abc
   ```

7. Configure the encryption algorithm and key.

   ```
   [edit security ipsec security-association sa1]
   user@host# set manual direction bidirectional encryption algorithm des key ascii-text
     cba210987654321
   ```

Results

Confirm your configuration by entering the **show security ipsec** command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

NOTE: After you configure the password, you do not see the password itself. The output displays the encrypted form of the password you configured.

```
[edit]
```

```
user@host# show security ipsec
security-association sa1 {
    mode transport;
    manual {
        direction bidirectional {
            protocol ah;
            spi 256;
            authentication {
                algorithm hmac-md5-96;
                key ascii-text "$9$AP5Hp1RcylMLxSygoZUHk1REhKMVwY2oJx7jHq.zF69A0OR";
                    ## SECRET-DATA
            }
            encryption {
                algorithm des;
                key ascii-text "$9$AP5Hp1RcylMLxSygoZUHk1REhKMVwY2oJx7jHq.zF69A0OR";
                    ## SECRET-DATA
            }
        }
    }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Enabling IPsec Authentication for an OSPF Interface

CLI Quick Configuration

To quickly apply a manual SA used for IPsec authentication to an OSPF interface, copy the following command, paste it into a text file, change any details necessary to match your network configuration, copy and paste the command into the CLI at the [**edit**] hierarchy level, and then enter **commit** from configuration mode.

```
[edit]
set protocols ospf area 0.0.0.0 interface so-0/2/0 ipsec-sa sa1
```

Step-by-Step Procedure

To enable IPsec authentication for an OSPF interface:

1. Create an OSPF area.

   

   NOTE: To specify OSPFv3, include the **ospf3** statement at the **[edit protocols]** hierarchy level.

   ```
   [edit]
   user@host# edit protocols ospf area 0.0.0.0
   ```

2. Specify the interface.

   ```
   [edit protocols ospf area 0.0.0.0]
   user@host# edit interface so-0/2/0
   ```

3. Apply the IPsec manual SA.

   ```
   [edit protocols ospf area 0.0.0.0 interface so-0/2/0.0]
   user@host# set ipsec-sa sa1
   ```

Results Confirm your configuration by entering the **show ospf interface detail** command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
[edit]
user@host# show protocols ospf
area 0.0.0.0 {
   interface so-0/2/0.0 {
      ipsec-sa sa1;
   }
}
```

To confirm your OSPFv3 configuration, enter the **show protocols ospf3** command.

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly.

- Verifying the IPsec Security Association Settings on page 74
- Verifying the IPsec Security Association on the OSPF Interface on page 74

Verifying the IPsec Security Association Settings

Purpose Verify the configured IPsec security association settings. Verify the following information:

- The Security association field displays the name of the configured security association.
- The SPI field displays the value you configured.
- The Mode field displays transport mode.
- The Type field displays manual as the type of security association.

Action From operational mode, enter the **show ospf interface detail** command.

Verifying the IPsec Security Association on the OSPF Interface

Purpose Verify that the IPsec security association that you configured has been applied to the OSPF interface. Confirm that the IPsec SA name field displays the name of the configured IPsec security association.

Action From operational mode, enter the **show ospf interface detail** command for OSPF, and enter the **show ospf3 interface detail** command for OSPFv3.

Related Documentation

- Understanding OSPF and OSPFv3 Authentication on SRX Series Devices on page 61

UTM

This topic includes the following sections:

License Enforcement

For features or scaling levels that require a license, you must install and properly configure the license to meet the requirements for using the licensable feature or scale level. The router or switch enables you to commit a configuration that specifies a licensable feature or scale without a license for a 30-day grace period. The grace period is a short-term grant that enables you to start using features in the pack or scale up to the system limits (regardless of the license key limit) without a license key installed. The grace period begins when the licensable feature or scaling level is actually used by the device (not when it is first committed). In other words, you can commit licensable features or scaling limits to the device configuration, but the grace period does not begin until the device uses the licensable feature or exceeds a licensable scaling level.

NOTE: Configurations might include both licensed and nonlicensed features. For these situations, the license is enforced up to the point where the license can be clearly distinguished. For example, an authentication-order configuration is shared by both Authentication, Authorization, and Accounting (AAA), which is licensed, and by Layer 2 Tunneling Protocol (L2TP), which is not licensed. When the configuration is committed, the device does not issue any license warnings, because it is not yet known whether AAA or L2TP is using the configuration. However, at runtime, the device checks for a license when AAA authenticates clients, but does not check when L2TP authenticates clients.

The device reports any license breach as a warning log message whenever a configuration is committed that contains a feature or scale limit usage that requires a license. Following the 30-day grace period, the device periodically reports the breach to syslog messages until a license is installed and properly configured on the device to resolve the breach.

NOTE: Successful commitment of a licensable feature or scaling configuration does not imply that the required licenses are installed or not required. If a required license is not present, the system issues a warning message after it commits the configuration.

Related Documentation

- *Junos OS Feature Licenses*
- *Installation and Upgrade Guide for Security Devices*

Understanding UTM Licensing

The majority of UTM features function as a subscription service requiring a license. You can redeem this license once you have purchased your subscription license SKUs. You redeem your license by entering your authorization code and chassis serial number into the Customer Service LMS interface. Once your entitlement is generated, you can use the CLI from your device to send a license update request to the LMS server. The LMS server then sends your subscription license directly to the device.

NOTE: UTM requires 1 GB of memory. If your J2320, J2350, or J4350 device has only 512 MB of memory, you must upgrade the memory to 1 GB to run UTM.

Table 12: UTM Feature Subscription Service License Requirements

UTM Feature	Requires License
Antispam	Yes
Antivirus: full	Yes
Antivirus: express	Yes
Antivirus: sophos	Yes
Content Filtering	No
Web Filtering: integrated	Yes
Web Filtering: redirect	No
Web Filtering: local	No
Web Filtering: enhanced	Yes

Related Documentation

- *Unified Threat Management Overview*
- *Junos OS UTM Library for Security Devices*
- *Understanding UTM Custom Objects*
- *Updating UTM Licenses (CLI Procedure)*
- *Understanding WELF Logging for UTM Features*
- *Example: Configuring WELF Logging for UTM Features*

Junos OS Feature License Model Number for J Series Services Routers and SRX Series Services Gateways

For information about how to purchase a software license, contact your Juniper Networks sales representative at http://www.juniper.net/in/en/contact-us/.

Each feature license is tied to exactly one software feature, and that license is valid for exactly one device. Table 13 on page 77 describes the Junos OS features that require licenses.

Table 13: Junos OS Feature Licenses

Junos OS License Requirements										
Feature	SRX100	SRX110	SRX210	SRX220	SRX240	SRX550	SRX650	SRX1400	SRX3000 line	SRX5000 line
Access Manager	X	X	X	X	X	X	X			
BGP Route Reflectors							X			
Dynamic VPN	X	X	X	X	X	X	X			
IDP Signature Update	X *	X	X *	X *	X *	X	X	X	X	X
Application Signature Update (Application Identification)	X	X	X	X	X	X	X	X	X	X
Juniper-Kaspersky Antivirus	X	X	X	X	X	X	X			
Juniper-Sophos Antivirus	X	X	X	X	X	X	X	X	X	X
Juniper-Sophos Antispam	X	X	X	X	X	X	X	X	X	X
Juniper-Enhanced Web filtering	X	X	X	X	X	X	X	X	X	X
Juniper-Websense Web filtering	X	X	X	X	X	X	X			
Logical Systems								X	X	X
SRX100 Memory Upgrade	X									
UTM	X*	X	X *	X	X *	X	X	X	X	X

* Indicates support on high-memory devices only.

Each license allows you to run the specified advanced software features on a single device.

Related Documentation

- *Junos OS License Overview*
- *Installation and Upgrade Guide for Security Devices*
- *Installation and Upgrade Guide for Security Devices*
- *Administration Guide for Security Devices*

VPNs

This topic includes the following sections:

HMAC-SHA-256-128 Authentication

authentication-algorithm (Security IPsec)

Syntax authentication-algorithm (hmac-md5-96 | hmac-sha-256-128 | hmac-sha-256-96 | hmac-sha1-96);

Hierarchy Level [edit security ipsec proposal *proposal-name*]

Release Information Statement modified in Junos OS Release 8.5. Support for **hmac-sha-256-128** added to high-end SRX Series devices in Junos OS Release 12.1X46-D20.

Description Configure the IPsec authentication algorithm.

Options The hash algorithm to authenticate data can be one of the following:

hmac-md5-96—Produce a 128-bit digest.

hmac-sha-256-128—Produce a 256-bit digest, truncated to 128 bits. This option is not supported on group VPNs.

hmac-sha-256-96—Produce a 256-bit digest, truncated to 96 bits. This option is not supported on group VPNs or high-end SRX Series devices.

hmac-sha1-96—Produce a 160-bit digest.

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

Related Documentation

- *AutoVPN Feature Guide for SRX Series Gateway Devices*
- *Dynamic VPN Feature Guide for SRX Series Gateway Devices*
- *IPsec VPN Feature Guide for Security Devices*

authentication (Security IPsec)

Syntax

```
authentication {
  algorithm (hmac-md5-96 | hmac-sha-256-128 | hmac-sha-256-96 | hmac-sha1-96);
  key (ascii-text key | hexadecimal key );
}
```

Hierarchy Level [edit security ipsec vpn *vpn-name* manual]

Release Information Statement modified in Junos OS Release 8.5. Support for **hmac-sha-256-128** added to high-end SRX Series devices in Junos OS Release 12.1X46-D20.

Description Configure IPsec authentication parameters for a manual security association. This statement is not supported on dynamic VPN implementations.

Options **algorithm**—Hash algorithm that authenticates packet data. It can be one of the following:

- **hmac-md5-96**—Produce a 128-bit digest.
- **hmac-sha-256-128**—Produce a 256-bit digest, truncated to 128 bits.
- **hmac-sha-256-96**—Produce a 256-bit digest, truncated to 96 bits. This option is not supported on high-end SRX Series devices.
- **hmac-sha1-96**—Produce a 160-bit digest.

key—Type of authentication key. It can be one of the following:

- **ascii-text *key***—ASCII text key. For **hmac-md5-96**, the key is 16 ASCII characters; for **hmac-sha1-96**, the key is 20 ASCII characters.
- **hexadecimal *key***—Hexadecimal key. For **hmac-md5-96**, the key is 32 hexadecimal characters; for **hmac-sha1-96**, the key is 40 hexadecimal characters.

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

Related Documentation

- *Group VPN Feature Guide for Security Devices*
- *IPsec VPN Feature Guide for Security Devices*

New Features in Junos OS Release 12.1X46-D15

Junos OS Release 12.1X46-D15 introduces the following new or enhanced functionality:

- IP Monitoring on page 81
- Routing Protocols on page 85

IP Monitoring

This topic includes the following section:

Next-Hop Functionality

- show services ip-monitoring status

show services ip-monitoring status

Syntax show services ip-monitoring status

Release Information Command modified in Junos OS Release 11.4 R2. Next-hop functionality added in Junos OS Release 12.1X46-D15.

Description Display a brief summary of IP monitoring status along with the current state for a given policy.

Required Privilege Level view

Related Documentation
- *IP Monitoring Feature Guide for Security Devices*
- *show services rpm probe-results (View)*

List of Sample Output
show services ip-monitoring status on page 83
show services ip-monitoring status on page 83
show services ip-monitoring status on page 84
show services ip-monitoring status on page 84
show services ip-monitoring status on page 84

Output Fields Table 14 on page 82 lists the output fields for the **show services ip-monitoring status** command. Output fields are listed in the approximate order in which they appear.

Table 14: show services ip-monitoring status Output Fields

Field Name	Field Description
Policy	Name of the policy configured.
Probe Name	Name of the probe configured.
Address	Displays the configured target address.
Status	Displays the status of the probe on the target address. If the status is PASS, then the target address is reached.
Route-Action	Displays route injection information configured for the policy and its failover status.
Route-Instance	Displays the routing instance of the route to be injected during failover.
Route	Routing address of the route to be injected during failover.
Next-Hop	Specifies the next-hop address of the route to be injected during failover. P2P interfaces only.
State	Displays the state of the route injection action. If the state is APPLIED, then the ip-monitoring policy is in failover state.

Table 14: show services ip-monitoring status Output Fields *(continued)*

Field Name	Field Description
Interface Action	Displays the interface action type as enable or disable.
Policy Action	Displays the policy action type as enable or disable.
Admin State	Displays the current admin state of the interface.
Action Status	Displays the current action status of the interface.

Sample Output

show services ip-monitoring status

```
user@host> show services ip-monitoring status

Policy - policy1 (Non-preemptive. Status: FAIL)
  RPM Probes:
    Probe name             Test Name       Address          Status
    ---------------------- --------------- ---------------- ---------
    probe_a                a1              15.1.1.10        FAIL
    probe_a                a2              200.1.1.1        FAIL
  Route-Action:
    route-instance    route             next-hop         State
    ----------------- ----------------- ---------------- -------------
    inet.0            200.1.1.0         150.1.1.1        APPLIED
  Interface-Action:
    interface         policy action   admin state action status
    ----------------- --------------- ----------- -----------------
    fe-0/0/5.2        Enable          UP          FAILOVER
    fe-0/0/5.4        Disable         DOWN        FAILOVER
    t1-1/0/0          Enable          UP          FAILOVER
    dl0               Enable          UP          FAILOVER
    ge-0/0/1          Enable          UP          FAILOVER
```

Sample Output

show services ip-monitoring status

In this example, the policy is in the failback state, and the no-preempt option is not configured.

```
user@host> show services ip-monitoring status

Policy - policy1 (Status: PASS)
  RPM Probes:
    Probe name             Test Name       Address          Status
    ---------------------- --------------- ---------------- ---------
    probe1                 a1              99.1.1.2         PASS
  Route-Action:
    route-instance    route             next-hop         state
    ----------------- ----------------- ---------------- -------------
    inet.0            99.1.1.0          12.12.12.2       NOT-APPLIED
  Interface-Action:
```

```
   interface         policy action   admin state action status
   ----------------- --------------- ----------- -----------------
   at-2/0/0          Enable          DOWN        MARKED-DOWN
   ge-0/0/2.2        Enable          DOWN        MARKED-DOWN
   ge-0/0/2.3        Enable          DOWN        MARKED-DOWN
```

Sample Output

show services ip-monitoring status

In this example, the policy is in the failover state, and the primary is restored. The no-preempt option is configured.

```
user@host> show services ip-monitoring status

Policy - policy1 (Non-preemptive. Status: FAILOVER-NO-PREEMPT)
  RPM Probes:
     Probe name             Test Name       Address          Status
    ---------------------- --------------- ---------------- ---------
     probe1                 a1              99.1.1.2         PASS
  Route-Action:
    route-instance    route              next-hop         state
    ----------------- ------------------ ---------------- -------------
    inet.0            99.1.1.0           12.12.12.2       APPLIED
  Interface-Action:
    interface         policy action   admin state action status
    ----------------- --------------- ----------- -----------------
    at-2/0/0          Enable          UP          FAILOVER
    ge-0/0/2.2        Enable          UP          FAILOVER
    ge-0/0/2.3        Enable          UP          FAILOVER
```

Sample Output

show services ip-monitoring status

When the probe succeeds and the policy is not applied, the output is as follows:

```
user@host> show services ip-monitoring status
Policy payment (Status: PASS)
 RPM Probes:
     Probe name             Test Name       Address          Status
    ---------------------- --------------- ---------------- ---------
     Probe-Payment-Server   paysvr          9.9.9.2          PASS
 Route-Action:
    route-instance    route              next-hop         state
    ----------------- ------------------ ---------------- -------------
    inet.0            9.9.9.0/24         e1-6/0/0.0       NOT-APPLIED
```

Sample Output

show services ip-monitoring status

When the probe fails and the policy is applied, the output is as follows:

```
user@host> show services ip-monitoring status
Policy payment (Status: FAIL)
 RPM Probes:
     Probe name             Test Name       Address          Status
    ---------------------- --------------- ---------------- ---------
     Probe-Payment-Server   paysvr          9.9.9.2          FAIL
```

```
Route-Action:
   route-instance    route             next-hop         state
   ----------------- ----------------- ---------------- -------------
   inet.0            9.9.9.0/24        e1-6/0/0.0       APPLIED
```

Routing Protocols

This topic includes the following section:

OSPF Nonbroadcast Multiaccess and Point-to-Multipoint Network Support

Example: Configuring an OSPF Interface on a Nonbroadcast Multiaccess Network

This example shows how to configure an OSPFv2 interface on a nonbroadcast multiaccess (NBMA) network on high-end SRX Series devices.

Requirements

Before you begin:

- Configure the router identifiers for the devices in your OSPF network. See *Example: Configuring an OSPF Router Identifier*.
- Control OSPF designated router election. See *Example: Controlling OSPF Designated Router Election*.
- Configure a multiarea OSPF network. See *Example: Configuring a Multiarea OSPF Network*.

Overview

When you configure OSPFv2 on an NBMA network, you can use nonbroadcast mode to interoperate with other equipment. Because there is no autodiscovery mechanism, you must configure each neighbor.

Nonbroadcast mode treats the NBMA network as a partially connected LAN, electing designated and backup designated routers. All routing devices must have a direct connection to both the designated and backup designated routers; otherwise, unpredictable results can occur.

When you configure the interface, specify either the IP address or the interface name. Using both the IP address and the interface name produces an invalid configuration. For

nonbroadcast interfaces, specify the IP address of the nonbroadcast interface as the interface name.

In this example, you configure the Ethernet interface xe-2/0/0.0 as an OSPFv2 interface in OSPF area 0.0.0.1 and specify the following settings:

- **interface-type nbma**—Sets the interface to run in NBMA mode. You must explicitly configure the interface to run in NBMA mode.
- **neighbor *address* <eligible>**—Specifies the IP address of the neighboring device as 192.0.1.2. If you want the neighbor to be a designated router, include the **eligible** keyword.

NOTE: OSPF routing devices normally discover their neighbors dynamically by listening to the broadcast or multicast hello packets on the network. Because an NBMA network does not support broadcast (or multicast), the device cannot discover its neighbors dynamically, so you must configure all the neighbors statically. To configure multiple neighbors, include multiple neighbor statements.

- **hello-interval**—Specifies the length of time, in seconds, before the device sends hello packets out of the interface before it establishes adjacency with a neighbor.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the [edit] hierarchy level, and then enter **commit** from configuration mode.

```
[edit]
set interfaces xe-2/0/0 unit 0 family inet address 192.0.2.1/24
set protocols ospf area 0.0.0.1 interface xe-2/0/0.0 interface-type nbma
set protocols ospf area 0.0.0.1 interface xe-2/0/0.0 neighbor 192.0.2.2 eligible
set protocols ospf area 0.0.0.1 interface xe-2/0/0.0 hello-interval 130
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode*.

To configure an OSPFv2 interface on an NBMA network:

1. Configure the interface.

   ```
   [edit]
   user@host# set interfaces xe-2/0/0 unit 0 family inet address 192.0.2.1/24
   ```

2. Create an OSPF area.

   ```
   [edit]
   user@host# edit protocols ospf area 0.0.0.1
   ```

3. Assign the interface to the area.

In this example, include the **eligible** keyword to allow the neighbor to be a designated router.

```
[edit protocols ospf area 0.0.0.1 ]
user@host# set interface xe-2/0/0.0 interface-type nbma neighbor 192.0.2.2 eligible
```

4. Configure the hello interval.

```
[edit protocols ospf area 0.0.0.1 ]
user@host# set interface xe-2/0/0.0 hello-interval 130
```

Results

From configuration mode, confirm your configuration by entering the **show interfaces** and the **show protocols ospf** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
user@host# show interfaces
xe-2/0/0 {
  unit 0 {
    family inet {
      address 192.0.2.1/24;
    }
  }
}
```

```
user@host# show protocols ospf
area 0.0.0.1 {
  interface xe-2/0/0.0 {
    interface-type nbma;
    neighbor 192.0.2.2 eligible;
    hello-interval 130;
  }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly.

Verifying the OSPF Interface

Purpose Verify the interface configuration by confirming that the **Type** field displays NBMA.

Action From operational mode, enter the **show ospf interface detail** command.

Related Documentation

- *OSPF Configuration Overview*
- *About OSPF Interfaces*
- *OSPF Timers Overview*

Example: Configuring an OSPF Interface on a Point-to-Multipoint Network

This example shows how to configure an OSPF interface on a point-to-multipoint network high-end SRX Series devices.

Requirements

Before you begin:

- Configure the router identifiers for the devices in your OSPF network. See *Example: Configuring an OSPF Router Identifier*.
- Control OSPF designated router election. See *Example: Controlling OSPF Designated Router Election*
- Configure a multiarea OSPF network. See *Example: Configuring a Multiarea OSPF Network*.

Overview

When you configure OSPFv2 on a nonbroadcast multiaccess (NBMA) network, such as a multipoint Asynchronous Transfer Mode (ATM) or Frame Relay, OSPFv2 operates by default in point-to-multipoint mode. In this mode, OSPFv2 treats the network as a set of point-to-point links. Because there is no autodiscovery mechanism, you must configure each neighbor.

When you configure the interface, specify either the IP address or the interface name. Using both the IP address and the interface name produces an invalid configuration.

In this example, you configure the Ethernet interface xe-2/0/0.0 as an OSPFv2 interface in OSPF area 0.0.0.1 and specify 192.0.2.1 as the neighbor's IP address.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the [edit] hierarchy level, and then enter **commit** from configuration mode.

```
[edit]
set interfaces xe-2/0/0 unit 0 family inet address 192.0.2.2/24
set protocols ospf area 0.0.0.1 interface xe-2/0/0 neighbor 192.0.2.1
set protocols ospf area 0.0.0.1 interface xe-2/0/0 interface-type p2mp
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode*.

To configure an OSPFv2 interface on a point-to-multipoint network:

1. Configure the interface.

   ```
   [edit]
   user@host# set interfaces xe-2/0/0 unit 0 family inet address 192.0.2.2/24
   ```

2. Create an OSPF area.

   ```
   [edit]
   user@host# edit protocols ospf area 0.0.0.1
   ```

3. Assign the interface to the area and specify the neighbor.

   ```
   [edit protocols ospf area 0.0.0.1]
   user@host# set interface xe-2/0/0 neighbor 192.0.2.1
   ```

   To configure multiple neighbors, include a **neighbor** statement for each neighbor.

4. Specify the interface type as **p2mp**.

   ```
   [edit protocols ospf area 0.0.0.1]
   user@host# set interface xe-2/0/0 interface-type p2mp
   ```

Results

From configuration mode, confirm your configuration by entering the **show interfaces** and the **show protocols ospf** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
user@host# show interfaces
xe-2/0/0 {
  unit 0 {
    family inet {
      address 192.0.2.2/24;
    }
  }
}
```

```
user@host# show protocols ospf
area 0.0.0.1 {
  interface xe-2/0/0 {
    interface-type p2mp;
    neighbor 192.0.2.1;
  }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly.

Verifying the OSPF Interface

Purpose Verify the interface configuration by confirming that the that the **Type** field displays P2MP.

Action From operational mode, enter the **show ospf interface detail** command.

Related Documentation

- *OSPF Configuration Overview*
- *About OSPF Interfaces*

Documentation Feedback

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

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

Requesting Technical Support

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

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

Self-Help Online Tools and Resources

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

- Find CSC offerings: http://www.juniper.net/customers/support/
- Search for known bugs: http://www2.juniper.net/kb/
- Find product documentation: http://www.juniper.net/techpubs/
- Find solutions and answer questions using our Knowledge Base: http://kb.juniper.net/
- Download the latest versions of software and review release notes: http://www.juniper.net/customers/csc/software/
- Search technical bulletins for relevant hardware and software notifications: http://kb.juniper.net/InfoCenter/
- Join and participate in the Juniper Networks Community Forum: http://www.juniper.net/company/communities/
- Open a case online in the CSC Case Management tool: http://www.juniper.net/cm/

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

Opening a Case with JTAC

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

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

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

Revision History

6 March 2014—Revision 1, Junos OS Release 12.1X46-D15 Feature Guide

21 May 2014—Revision 2, Junos OS Release 12.1X46-D20 Feature Guide

10 September 2014—Revision 3, Junos OS Release 12.1X46-D25 Feature Guide

8 January 2015—Revision 4, Junos OS Release 12.1X46-D30, Feature Guide